


Universidad Complutense de Madrid

Facultad de Ciencias Matemáticas



TESIS DOCTORAL

Global aspects of bracket-generating
distributions.
Aspectos globales de distribuciones
generadoras por corchete.

Memoria para optar al grado de doctor presentada por

Francisco Javier Martínez Aguinaga

Directores:


Francisco Presas Mata y Álvaro del Pino Gómez

Universidad Complutense de Madrid

Facultad de Ciencias Matemáticas

Programa de Doctorado en Investigación Matemática

Instituto de Ciencias Matemáticas (ICMAT)



TESIS DOCTORAL

Global aspects of bracket-generating
distributions.

Aspectos globales de distribuciones
generadoras por corchete.

Memoria para optar al grado de doctor presentada por

Francisco Javier Martínez Aguinaga

Directores:

Francisco Presas Mata y Álvaro del Pino Gómez

A mis padres.



U N I V E R S I D A D
COMPLUTENSE
M A D R I D

DECLARACIÓN DE AUTORÍA Y ORIGINALIDAD DE LA TESIS PRESENTADA PARA OBTENER EL TÍTULO DE DOCTOR

D./Dña. Francisco Javier Martínez Aguinaga,
estudiante en el Programa de Doctorado Doctorado en Investigación Matemática,
de la Facultad de Ciencias Matemáticas de la Universidad Complutense de
Madrid, como autor/a de la tesis presentada para la obtención del título de Doctor y
titulada:

Global aspects of bracket-generating distributions.

Aspectos globales de distribuciones generadoras por corchete.

y dirigida por: Francisco Presas Mata y Álvaro del Pino Gómez.

DECLARO QUE:

La tesis es una obra original que no infringe los derechos de propiedad intelectual ni los derechos de propiedad industrial u otros, de acuerdo con el ordenamiento jurídico vigente, en particular, la Ley de Propiedad Intelectual (R.D. legislativo 1/1996, de 12 de abril, por el que se aprueba el texto refundido de la Ley de Propiedad Intelectual, modificado por la Ley 2/2019, de 1 de marzo, regularizando, aclarando y armonizando las disposiciones legales vigentes sobre la materia), en particular, las disposiciones referidas al derecho de cita.

Del mismo modo, asumo frente a la Universidad cualquier responsabilidad que pudiera derivarse de la autoría o falta de originalidad del contenido de la tesis presentada de conformidad con el ordenamiento jurídico vigente.

En Madrid, a 18 de diciembre de 2022

Fdo.:

Esta DECLARACIÓN DE AUTORÍA Y ORIGINALIDAD debe ser insertada en
la primera página de la tesis presentada para la obtención del título de Doctor.

Contents

Contents	VIII
Abstract	X
Resumen	XII
1 Introduction	1
1.1 Introducción	1
1.2 Introduction	3
1.3 Spaces of horizontal and transverse curves into bracket-generating distributions	5
1.4 Legendrian embedding spaces	8
1.5 Preliminaries about bracket-generating distributions	9
1.6 Convex integration with avoidance and classification of bracket-generating distributions.	12
1.7 Guide to the contents of the thesis.	16
Part I Legendrian and Formal Legendrian embedding spaces	
2 The fundamental group of the space of Formal Legendrian embeddings.	23
2.1 Spaces of embeddings of the circle into euclidean space.	23
2.2 Formal Legendrian Embeddings in \mathbb{R}^3	24
2.3 Classification of formal Legendrian embeddings in \mathbb{R}^3	26
2.4 Computation of the Fundamental group of formal Legendrian Embeddings in \mathbb{R}^3	27
2.5 Main application: formal non-triviality of Kálmán's loop.	29
3 Connected sum of parametric families of Legendrian embeddings.	37
3.1 Topological relations between $\pi_1(\mathcal{L}eg(\mathbb{R}^3, \xi_{\text{std}}))$ and $\pi_1(\mathcal{L}eg(\mathbb{S}^3, \xi_{\text{std}}))$	37
3.2 Connected sum at the π_0 -level.	40
3.3 Connected sum at the π_1 -level.	41
3.4 Invariants for one-parameter families of Legendrians.	45

3.5	Legendrian (p, q) -tangles and associated words.	49
3.6	New examples of loops with non-trivial monodromy.	52

Part II Horizontal and transverse embedding spaces

4	h–Principle for horizontal and transverse curves.	61
4.1	Introduction	61
4.2	ε -horizontal and ε -transversality	62
4.3	Graphical models	68
4.4	Microflexibility of curves	72
4.5	Tangles	81
4.6	Controllers	96
4.7	h -Principles for horizontal embeddings	101
4.8	h -Principle for transverse embeddings	110
4.9	Appendix: Technical lemmas on commutators	114

Part III Spaces of bracket-generating distributions

5	Convex integration with avoidance	121
5.1	Jets and relations	121
5.2	Ampleness and convex integration	126
5.3	Avoidance	133
5.4	Proof of the main Theorem	139
5.5	An example: exact, linearly-independent differential forms	149
5.6	Relations involving functions	155
6	h-Principle for step-2 distributions	157
6.1	h -Principle for step-2 distributions	157
7	h–Principle for hyperbolic $(4, 6)$-distributions.	163
7.1	Maximal non-involutivity	163
7.2	h -Principle for hyperbolic $(4, 6)$ -distributions	165

Part IV Future work

8	Future work	177
8.1	Transverse submanifolds.	177
8.2	Spaces of distributions.	177
	References	185

Abstract

The central topic of this PhD thesis is the study of global properties of bracket-generating distributions. In the first two parts we focus our attention on the study of tangent and transverse embeddings to these type of distributions. We have classified regular embedded and transverse curves into any manifold of dimension greater than 3 (i.e. we have proved that there exists a complete h -principle). This contrasts with the 3-dimensional contact case, where it is well known that an h -principle for legendrian/transverse embedded curves does not hold.

We have also studied parametric families of Legendrian embedding in the case of contact manifolds of dimension 3. We have computed the fundamental group of the space of formal Legendrian embeddings and as a consequence we have shown that previous examples of non-trivial loops of Legendrian embeddings in the literature were already non trivial at the formal level. Continuing with the study of loops of Legendrian embeddings, we have defined a connected sum for 1-parametric families of legendrian embeddings. As the main application of this construction, we have found infinitely many new examples of non-trivial loops of Legendrian embeddings with non-trivial monodromy.

In the third part of this thesis we have studied the classification, up to homotopy, of tangent distributions satisfying various non-involutivity conditions. On one hand, we have proved that the full h -principle holds for step-2 bracket-generating distributions. This result follows from an application of the method of convex integration developed by M. Gromov. The classification of $(3, 5)$ and $(3, 6)$ distributions follows as a particular case.

The main result of the third part of this thesis is the development of a new h -principle technique called convex integration up to avoidance. This technique refines the classical method of convex integration by implementing an “avoidance trick”.

The goal of this trick is to avoid some principal subspaces where the differential relation fails to be ample. Given any differential relation, this method produces an associated object called an “avoidance template”. If the process is successful, we say that the relation is “ample up to avoidance” and we prove that convex integration applies.

Using this technique we have found the first example of a differential relation that is ample in coordinate directions but not in all directions, answering a question of Eliashberg and Mishachev. Our main application is the proof, by using this method, of a complete h -principle for hyperbolic $(4, 6)$ distributions. This example shows that this new technique is capable of addressing differential relations beyond the applicability of classical convex integration.

Resumen

El tema central e hilo conductor de esta tesis doctoral es el estudio de propiedades globales de distribuciones generadoras por corchete. En las dos primeras partes hemos centrado nuestra atención en el estudio de encajes regulares tangentes y transversos a este tipo de distribuciones. Hemos clasificado las curvas embebidas horizontales y transversas en cualquier variedad de dimensión mayor que 3 (i.e. hemos probado que existe un h -principio completo). Esto contrasta con el caso 3-dimensional de contacto, donde es bien sabido que no existe h -principio para curvas legendrianas/transversas.

Paralelamente, hemos estudiado familias paramétricas de nudos legendrianos en el caso de variedades de contacto de dimensión 3. Hemos computado el grupo fundamental del espacio de nudos legendrianos formales y como consecuencia hemos demostrado que ejemplos previos en la literatura de lazos de encajes Legendrianos no triviales eran ya no triviales a nivel formal. Continuando con el estudio de lazos de encajes legendrianos, hemos definido una suma conexa de familias 1-paramétricas de encajes legendrianos. Como principal aplicación de esta construcción hemos encontrado infinitos nuevos ejemplos de lazos de encajes legendrianos no triviales con monodromía no trivial.

En la tercera parte de esta tesis hemos estudiado el problema de clasificación, módulo homotopía, de distribuciones tangentes satisfaciendo varias condiciones de no-involutividad. Por un lado, hemos demostrado que existe un h -principio completo para distribuciones generadoras por corchete de paso 2. Este resultado se sigue de una aplicación del método de integración convexa desarrollado por M. Gromov. La clasificación de distribuciones $(3, 5)$ y $(3, 6)$ se sigue como caso particular.

El principal resultado de la tercera parte de esta tesis es el desarrollo de una técnica de h -principio novedosa llamada integración convexa módulo evitación. Esta técnica refina el método de integración convexa clásico implementando un “truco de evitación”.

El objetivo de este truco es evitar algunos subespacios principales donde la relación diferencial estudiada no satisface amplitud. Dada una relación diferencial, este método produce un objeto asociado llamado “plantilla de evitación”. Si el proceso es exitoso, decimos que la relación es “amplia módulo evitación” y demostramos que el método de integración convexa aplica.

Usando esta nueva técnica hemos encontrado el primer ejemplo de una relación diferencial que es amplia en direcciones coordinadas pero no es amplia en todas las direcciones, dando respuesta a una pregunta planteada por Eliashberg y Mishachev. Nuestra principal aplicación es la prueba, usando este método, de que existe un h -principio completo para distribuciones $(4, 6)$ hiperbólicas. Este ejemplo pone de manifiesto que esta nueva técnica es capaz de abordar relaciones diferenciales más allá del campo de aplicación del método de integración convexa clásico.

Chapter 1

Introduction

1.1 Introducción

El tema central de esta tesis doctoral es el estudio de propiedades globales de distribuciones generadoras por corchete. Está dividida en tres partes bien diferenciadas que conjuntamente abordan el estudio de propiedades globales de estos objetos. Las dos primeras partes están dedicadas al estudio de la topología global de espacios de encajes tangentes mientras que la tercera parte trata acerca de la topología global de espacios de distribuciones generadoras por corchete.

Una *q-distribución* en una variedad diferenciable M es una sección diferenciable \mathcal{D} al fibrado de Grassmann de q -planos. Existe una motivación desde el punto de vista de teoría de control para considerar estos objetos: podemos pensar en M como el espacio de configuraciones y en \mathcal{D} como las *direcciones admisibles de movimiento*. Centramos nuestra atención en distribuciones generadoras por corchete, lo que significa que cualquier vector en TM puede expresarse como combinación lineal de corchetes de Lie involucrando secciones de \mathcal{D} .

Las estructuras de contacto son ejemplos prototípicos de distribuciones generadoras por corchete y han sido ampliamente estudiadas. La Topología de Contacto es un área muy activa en Matemáticas que tiene intersección no vacía con muchas ramas como el estudio de sistemas dinámicos, geometría algebraica, análisis, topología diferencial o topología simpléctica. No obstante, no se sabe tanto acerca de la topología global de distribuciones generadoras por corchete más generales.

Una pregunta natural es si dos puntos cualesquiera en M pueden ser conectados por un *camino horizontal*, i.e. un camino cuyos vectores velocidad tomen valores en \mathcal{D} . Una condición suficiente está dada por un teorema clásico de Chow [27]: dos puntos cualesquiera en M pueden ser conectados si \mathcal{D} es *generadora por corchete*. Esto es, el teorema de Chow es un enunciado *infinitesimal* a *global*.

Aunque las pruebas clásicas del teorema de Chow producen caminos horizontales que son diferenciables a trozos, también se pueden construir caminos C^∞ mediante suavizados apropiados, ver [60, Subsection 1.2.B]. Se sigue que cualquier clase de homotopía de lazos en M puede ser representada mediante un lazo horizontal diferenciable. Esto es, la inclusión

$$\iota : \mathcal{L}(M, \mathcal{D}) \longrightarrow \mathcal{L}(M)$$

es una π_0 -sobreyección. Aquí $\mathcal{L}(M)$ denota el espacio de lazos (no basados) de M (equipado con la topología C^∞) y $\mathcal{L}(M, \mathcal{D})$ es el subespacio de lazos horizontales. Más recientemente, Z. Ge [51] ha probado que la inclusión análoga para lazos H^1 es una equivalencia débil de homotopía; ver también [17].

En esta tesis consideramos una variación de este tema, probando resultados de clasificación para espacios de encajes horizontales. Nuestros teoremas relacionan estos espacios con sus homólogos *formales* (grosso modo, espacios de encajes diferenciables más datos homotópicos adicionales). Ser

cuidadosos con la condición de encaje es delicado (como suele ocurrir con este tipo de h -principios) y gran parte de la Parte II de esta tesis está dedicada a abordar esta cuestión. También probamos teoremas de clasificación análogos para inmersiones horizontales mediante pruebas más simples. Asimismo, deducimos que la aplicación ι mencionada anteriormente es una equivalencia débil de homotopía (i.e. el análogo smooth del teorema de Ge). Finalmente, nuestras técnicas generalizan al caso de curvas transversas inmersas/embebidas, dando lugar a resultados de clasificación análogos.

Así pues, mostramos que estos objetos tienen una naturaleza flexible (satisfacen un h -principio completo) excepto en el caso de encajes legendrianos y transversos en variedades 3-dimensionales. A continuación tratamos el caso de estas familias de curvas particulares en la Parte II de esta tesis.

Estudiamos los espacio homólogos formales a los espacios de encajes legendrianos en el espacio 3-dimensional de contacto standard. Computamos el grupo fundamental y mostramos que ejemplos previos de la literatura de elementos no triviales del grupo fundamental son ya no triviales a nivel formal. Esto aporta una perspectiva algebro-geométrica al estudio de estos espacios que produce invariantes algebraicos no triviales de familias 1-paramétricas de encajes legendrianos. También introducimos una noción de suma conexa de familias 1-paramétricas de legendrianas que da lugar a la construcción de infinitos nuevos ejemplos de lazos no triviales de legendrianas.

En la tercera parte de esta tesis estudiamos la clasificación, módulo homotopía, de distribuciones tangentes satisfaciendo varias condiciones de no-involutividad. Por una parte, hemos probado que existe h -principio para distribuciones generadoras por corchete de paso 2. Este resultado se sigue de una aplicación del método de integración convexa desarrollado por M. Gromov. La clasificación de distribuciones $(3, 5)$ y $(3, 6)$ se sigue como caso particular.

La principal contribución de la tercera parte de esta tesis es el desarrollo de una técnica de h -principio novedosa llamada integración convexa módulo evitación. Esta técnica refina el método de integración convexa clásico implementando un “truco de evitación”.

Usando esta nueva técnica hemos encontrado el primer ejemplo de una relación diferencial que es amplia en direcciones coordinadas pero no es amplia en todas las direcciones, dando respuesta a una pregunta planteada por Eliashberg y Mishachev. Nuestra principal aplicación es la prueba, usando este método, de que existe un h -principio completo para distribuciones $(4, 6)$ hiperbólicas. Este ejemplo pone de manifiesto que esta nueva técnica es capaz de abordar relaciones diferenciales más allá del campo de aplicación del método de integración convexa clásico.

Trabajamos bajo el siguiente supuesto:

Supuesto 1.1.1 *Todas las distribuciones generadoras por corchete que consideramos en esta tesis son de vector de crecimiento constante (i.e. el vector de crecimiento no depende del punto). Ver Subsección 1.5.0.2.*

1.1.1 Lista de artículos

Parte de los resultados de esta Tesis Doctoral han sido recogidos en la siguiente lista de artículos:

- [1] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Fundamental groups of formal legendrian and horizontal embedding spaces*. Algebraic & Geometric Topology 20 (2020), 3219–3312.
- [2] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Parametric connected sums in the space of legendrian embeddings*. En preparación.
- [3] J. Martínez-Aguinaga. *Existence and classification of maximal growth distributions*. En preparación.
- [4] J. Martínez-Aguinaga, A. del Pino. *Convex integration with avoidance and hyperbolic $(4,6)$ distributions*. arXiv:2112.14632.
- [5] J. Martínez-Aguinaga, A. del Pino. *Classification of tangent and transverse knots in bracket-generating distributions*. arXiv:2210.00582

Durante el desarrollo de esta Tesis Doctoral también se han elaborado los siguientes artículos:

- [6] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Loops of Legendrians in contact 3-manifolds*. Classical and Quantum Physics. 60 Years Alberto Ibort Fest Geometry, Dynamics and Control. Springer Proceedings in Physics 229, pp 361–372. (2019).
- [7] E. Fernández, J. Martínez-Aguinaga, F. Presas. *The homotopy type of the contactomorphism groups of tight contact 3-manifolds, part I*. arXiv:2012.14948.

1.2 Introduction

The central topic of this PhD thesis is the study of global properties of bracket-generating distributions. It is divided in three well differentiated parts that coalesce into the global study of such objects. The first two parts are devoted to the study of the global topology of spaces of tangent embedding while the third one deals with the study of the global topology of spaces of bracket-generating distributions.

A q -distribution on a smooth manifold M is a smooth section \mathcal{D} of the Grassmann bundle of q -planes. There is a control-theoretic motivation for considering such objects: we may think of M as configuration space and of \mathcal{D} as the *admissible directions of motion*. We focus our attention in bracket-generating distributions, which means that any vector in TM can be written as a linear combination of Lie brackets involving sections of \mathcal{D} .

Contact structures are prototypical examples of bracket-generating distributions and have been widely studied. Contact Topology is a very active area in Mathematics that has non-empty intersection with many other branches such as the study of dynamical systems, algebraic geometry, analysis, differential topology or symplectic topology. Nonetheless, not so much is known about the global topology of more general bracket-generating distributions.

A natural question is whether any two points in M can be connected by a *horizontal path*, i.e. a path whose velocity vectors take values in \mathcal{D} . A sufficient condition is given by a classic theorem of Chow [27]: any two points in M can be connected if \mathcal{D} is *bracket-generating*. Being bracket-generating means that any vector in TM can be written as a linear combination of Lie brackets involving sections of \mathcal{D} . That is, Chow’s theorem is an *infinitesimal* to *global* statement.

Even though classic proofs of Chow’s theorem produce horizontal paths that are piecewise smooth, C^∞ -paths can be constructed by suitable smoothing, see [60, Subsection 1.2.B]. It follows that every homotopy class of loops on M can be represented by a smooth horizontal loop. That is, the inclusion

$$\iota : \mathcal{L}(M, \mathcal{D}) \longrightarrow \mathcal{L}(M)$$

is a π_0 -surjection. Here $\mathcal{L}(M)$ is the (unbased) loop space of M (endowed with the C^∞ -topology) and $\mathcal{L}(M, \mathcal{D})$ is the subspace of horizontal loops. More recently, Z. Ge [51] proved that the analogous inclusion for H^1 -loops is a weak homotopy equivalence; see also [17].

In this thesis we consider a variation on this theme, proving classification statements for spaces of horizontal embeddings. Our theorems relate these spaces to their *formal* counterparts (roughly speaking, spaces of smooth embeddings plus some additional homotopical data). Taking care of the embedding condition is rather delicate (as is often the case for h -principles of this type) and much of Part II of this thesis is dedicated to handling it. Analogous classification statements hold for horizontal immersions, with simpler proofs. We also deduce that the map ι above is a weak homotopy equivalence (i.e. the smooth analogue of Ge’s theorem). Lastly, our techniques translate to the setting of embedded/immersed transverse curves, yielding similar classification results.

Thus, we show that these objects have a flexible nature (they abide by a complete h -principle) except in the case of legendrian and transverse embeddings in contact 3-dimensional manifolds. We then move to these particular families of curves in Part II of this thesis.

We study the formal counterpart of spaces of legendrian embeddings in the standard contact 3-dimensional space. We compute the fundamental group and we show that previous examples in the literature of non-trivial elements in the fundamental group are already non-trivial at the formal level. This provides an algebraic-geometrical perspective to the study of these spaces that produces non-trivial algebraic invariants of 1-parametric families of Legendrian embeddings. We also introduce a notion of connected sum of 1-parametric families of Legendrians that leads to the construction of infinitely many non-trivial examples of loops of Legendrians.

In the third part of this thesis we study the classification, up to homotopy, of tangent distributions satisfying various non-involutivity conditions. On one hand, we have proved that the full h -principle holds for step-2 bracket-generating distributions. This result follows from an application of the method of convex integration developed by M. Gromov. The classification of $(3, 5)$ and $(3, 6)$ distributions follows as a particular case.

The main contribution of the third part of this thesis is the development of a new h -principle technique called convex integration up to avoidance. This technique refines the classic method of convex integration by implementing an “avoidance trick”.

The goal of this trick is to avoid some principal subspaces where the differential relation fails to be ample. Given any differential relation, this method produces an associated object called an “avoidance template”. If the process is successful, we say that the relation is “ample up to avoidance” and we prove that convex integration applies.

Using this technique we have found the first example of a differential relation that is ample in coordinate directions but not in all directions, answering a question of Eliashberg and Mishachev. Our main application is the proof, by using this method, of a complete h -principle holds for hyperbolic $(4, 6)$ distributions. This example shows that this new technique is capable of addressing differential relations beyond the applicability of classic convex integration.

We now state our theorems. We work under the following assumption:

Assumption 1.2.1 *All the bracket-generating distributions we consider in this thesis are of constant growth (i.e. the growth vector does not depend on the point). See Subsection 1.5.0.2.*

1.2.1 List of articles

Part of the results of this PhD thesis have been covered in the following list of articles:

- [1] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Fundamental groups of formal legendrian and horizontal embedding spaces*. Algebraic & Geometric Topology 20 (2020), 3219–3312.
- [2] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Parametric connected sums in the space of legendrian embeddings*. In preparation.
- [3] J. Martínez-Aguinaga. *Existence and classification of maximal growth distributions*. In preparation.
- [4] J. Martínez-Aguinaga, A. del Pino. *Convex integration with avoidance and hyperbolic $(4,6)$ distributions*. arXiv:2112.14632.
- [5] J. Martínez-Aguinaga, A. del Pino. *Classification of tangent and transverse knots in bracket-generating distributions*. arXiv:2210.00582

During the development of this PhD Thesis the following articles have also been developed:

- [6] E. Fernández, J. Martínez-Aguinaga, F. Presas. *Loops of Legendrians in contact 3-manifolds*. Classical and Quantum Physics. 60 Years Alberto Ibort Fest Geometry, Dynamics and Control. Springer Proceedings in Physics 229, pp 361–372. (2019).
- [7] E. Fernández, J. Martínez-Aguinaga, F. Presas. *The homotopy type of the contactomorphism groups of tight contact 3-manifolds, part I*. arXiv:2012.14948.

1.3 Spaces of horizontal and transverse curves into bracket-generating distributions

1.3.1 Immersed horizontal curves

Let us write $\mathfrak{Imm}(M, \mathcal{D}) \subset \mathcal{L}(M, \mathcal{D})$ for the subspace of immersed horizontal loops. In order to study it, we introduce the so-called *scanning map*:

$$\mathfrak{Imm}(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}^f(M, \mathcal{D}),$$

taking values in the space of formal horizontal immersions

$$\mathfrak{Imm}^f(M, \mathcal{D}) := \{(\gamma, F) \mid \gamma \in \mathcal{L}(M), \quad F \in \text{Mon}(TS^1, \gamma^*\mathcal{D})\}.$$

The question to be addressed is whether the scanning map is a weak homotopy equivalence. The answer is positive if \mathcal{D} is a contact structure [39, Section 14.1] but, for other distributions, the answer may be negative due to the presence of *rigid curves* [14].

A horizontal curve is rigid if possesses no C^∞ -deformations relative to its endpoints, up to reparametrisation. These curves are isolated and conform exceptional components within the space of all horizontal maps with given boundary conditions. *Rigid loops* also exist. Because of this, the inclusion $\mathfrak{Imm}(M, \mathcal{D}) \rightarrow \mathfrak{Imm}^f(M, \mathcal{D})$ can fail to be bijective at the level of connected components; see [90, Remark 23]. Being rigid is the most extreme case of being *singular*. This means that the *endpoint map* of the curve is not submersive, so the curve has fewer deformations than expected; see Subsection 4.1.1. This phenomenon does not happen in the Transverse case (see 1.3.4).

The subspaces of rigid and singular curves have a *geometric* and not a *topological* nature. By this, we mean that small perturbations of \mathcal{D} can radically change their homotopy type; see [81] or [90, Theorem 27]. This motivates us to discard singular curves and focus on $\mathfrak{Imm}^r(M, \mathcal{D})$, the subspace of regular horizontal immersions. In doing so, the subspace that we discard is not too large: germs of singular horizontal curves were shown to form a subset of infinite codimension among all horizontal germs, first in the analytic case [89] and then in general [16]. Earlier, it had already been observed [68, Corollary 7] that *regular* (i.e. non-singular) germs are C^∞ -generic.

Our first result reads:

Theorem 1.1. *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathfrak{Imm}^r(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}^f(M, \mathcal{D}).$$

Apart from the aforementioned contact case, in which there are no singular curves, this was already known in the Engel case [90].

Corollary 1.3.1 *Let \mathcal{D}_0 and \mathcal{D}_1 be bracket-generating distributions on a manifold M , homotopic as subbundles of TM . Then, the spaces $\mathfrak{Imm}^r(M, \mathcal{D}_0)$ and $\mathfrak{Imm}^r(M, \mathcal{D}_1)$ are weakly homotopy equivalent.*

This follows immediately from Theorem 1.1 and the analogous fact about $\mathfrak{Imm}^f(M, \mathcal{D}_0)$ and $\mathfrak{Imm}^f(M, \mathcal{D}_1)$. It follows that all the data about \mathcal{D} encoded in $\mathfrak{Imm}^r(M, \mathcal{D})$ is purely formal.

1.3.2 Embedded horizontal curves

We now consider the subspace of embedded horizontal loops $\mathfrak{Emb}(M, \mathcal{D}) \subset \mathfrak{Imm}(M, \mathcal{D})$, together with its scanning map

$$\mathfrak{Emb}(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}^f(M, \mathcal{D}),$$

into the space of *formal horizontal embeddings*:

$$\begin{aligned} \mathfrak{Emb}^f(M, \mathcal{D}) := \{ (\gamma, (F_s)_{s \in [0,1]}) : \quad & \gamma \in \mathfrak{Emb}(M), \quad F_s \in \text{Mon}_{\mathbb{S}^1}(T\mathbb{S}^1, \gamma^*TM), \\ & F_0 = \gamma', \quad F_1 \in \gamma^*\mathcal{D} \}, \end{aligned}$$

i.e. the homotopy pullback of $\mathfrak{Emb}(M)$ and $\mathfrak{Imm}^f(M, \mathcal{D})$ mapping into $\mathfrak{Imm}^f(M)$. Reasoning as above leads us to introduce $\mathfrak{Emb}^r(M, \mathcal{D})$, the subspace of regular horizontal embeddings. Our second (and main) result reads:

Theorem 1.2. *Let (M, \mathcal{D}) be a bracket-generating distribution with $\dim(M) \geq 4$. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathfrak{Emb}^r(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}^f(M\mathcal{D}).$$

Note that the dimensional assumption is sharp, since the result is known to be false in 3-dimensional Contact Topology [11].

Theorem 1.2 was already known in the Engel case [24] and in the higher-dimensional contact setting [39, p. 128]. Our arguments differ considerably from both. The proof in [39] is contact-theoretical in nature, relying on isocontact immersions. The one in [24] uses the so-called *Geiges projection*, which is particular to the Engel case. The methods in the present thesis use instead

local charts in which the distribution can be understood as a connection; see Subsection 4.3. This is reminiscent of the Lagrangian projection in Contact Topology and closely related to methods used in the Geometric Control Theory [60, 81] (with the added difficulty of tracking the embedding condition).

Much like earlier:

Corollary 1.3.2 *Fix a manifold M with $\dim(M) \geq 4$. Let \mathcal{D}_0 and \mathcal{D}_1 be bracket-generating distributions on M , homotopic as subbundles of TM . Then, the spaces $\mathfrak{Emb}^r(M, \mathcal{D}_0)$ and $\mathfrak{Emb}^r(M, \mathcal{D}_1)$ are weakly homotopy equivalent.*

1.3.3 Horizontal loops

Now we go back to the problem we started with:

Theorem 1.3. *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathcal{L}(M, \mathcal{D}) \longrightarrow \mathcal{L}(M).$$

This also holds for the (based) loop space $\Omega_p(M)$ and its subspace of horizontal loops $\Omega_p(M, \mathcal{D})$, for all $p \in M$. Observe that the statement uses no regularity assumptions. The reason is that singularity issues can be bypassed thanks to what we call the *stopping-trick* (namely, one can slow the parametrisation of a horizontal curve down to zero locally in order to guarantee that enough compactly-supported variations exist). See Subsection 4.7.5.

1.3.4 Immersed transverse curves

The other geometrically interesting notion for curves in bracket-generating distributions is that of transversality. We define $\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D})$ to be the space of immersed loops that are everywhere transverse to \mathcal{D} . Like in the horizontal setting, one can introduce formal transverse immersions

$$\mathfrak{Imm}_{\mathcal{T}}^f(M, \mathcal{D}) = \{(\gamma, F) : \gamma \in \mathcal{L}(M), \quad F \in \text{Mon}_{\mathbb{S}^1}(TS^1, \gamma^*(TM/SD))\},$$

and see that there is a scanning map

$$\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}_{\mathcal{T}}^f(M, \mathcal{D}).$$

Being transverse is an open condition and therefore rigidity/singularity is not a phenomenon we encounter. We prove:

Theorem 1.4. *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then the inclusion*

$$\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}_{\mathcal{T}}^f(M, \mathcal{D})$$

is a weak homotopy equivalence.

This result is not new. The h -principle for smooth immersions (of any dimension!) transverse to analytic bracket-generating distributions was proven in [89]. The analyticity assumption was later dropped by A. Bhowmick in [16], using Nash-Moser methods. Both articles rely on an argument

due to Gromov relating the flexibility of transverse maps to the microflexibility of (micro)regular horizontal curves. The approach in this thesis is independent.

Once again, a corollary is that the weak homotopy type of $\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D})$ depends on \mathcal{D} only formally.

1.3.5 Embedded transverse curves

Lastly, we address embedded transverse loops $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})$ and their scanning map into the analogous formal space:

$$\mathfrak{Emb}_{\mathcal{T}}^f(M, \mathcal{D}) = \left\{ (\gamma, (F_s)_{s \in [0,1]}) : \begin{array}{l} \gamma \in \mathfrak{Emb}(M), \quad F_s \in \text{Mon}_{\mathbb{S}^1}(T\mathbb{S}^1, \gamma^*TM), \\ F_0 = \gamma', \quad F_1 : T\mathbb{S}^1 \rightarrow \gamma^*TM \rightarrow \gamma^*(TM/\mathcal{D}) \text{ is injective} \end{array} \right\}.$$

Our fourth result reads:

Theorem 1.5. *Let (M, \mathcal{D}) be a bracket-generating distribution with $\dim(M) \geq 4$. Then the inclusion*

$$\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}_{\mathcal{T}}^f(M, \mathcal{D})$$

is a weak homotopy equivalence. In particular, $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})$ depends only on the formal class of \mathcal{D} .

The dimension condition is sharp, since transverse embeddings into 3-dimensional contact manifolds do not satisfy a complete h -principle. Indeed, there are examples of transverse knots that have the same formal invariants but are not transversely isotopic [15]. Furthermore, Theorem 1.5 is only interesting in corank 1. Indeed, it is a classical result [39, 4.6.2] that closed n -dimensional submanifolds transverse to corank k distributions abide by all forms of the h -principle if $k > n$.

1.4 Legendrian embedding spaces

The computation of the homotopy type of the space of Legendrian embeddings into a contact 3-fold has a long story. For a while, it was thought that the computation could be made at the formal level. We mean by that that the inclusion of the space of Legendrian embeddings into the space of formal Legendrian embeddings, i.e. the space of pairs made of a smooth embedding and a formal Legendrian derivative, was a weak homotopy equivalence.

This was proven to be wrong in the key article of D. Bennequin [11]; in which it was shown that the formal space associated to the standard contact \mathbb{R}^3 possesses some connected components that are not representable by Legendrian knots. In other words, the restriction of the induced map of the inclusion at π_0 -level was not surjective. This was the first hint of rigidity phenomena in contact topology.

Later on, there has been an industry checking how far the inclusion map is from being injective or surjective at π_0 -level, see, eg, the work of Chekanov [26], Ding and Geiges [32], Eliashberg and Fraser [37], Etnyre and Honda [43] or Osváth, Szabó and Thurston [98].

The next step was the study of higher homotopy groups. This was developed by Kálmán [71] in dimension 3 using pseudoholomorphic curves invariants and by Sabloff and Sullivan [93] in dimension $2n + 1$, $n > 1$, using generating function invariants. However, they just checked that several non trivial loops in the space of Legendrian embeddings were trivial as elements in the

fundamental group of the space of smooth embeddings. We show that all Kálmán's examples are non trivial in the space of formal Legendrian embeddings, see Section 2.5. This makes unnecessary the use of sophisticated invariants to compute these examples. In order to do that, we compute the fundamental group of the space of formal Legendrian embeddings. This is the content of Section 2.2.

1.5 Preliminaries about bracket-generating distributions

The following definition generalises the notion of distribution:

Definition 1.5.1 *Let M be a smooth manifold. A **differential system** \mathcal{D} is a C^∞ -submodule of the space of smooth vector fields.*

Given a smooth distribution on M , we can construct a differential system by taking its smooth sections. Conversely, a differential system \mathcal{D} gives a distribution if the dimension of its pointwise span $\mathcal{D}(p) \subset T_p M$ is independent of $p \in M$. Nonetheless, note that this is not a one to one correspondence. In this manner, we think of differential systems as singular distributions; we will often abuse notation and use \mathcal{D} to denote both the distribution and its sections.

Remark 1.5.2 *When M is not compact, it is convenient to impose that \mathcal{D} satisfies the sheaf condition. The reason is that there may be differential systems that only differ from one another due to their behaviour at infinity; imposing the sheaf condition removes this redundancy. These subtleties will not be relevant for us.*

1.5.0.1 Lie flag

Let us introduce some terminology. We say that the string a , depending on the variable a , is a **formal bracket expression** of length 1. Similarly, we say that the string $[a_1, a_2]$, depending on the variables a_1 and a_2 , is a formal bracket expression of length 2. Inductively, we define a formal bracket expression of length n to be a string of the form $[A(a_1, \dots, a_j), B(a_{j+1}, \dots, a_n)]$ with $0 < j < n$ and A and B formal bracket expressions of lengths j and $n - j$, respectively.

Given a differential system \mathcal{D} , we define its **Lie flag** as the sequence of differential systems

$$\mathcal{D}_1 \subset \mathcal{D}_2 \subset \mathcal{D}_3 \subset \dots$$

in which \mathcal{D}_i is the C^∞ -span of vector fields of the form $A(v_1, \dots, v_j)$, $j \leq i$, where the v_k are vector fields in \mathcal{D} and A is a formal bracket expression of length j . As such, $\mathcal{D}_1 = \mathcal{D}$.

1.5.0.2 Growth vector

Given a point $p \in M$, one can use the Lie flag to produce a flag of vector spaces:

$$\mathcal{D}_1(p) \subset \mathcal{D}_2(p) \subset \mathcal{D}_3(p) \subset \dots$$

Here $\mathcal{D}_i(p)$ denotes the span of \mathcal{D}_i at p . This yields a non-decreasing sequence of integers

$$(\dim(\mathcal{D}_1(p)), \dim(\mathcal{D}_2(p)), \dim(\mathcal{D}_3(p)), \dots)$$

which in general depends on p . This sequence is called the **growth vector** of \mathcal{D} at p .

If the growth vector does not depend on the point, we will say that the differential system \mathcal{D} is of **constant growth**. If this is the case, all the differential systems in the Lie flag arise as spaces of sections of distributions. Some examples of distributions of constant growth are (regular) foliations, contact structures, and Engel structures.

The following notion is central to us:

Definition 1.5.3 *A differential system (M, \mathcal{D}) is **bracket-generating** if, for every $p \in M$ and every $v \in T_p M$, there is an integer m such that $v \in \mathcal{D}_m(p)$. This integer is called the **step**.*

This definition a priori depends on the choice of point. Nonetheless, it will not for our purposes since we will work under the following assumption.

Assumption 1.5.4 *As stated in Assumption 1.2.1: we will henceforth assume that the differential system \mathcal{D} we start with is a distribution of constant growth.*

1.5.0.3 The nilpotentisation and the curvature

Note that the Lie Flag defined earlier can be understood as follows. Since the Lie bracket is a (naturally defined) first order operator acting on vector fields, we can apply it to the sections $\Gamma(\xi)$ of ξ . This defines for us the sequence of modules that we call *Lie flag*:

$$\begin{aligned} \Gamma^1(\xi) &\subset \Gamma^2(\xi) \subset \Gamma^3(\xi) \subset \dots \\ \Gamma^1(\xi) &:= \Gamma(\xi), \quad \Gamma^{i+1}(\xi) := [\Gamma^1(\xi), \Gamma^i(\xi)]. \end{aligned}$$

For simplicity, we will always assume that ξ is *regular*, i.e. there is a distribution ξ_i such that $\Gamma^i(\xi) = \Gamma(\xi_i)$. The rank of ξ_i is then a measurement of the non-involutivity of ξ .

We then have

$$\xi_1 = \xi \subset \xi_2 \subset \xi_3 \subset \dots$$

one notes that it stabilises: i.e. there exists some smallest i_0 such that $\xi_i = \xi_{i_0}$ for all $i \geq i_0$. This means that $\Gamma^{i_0}(\xi)$ is involutive and thus ξ_{i_0} is the tangent bundle of a foliation \mathcal{F} on M .

We define the **nilpotentisation** $\mathcal{L}(\xi)$ of ξ as the graded vector bundle

$$\xi_1 \oplus \xi_2/\xi_1 \oplus \dots \oplus \xi_i/\xi_{i-1} \oplus \dots \oplus \xi_{i_0}/\xi_{i_0-1}.$$

One can then observe that the composition

$$\Gamma^j(\xi) \times \Gamma^i(\xi) \longrightarrow \Gamma^{i+j}(\xi) \longrightarrow \Gamma^{i+j}(\xi)/\Gamma^{i+j-1}(\xi)$$

of the Lie bracket with the projection is C^∞ -linear. In particular, it descends to a bilinear map

$$\Omega_{i,j}(\xi) : \xi_j/\xi_{j-1} \times \xi_i/\xi_{i-1} \longrightarrow \xi_{i+j}/\xi_{i+j-1}$$

that is called the (i,j)-**curvature**. All the curvatures together endow $\mathcal{L}(\xi)$ with a fibrewise Lie bracket compatible with the grading. We will say that $\mathcal{L}(\xi)$ is a bundle of positively graded Lie algebras. These algebras need not be modelled on a single graded Lie algebra (since the Lie bracket is allowed to vary smoothly from fibre to fibre) and therefore the bundle need not be locally trivial.

Note that $\mathcal{L}(\xi)$ (regarded as a graded vector bundle) is the graded version of ξ_{i_0} (regarded as a vector bundle filtered by the ξ_i). In particular, there is a vector bundle isomorphism between the two that is unique up to homotopy. A concrete way of defining such an isomorphism is by selecting a metric on TM .

1.5.0.4 Formal distributions

$\mathcal{L}(\xi)$ captures the non-involutivity of ξ in a more refined manner than the Lie flag. We can think of it as a partial formal datum associated to ξ . Indeed, by construction, $\mathcal{L}(\xi)$ is uniquely determined by the $(i_0 - 1)$ -jet of ξ at each individual point. Because we are interested in differential relations that depend only on the curvatures, we are happy to forget the full jet and focus on $\mathcal{L}(\xi)$ instead; this motivates the upcoming definitions.

In general, given a foliation \mathcal{F} , we will say that a *formal \mathcal{F} -generating distribution* is a positively graded Lie algebra bundle structure on $T\mathcal{F}$ such that the degree-1 part is a generating set. The space of formal \mathcal{F} -generating distributions is denoted by $\text{Dist}^f(\mathcal{F})$; we topologise it using the (weak) C^∞ -topology. In particular, in families, each of the graded pieces and the bracket vary smoothly and therefore the rank of each graded piece remains constant.

We denote by $\text{Dist}(\mathcal{F})$ the space of regular distributions that are contained in and generate by Lie brackets \mathcal{F} . We similarly topologise it using the C^∞ -topology, turning the nilpotentisation procedure described above into a continuous inclusion:

$$\mathcal{L} : \text{Dist}(\mathcal{F}) \longrightarrow \text{Dist}^f(\mathcal{F}).$$

As we pointed out before, \mathcal{L} is only defined up to homotopy, but a concrete and consistent choice for all distributions at once can be made by choosing a metric on TM .

1.5.0.5 Formal distributions with constraints

The inclusion \mathcal{L} becomes more interesting once we introduce some natural differential constraints. Fix a GL-invariant open \mathcal{U} in the space of positively graded Lie algebras of dimension $\text{rank}(\mathcal{F})$. Recall that we are interested in distributions whose first layer bracket-generates the rest. This implies that the smallest \mathcal{U} we want to look at consists of those Lie algebras generated by their degree-1 part.

We will say that $F \in \text{Dist}^f(\mathcal{F})$ is a **formal \mathcal{U} -distribution** if $F(p) \in \mathcal{U}$ for all $p \in M$. Note that an identification of \mathcal{F}_p with $\mathbb{R}^{\text{rank}(\mathcal{F})}$ is needed for this to make sense, but the concrete choice we make is irrelevant due to GL-invariance. The subspace of all such F is denoted by $\text{Dist}^f(\mathcal{F}, \mathcal{U})$. This process effectively lifts \mathcal{U} to a Diff-invariant differential relation $\mathcal{R}_{\mathcal{U}}$ contained in the space of $(i_0 - 1)$ -jets of distributions. Its solutions (i.e. those distributions whose nilpotentisation takes values in \mathcal{U}) will be denoted by $\text{Dist}(\mathcal{F}, \mathcal{U})$.

The nilpotentisation map can be regarded then as an inclusion

$$\mathcal{L}_{\mathcal{U}} : \text{Dist}(\mathcal{F}, \mathcal{U}) \longrightarrow \text{Dist}^f(\mathcal{F}, \mathcal{U}),$$

that we sometimes call the **scanning map**. The main question in the topological study of distributions reads:

Question 1.6. Fix \mathcal{U} . Is $\mathcal{L}_{\mathcal{U}}$ a weak homotopy equivalence (for any foliated manifold and relative to boundary conditions)?

A positive answer to this question is often phrased by saying that the differential relation $\mathcal{R}_{\mathcal{U}}$ satisfies the full h -principle.

1.5.0.6 Two results of Gromov

We remind the reader that Gromov's method of flexible sheaves [57] applies to open and Diff-invariant differential relations to provide a full h -principle over *open manifolds*. All the *non-foliated* examples $\text{Dist}(TM, \mathcal{U})$ described above fit within this scheme due to the openness and GL-invariance of \mathcal{U} , as long as M is open.

Another well-known remark of Gromov says that the foliated case can be regarded as a parametric version of the standard case [59]. More precisely, the following claims are equivalent:

- The full h -principle holds for $\text{Dist}(T\mathcal{F}, \mathcal{U})$, for all foliations \mathcal{F} of rank n .
- The full h -principle holds for $\text{Dist}(TM, \mathcal{U})$, for all manifolds M of dimension n .

Due to these observations, we will tackle Question 1.9 for the case $T\mathcal{F} = TM$, where M is a closed manifold.

1.6 Convex integration with avoidance and classification of bracket-generating distributions.

Convex integration appeared first in the work of J. Nash on C^1 isometric immersions/embeddings [85]. Roughly speaking, the idea is that a short immersion can be corrected, one codirection at a time, by introducing oscillations that increase its length. This process can be iterated in such a way that, after infinitely many corrections at progressively smaller scales, one obtains an isometric map that is only C^1 .

In [58], M. Gromov turned the ideas of Nash into a scheme capable of constructing and classifying solutions of more general differential relations¹. The implementation is rather involved (particularly for differential relations of order higher than one), but the rough idea remains the same: We start with an arbitrary section f , which we correct one derivative at a time, inductively in the order of the derivatives. Namely, given a locally defined codirection λ and an order k , we add oscillations to f in order to adjust its pure derivative of order k along λ . Once we iterate over all orders and, for each order, over a well-chosen collection of codirections, we will have corrected all the derivatives of f , yielding a solution of our differential relation \mathcal{R} . There are several subtleties one must deal with:

Why openness?

On a given step, we correct the pure derivative of order k along some λ . As we do so, we may introduce errors in all other derivatives, potentially destroying what we had achieved in previous steps. Therefore, a key part of the argument is proving that oscillations along λ can be added at the expense of adding arbitrarily small errors in all other derivatives of order at most k .

This leads us to restrict our attention to open relations, because the errors will then be absorbed by openness. Note that, under this assumption of openness, we do not need to introduce infinitely many corrections at different scales anymore. This differs from the isometric immersion case.

¹ The crucial observation of Gromov is that the arguments presented in [85] and [94, 67] can be understood as integration processes. Recent work of M. Theillière [99, 100] connects Gromov's approach to this earlier literature, showing that, for certain differential relations, one can perform convex integration *without integrating*, relying instead on explicit corrugations. This yields solutions with self-similarity properties.

The formal datum

Another key point is that we need to make sense of what “correcting” is. Indeed, at each step we must study the space of all possible oscillations of f along λ and select one that is closer to being a solution of \mathcal{R} . In order to do this, our initial data will not be f but a pair (f, F) , where F is a formal solution of \mathcal{R} (i.e. a choice of Taylor polynomial solving \mathcal{R} at each point). The formal datum F guides the convex integration process: at each step we add oscillations to both f and F so that their derivatives along λ agree. The process terminates when we produce a holonomic pair (g, G) (i.e. G is the Taylor polynomial of g at all points and, since G is a formal solution, g is thus a solution).

Ampleness

The argument we are outlining only works if, at each step, we can find suitable oscillations for f and F . The way to do this is to consider $\text{Pr}_{\lambda, F}$; the space consisting of all Taylor polynomials that differ from F only in the direction of λ (and in order k); we call this the *principal subspace* associated to F and λ . Inside of $\text{Pr}_{\lambda, F}$ we can find $\mathcal{R}_{\lambda, F}$, the subset of Taylor polynomials that are still solutions of \mathcal{R} . Our oscillations will be chosen within this subset.

Crucially, we know that $\mathcal{R}_{\lambda, F}$ is non-empty, because it contains F . It is also open by assumption. To carry out the proof we also require that it is *ample*: this means that the connected component $\tilde{\mathcal{R}}_{\lambda, F} \subset \mathcal{R}_{\lambda, F}$ containing F has the whole of $\text{Pr}_{\lambda, F}$ as its convex hull. The geometric way of interpreting this condition is that the space of admissible order- k derivatives along λ is large and, upon integration, can be used to approximate any Taylor polynomial of order one less.

A relation \mathcal{R} is said to be ample if each $\mathcal{R}_{\lambda, F}$ is ample.

Ampleness in coordinate directions

We are then interested in open and ample relations \mathcal{R} . In practice, openness is readily checked, but ampleness takes some effort: a priori, it is a condition that has to be verified for each formal solution F , each order k , and each codirection λ . However, as is apparent from the explanations above, one need not study all λ , but only sufficiently many of them (finitely many per chart) to correct all derivatives. A pair $(\mathcal{R}, \{\lambda_i\})$ consisting of a relation \mathcal{R} and a suitable collection of codirections $\{\lambda_i\}$ is said to be *ample in coordinate directions* if this weaker condition holds.

In [39, p. 171], Y. Eliashberg and N. Mishachev posed the following question:

Question 1.7. Is there a (geometrically meaningful) differential relation \mathcal{R} that is ample in coordinate directions but not ample?

The present thesis provides the first such examples. We construct them using a convex integration scheme that we call *convex integration up to avoidance*. This is our main contribution and we introduce it next.

1.6.1 Statement of the main result

This thesis extends the applicability of convex integration to open relations that may not be ample nor ample in coordinate directions. The relevant (weaker) condition that they must satisfy instead is called *ampleness up to avoidance* (Definition 5.3.2). Our main theorem reads:

Theorem 1.8. *The full C^0 -close h -principle holds for differential relations that are open and ample up to avoidance.*

This result is restated in slightly more generality in Theorem 5.4, Section 5.3. We emphasise that we do not require our differential relations to be of first order.

Ampleness up to avoidance effectively allows us to take the relation of interest \mathcal{R} and a formal solution $F : M \rightarrow \mathcal{R}$, and find a smaller relation $\mathcal{R}(F) \subset \mathcal{R}$ that is now ample along coordinate directions and has F as a formal solution. In particular we can now answer Question 1.7:

Corollary 1.6.1 *There are open relations \mathcal{R} such that:*

- \mathcal{R} is ample up to avoidance and Diff-invariant.
- \mathcal{R} is not ample nor ample in coordinate directions.
- Each relation $\mathcal{R}(F)$ is ample in coordinate directions, but not necessarily ample nor Diff-invariant.

A concrete example is given in Theorem 1.12 below, which is our main application. To get there and introduce the rest of our applications, we go into the theory of tangent distributions.

1.6.2 Applications in the study of distributions

We now review what is known about Question 1.9 for various choices of \mathcal{U} . We will explain what the contributions of this thesis are as we go along. Our main application is Theorem 1.12 below.

h -Principle for step 2

We denote by $\text{Dist}(\mathcal{F})$ the space of smooth distributions that are contained in $\mathcal{F} \subset TM$, a smooth foliation, and generate by Lie brackets. We can consider the inclusion of spaces of distributions in their formal counterpart $\text{Dist}^f(\mathcal{F})$. These are spaces that encode the underlying algebraic topology of the spaces of distributions in the sense of the philosophy of the h -principle. (See subsection 1.5.0.4 for further details)

$$\mathcal{L} : \text{Dist}(\mathcal{F}) \longrightarrow \text{Dist}^f(\mathcal{F}).$$

Formal distributions with constraints

The inclusion \mathcal{L} becomes more interesting once we introduce some natural differential constraints. Fix a GL-invariant open \mathcal{U} in the space of positively graded Lie algebras of dimension $\text{rank}(\mathcal{F})$. Recall that we are interested in distributions whose first layer bracket-generates the rest. This implies that the smallest \mathcal{U} we want to look at consists of those Lie algebras generated by their degree-1 part.

We will say that $F \in \text{Dist}^f(\mathcal{F})$ is a **formal \mathcal{U} -distribution** if $F(p) \in \mathcal{U}$ for all $p \in M$. Those distributions whose nilpotentisation takes values in \mathcal{U} will be denoted by $\text{Dist}(\mathcal{F}, \mathcal{U})$.

We can then consider the analogous inclusion of these spaces of distributions into their formal counterpart:

$$\mathcal{L}_{\mathcal{U}} : \text{Dist}(\mathcal{F}, \mathcal{U}) \longrightarrow \text{Dist}^f(\mathcal{F}, \mathcal{U}),$$

that we sometimes call the **scanning map**. The main question in the topological study of distributions reads:

Question 1.9. Fix \mathcal{U} . Is $\mathcal{L}_{\mathcal{U}}$ a weak homotopy equivalence (for any foliated manifold and relative to boundary conditions)?

A positive answer to this question is often phrased by saying that the differential relation $\mathcal{R}_{\mathcal{U}}$ satisfies the full h -principle.

Let M be a smooth manifold. One expects the answer to Question 1.9 to be positive if the differential constraints we introduce are rather weak (i.e. if \mathcal{U} is large). As we stated above, the weakest assumption we are interested in is that \mathcal{U} consists of Lie algebras generated by their first layer, so $\text{Dist}(TM, \mathcal{U})$ is the space of bracket-generating distributions in M .

Under this weak assumption we prove:

Theorem 1.10. *Let M be a smooth manifold of dimension at least 4. The complete C^0 -close h -principle holds for bracket-generating distributions of step 2 in M .*

The result is sharp, since the 3-dimensional case corresponds to contact structures, which are known not to abide by the full h -principle [11]. The proof is presented in Section 6.1 and is, in fact, a routinary application of convex integration.

Remark 1.6.2 *Theorem 1.10 partially answers an open question raised during the workshop on Engel Structures held in April 2017 at AIM (American Institute of Mathematics, San Jose, California). Concretely, [38, Problem 6.2] asks whether any parallelizable n -manifold admits a k -plane field $\xi \subset TM$ with maximal growth vector. This question is further refined to ask whether any formal distribution of maximal growth admits a holonomic representative up to homotopy.*

For step 2, our result answers the question and its refinement positively and goes a bit beyond. Indeed, for $k > 3$ we provide a full classification in terms of formal data and not just a existence statement. On the other hand, we do not tackle the higher step case. This is left as an interesting open question.

Maximal non-involutivity

Theorem 1.10 says that being bracket-generating is a very flexible condition (in dimension 4 onwards). As such, we would like to consider more restrictive assumptions on \mathcal{U} . Our guiding example is Contact Topology, the study of contact structures. These are distributions whose nilpotentisation is non-degenerate, in the sense that the first curvature is a non-degenerate two-form. This is equivalent to the fact that a contact structure has as many non-trivial Lie brackets as possible. This non-degeneracy is, ultimately, responsible for the contact scanning map not being a homotopy equivalence in general [11], even though partial flexibility results do hold [34]. In the last few years we have seen spectacular progress in our understanding of higher-dimensional contact structures [12, 84].

We will henceforth focus on distributions presenting a similar flavour of non-degeneracy. We will call this **maximal non-involutivity**; the precise meaning of this will be explained for each dimension and rank as we go along.

Even-contact structures

In even dimensions, a hyperplane field is maximally non-involutive if its curvature has corank 1 (i.e. it has a 1-dimensional kernel). Such distributions are called even-contact structures. For them, Question 1.9 was answered positively by McDuff [83], proving that they are (topologically) much more flexible than contact structures. However, interesting questions about them from a geometric perspective remain open [86].

Dimensions 3 and 4

In dimension 3, a bracket-generating distribution is necessarily a contact structure; we have already mentioned that the h -principle fails for them. In dimension 4, a corank-1 regular bracket-generating distribution is an even-contact structure.

The remaining case in dimension 4 corresponds to rank 2. In this situation, a maximally non-involutive distribution is a regular bracket-generating distribution of step 3; these are called Engel

structures. Various results have appeared in the last few years regarding their classification [103, 29, 25, 92] and the classification of their submanifolds [52, 90, 24] but a definite answer to Question 1.9 is still open.

Dimension 5

In dimension 5, maximally non-involutive hyperplanes are contact structures.

Rank-3 distributions are maximally non-involutive if they are of step 2. In particular, as a corollary of Theorem 1.10, we have:

Theorem 1.11. *Let M be 5 or 6 dimensional. The complete C^0 -close h -principle holds for maximally non-involutive rank-3 distributions.*

Maximally non-involutive distributions of rank-2 are the so-called (2,3,5) distributions of Cartan [21], which have been classified only in open manifolds [31]. If we replace maximal non-involutivity by some concrete closed growth-vector condition, there are other interesting classes of distributions (e.g. Goursat structures) whose classification is open as well.

Dimension 6

Our main application concerns rank 4 distributions in 6-dimensional manifolds. It turns out that maximally non-involutive (4,6)-distributions come in two families, *elliptic* and *hyperbolic*. The statement reads:

Theorem 1.12. *Let M be a 6-dimensional manifold. The complete C^0 -close h -principle holds for rank-4 distributions of hyperbolic type.*

The proof can be found in Section 7.2 and it is a consequence of our main result Theorem 1.8. We emphasise that this result requires ameness up to avoidance and is beyond the scope of classic convex integration.

Remark 1.6.3 *We conjecture that the answer to Question 1.9 is negative for elliptic (4,6) distributions. In Corollary 7.2.5 we will show that ameness does not hold for the differential relation that defines them.*

The remaining cases are: Corank-1 (which are even-contact structures), rank-3 (classified by Theorem 1.11) and rank-2 (the so-called (2,3,5,6) structures, for which nothing is known).

1.7 Guide to the contents of the thesis.

1.7.1 Chapter 1. Introduction.

We introduce the general framework of bracket-distributions together with an overview of the main problems solved in this thesis.

1.7.2 Chapter 2. The fundamental group of Formal Legendrian embeddings.

We first reprove a folklore result in the area, the classification of formal legendrian embeddings by their formal invariants:

Theorem 1.13 (Theorem 2.3). *Formal Legendrian embeddings are classified by their parametrized knot type, rotation number and Thurston–Bennequin invariant.*

We then move to the main result in the chapter: the computation of the fundamental group of the space of Formal Legendrian embeddings.

Theorem 1.14 (Theorem 2.4). *The sequence*

$$0 \longrightarrow \mathbb{Z} \oplus \mathbb{Z}_m \longrightarrow \pi_1(\mathfrak{FLeg}(\mathbb{R}^3)) \longrightarrow \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \longrightarrow 0$$

is exact, where $m \geq 0$ (this integer depends on the connected component). In particular, if we fix the connected component $\mathfrak{Emb}'(\mathbb{S}^1, \mathbb{R}^3) \subseteq \mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ of the parametrized unknot or of the parametrized (p, q) torus knot we have that $m = 0$ and so

$$0 \longrightarrow \mathbb{Z} \oplus \mathbb{Z} \longrightarrow \pi_1(\mathfrak{FLeg}(\mathbb{R}^3)) \longrightarrow \pi_1(\mathfrak{Emb}'(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \longrightarrow 0$$

is exact.

Finally we get the main application: the non-triviality of Kálmán's loop at the formal level:

Proposition 1.7.1 (Proposition 2.5.1) *The loop of formal Legendrian embeddings $(\gamma^{\theta, k}, F_s^{\theta, k})$ of Kálmán is non trivial for any $k \in \mathbb{Z}$.*

1.7.3 Chapter 3. Connected sum of loops of Legendrian embeddings.

In this chapter we define a **connected sum of loops of legendrian embeddings** in \mathbb{S}^3 that allow us construct new loops of Legendrian embeddings. Along the process get the following commutative diagram:

$$\begin{array}{ccccccc} 0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) & \longrightarrow & \pi_1(\mathfrak{Leg}(\mathbb{R}^3, \xi_{\text{std}})) & \longrightarrow & \pi_1(\mathfrak{Leg}(\mathbb{S}^3, \xi_{\text{std}})) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) / \pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) & \longrightarrow & \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) & \longrightarrow & \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) \longrightarrow 0 \end{array}$$

By using this diagram we can explain how to perform the aforementioned parametric connected sum for loops of Legendrians in \mathbb{R}^3 . As the main application we construct infinitely many new examples of loops of Legendrians which are non-trivial:

Theorem 1.15 (Theorem 3.4. Infinitely many new examples of loops with non-trivial monodromy.).

The monodromy at the level of H_0 associated to the following families of loops is not the identity and, thus, all these families of loops are not contractible in the space $\mathfrak{Leg}(\mathbb{R}^3, \xi_{\text{std}})$:

For $m \geq 1$, the connected sum loop of Kálmán's loop based at a trefoil and m loops K_{p_i, q_i}^θ where for $1 \leq i \leq m$ is one of the following:

- a) Kálmán's loop based at the torus knot K_{p_i, q_i} ,
- b) Kálmán's inverse loop based at the torus knot K_{p_i, q_i} , or
- c) The constant loop $K_{p_i, q_i}^{\text{const}}$ based at the torus knot K_{p_i, q_i} .

1.7.4 Chapter 4. h -Principle for horizontal and transverse curves.

The h -principles for horizontal curves are proven in Section 4.7. The h -principles for transverse curves in Section 4.8. Along the way we state and prove the appropriate relative versions. We will put all our emphasis on the embedding cases; the other statements (immersions and smooth curves) follow from the same arguments with considerable simplifications.

We state the h -principle for regular immersions first:

Theorem 1.16 (Theorem 1.1). *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathfrak{Imm}^r(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}^f(M, \mathcal{D}).$$

We then deduce as a corollary:

Corollary 1.7.2 *Let \mathcal{D}_0 and \mathcal{D}_1 be bracket-generating distributions on a manifold M , homotopic as subbundles of TM . Then, the spaces $\mathfrak{Imm}^r(M, \mathcal{D}_0)$ and $\mathfrak{Imm}^r(M, \mathcal{D}_1)$ are weakly homotopy equivalent.*

We prove that the h -principle holds for regular embeddings into manifolds of dimension at least 4:

Theorem 1.17 (Theorem 1.2). *Let (M, \mathcal{D}) be a bracket-generating distribution with $\dim(M) \geq 4$. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathfrak{Emb}^r(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}^f(M, \mathcal{D}).$$

And we get the following corollary:

Corollary 1.7.3 *Fix a manifold M with $\dim(M) \geq 4$. Let \mathcal{D}_0 and \mathcal{D}_1 be bracket-generating distributions on M , homotopic as subbundles of TM . Then, the spaces $\mathfrak{Emb}^r(M, \mathcal{D}_0)$ and $\mathfrak{Emb}^r(M, \mathcal{D}_1)$ are weakly homotopy equivalent.*

We also address the homotopy type of the smooth horizontal loop space, proving the following h -principle:

Theorem 1.18 (Theorem 1.3). *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then, the following inclusion is a weak homotopy equivalence:*

$$\mathcal{L}(M, \mathcal{D}) \longrightarrow \mathcal{L}(M).$$

Finally, we deal with the transverse case. We prove an h -principle for transverse immersions:

Theorem 1.19. *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then the inclusion*

$$\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}_{\mathcal{T}}^f(M, \mathcal{D})$$

is a weak homotopy equivalence.

And we also show that the h -principle holds in the transverse embedded case:

Theorem 1.20 (Theorem 1.5). *Let (M, \mathcal{D}) be a bracket-generating distribution with $\dim(M) \geq 4$. Then the inclusion*

$$\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}_{\mathcal{T}}^f(M, \mathcal{D})$$

is a weak homotopy equivalence. In particular, $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})$ depends only on the formal class of \mathcal{D} .

1.7.5 Chapter 5. Convex integration with avoidance.

In this Chapter we prove the Theorem of convex integration up to avoidance:

Theorem 1.21 (Theorem 1.8). *The full C^0 -close h -principle holds for differential relations that are open and ample up to avoidance.*

This result is restated in slightly more generality in Theorem 5.4, Section 5.3.

We can now answer Question 1.7:

Corollary 1.7.4 *There are open relations \mathcal{R} such that:*

- \mathcal{R} is ample up to avoidance and Diff-invariant.
- \mathcal{R} is not ample nor ample in coordinate directions.
- Each relation $\mathcal{R}(F)$ is ample in coordinate directions, but not necessarily ample nor Diff-invariant.

1.7.6 Chapter 6. h -Principle for step-2 distributions.

As an application of classic convex integration we prove:

Theorem 1.22 (Theorem 1.10). *Let M be a smooth manifold of dimension at least 4. The complete C^0 -close h -principle holds for bracket-generating distributions of step 2 in M .*

1.7.7 Chapter 7. h -Principle for hyperbolic $(4, 6)$ -distributions.

Finally we devote this chapter to the main application of Convex integration up to avoidance (Theorem 1.8): the h -principle for hyperbolic $(4, 6)$ distributions.

Theorem 1.23 (Theorem 1.12). *Let M be a 6-dimensional manifold. The complete C^0 -close h -principle holds for rank-4 distributions of hyperbolic type.*

1.7.8 Chapter 8. Future work.

In this last chapter we address future work and open questions in the field of bracket-generating distributions.

Part I

Legendrian and Formal Legendrian embedding spaces

Chapter 2

The fundamental group of the space of Formal Legendrian embeddings.

2.1 Spaces of embeddings of the circle into euclidean space.

Denote by $\mathfrak{Emb}(N, M)$ the space of embeddings of a manifold N into a manifold M equipped with the C^∞ -topology.

2.1.1 The space $\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)$.

Theorem 2.1 (Hatcher, [63] Appendix: equivalence (15)). *The space of parametrized unknotted circles in \mathbb{R}^3 has the homotopy type of $\mathrm{SO}(3)$.*

The group $\mathrm{SO}(4)$ acts freely on the connected component $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \subseteq \mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)$ of the parametrized (p, q) torus knots as

$$\begin{aligned} \mathrm{SO}(4) \times \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) &\longrightarrow \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \\ (A, \gamma) &\longmapsto A \cdot \gamma. \end{aligned}$$

Thus, we have an inclusion $\mathrm{SO}(4) \hookrightarrow \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$. The following result holds:

Theorem 2.2 (Hatcher, [65] Theorem 1). *The inclusion $\mathrm{SO}(4) \hookrightarrow \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$ is a homotopy equivalence.*

As a consequence of these results we obtain that

Corollary 2.1.1 *Let $\mathfrak{Emb}_0(\mathbb{S}^1, \mathbb{R}^3), \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3) \subseteq \mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ be the connected component of the parametrized unknots or of the parametrized (p, q) torus knots, respectively. The fundamental groups of these spaces are given by*

- $\pi_1(\mathfrak{Emb}_0(\mathbb{S}^1, \mathbb{R}^3)) \cong \mathbb{Z}_2$,
- $\pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)) \cong G_{p,q} \rtimes \mathbb{Z}_2$,

where $G_{p,q}$ is the knot group of the (p, q) torus knot.

Proof. The case of the connected component $\mathfrak{Emb}_0(\mathbb{S}^1, \mathbb{R}^3)$ follows from Theorem 2.1.

We need to study the connected component $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)$ to conclude the proof. Consider the following space

$$\mathcal{S}tereo_{p,q} = \{(\gamma, x) : x \notin \text{Im}(\gamma)\} \subseteq \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \times \mathbb{S}^3.$$

We have two natural fibrations associated to the projection maps

$$\begin{array}{ccc} \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3) & \hookrightarrow & \mathcal{S}tereo_{p,q} \\ \downarrow & & \downarrow \\ \mathbb{S}^3 & & \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \end{array} \quad \begin{array}{ccc} \mathbb{S}^3 \setminus K_{p,q} & \hookrightarrow & \mathcal{S}tereo_{p,q} \\ & & \downarrow \\ & & \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \end{array}$$

where $K_{p,q}$ is the image of the standard (p, q) torus knot in \mathbb{S}^3 . From the first fibration we obtain

$$\pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)) \cong \pi_1(\mathcal{S}tereo_{p,q}).$$

Moreover, from the second one and the fact that $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$ has the homotopy type of $\text{SO}(4)$ (Theorem 2.2), we obtain that the sequence

$$0 \longrightarrow G_{p,q} \longrightarrow \pi_1(\mathcal{S}tereo_{p,q}) \longrightarrow \mathbb{Z}_2 \longrightarrow 0$$

is exact. Now, it is a simple exercise to check that this sequence is right split.

2.2 Formal Legendrian Embeddings in \mathbb{R}^3 .

We denote by ξ the standard contact structure in \mathbb{R}^3 (using coordinates (x, y, z)) given by $\xi = \text{Ker}(dz - ydx)$. Throughout the Section we fix the Legendrian framing ∂_y .¹

2.2.1 Formal Legendrian Embeddings in \mathbb{R}^3 .

Definition 2.2.1 *An immersion $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is said to be Legendrian if $\gamma'(t) \in \xi_{\gamma(t)}$ for all $t \in \mathbb{S}^1$. If γ is an embedding, we say it is a Legendrian embedding.*

Definition 2.2.2

(a) A formal Legendrian immersion in \mathbb{R}^3 is a pair (γ, F) such that:

- (i) $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is a smooth map.
- (ii) $F : \mathbb{S}^1 \rightarrow \gamma^*(T\mathbb{R}^3 \setminus \{0\})$ satisfies $F(t) \in \xi_{\gamma(t)}$, where 0 is the zero section.

(b) A formal Legendrian embedding in \mathbb{R}^3 is a pair (γ, F_s) , satisfying:

- (i) $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is an embedding.
- (ii) $F_s : \mathbb{S}^1 \rightarrow \gamma^*(T\mathbb{R}^3 \setminus \{0\})$, is a 1-parametric family, $s \in [0, 1]$, such that $F_0 = \gamma'$ and $F_1(t) \in \xi_{\gamma(t)}$.

Use the framing $\langle \partial_y \rangle$ (see Footnote) to trivialize the contact distribution understood as a bundle. This provides a bundle isomorphism $\xi \simeq \mathbb{R}^2$. From now on, we will understand the map $F : \mathbb{S}^1 \rightarrow$

¹ Observe that ξ is naturally (co)oriented by the contact form $dz - ydx$ and, thus, ∂_y determines a unique oriented framing up to homotopy.

$\mathbb{S}^1 \equiv \mathbb{S}^2 \cap \mathbb{R}^2$ and the family $F_s : \mathbb{S}^1 \rightarrow \mathbb{S}^2$ with $F_1 : \mathbb{S}^1 \rightarrow \mathbb{S}^1 \equiv \mathbb{S}^2 \cap \mathbb{R}^2$. We say that an immersion is *strict* if it is a non injective map.

Denote by $\mathbf{LegImm}(\mathbb{R}^3)$ the space of Legendrian immersions in \mathbb{R}^3 and by $\mathbf{Leg}(\mathbb{R}^3)$ the space of Legendrian embeddings in \mathbb{R}^3 . Denote also by $\mathfrak{F}\mathbf{LegImm}(\mathbb{R}^3)$ the space of formal Legendrian immersions and by $\mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)$ the space of formal Legendrian embeddings. These definitions make sense for immersions and embeddings of the interval. We define $\mathbf{LegImm}([0, 1], \mathbb{R}^3)$, $\mathbf{Leg}([0, 1], \mathbb{R}^3)$, $\mathfrak{F}\mathbf{LegImm}([0, 1], \mathbb{R}^3)$ and $\mathfrak{F}\mathbf{Leg}([0, 1], \mathbb{R}^3)$ analogously.

All the spaces of Legendrians are equipped with the C^∞ -topology. On the other hand, the spaces of formal Legendrians are equipped with the product topology that is the C^∞ -topology for the first factor (the smooth immersion/embedding) and the C^∞ -topology for the second factor (the formal derivative).

Remark 2.2.3 *It is well-known that h -principle holds for Legendrian immersions (see, eg, Eliashberg and Mishachev [39]). Hence, $\pi_0(\mathbf{LegImm}(\mathbb{R}^3)) \cong \mathbb{Z}$, $\pi_1(\mathbf{LegImm}(\mathbb{R}^3)) \cong \mathbb{Z}$ and $\pi_k(\mathbf{LegImm}(\mathbb{R}^3)) = 0$, for all $k \geq 2$. The connected components of $\mathbf{LegImm}(\mathbb{R}^3)$ are given by the rotation number. The rotation number of an Legendrian immersion γ is $\text{Rot}(\gamma) = \deg(\gamma' : \mathbb{S}^1 \rightarrow \mathbb{S}^1)$. Let us explain the group $\pi_1(\mathbf{LegImm}(\mathbb{R}^3)) \cong \mathbb{Z}$. Take a loop γ^θ in $\mathbf{LegImm}(\mathbb{R}^3)$, the integer is just $\text{Rot}_L(\gamma^\theta) = \deg(\theta \mapsto (\gamma^\theta)'(0))$, we call this number rotation number of the loop. These invariants make sense in the formal case and the definitions are the obvious ones.*

2.2.2 The space $\mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)$.

Consider the space $\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3) = \{(\gamma, F) | \gamma \in \mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3), F \in \mathbf{Maps}(\mathbb{S}^1, \mathbb{S}^1)\}$. We have a natural fibration $\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3) \rightarrow \mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)$. In order to compute the homotopy groups of $\mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)$, take $\gamma \in \mathbf{Leg}(\mathbb{R}^3)$ and fix (γ, γ') as base point. The fiber over this point is $\mathcal{F} = \mathcal{F}_{(\gamma, \gamma')} = \Omega_{\gamma'}(\mathbf{Maps}(\mathbb{S}^1, \mathbb{S}^2))$, where this denotes the space of loops in the space $\mathbf{Maps}(\mathbb{S}^1, \mathbb{S}^2)$ based at γ' . We have the following exact sequence of homotopy groups associated to the fibration:

$$\begin{array}{ccccccc}
 & & \cdots & \longrightarrow & \pi_2(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) & & \\
 & & & \swarrow & & & \\
 \pi_1(\mathcal{F}) & \longrightarrow & \pi_1(\mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)) & \longrightarrow & \pi_1(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) & & \\
 & & & \swarrow & & & \\
 \pi_0(\mathcal{F}) & \longrightarrow & \pi_0(\mathfrak{F}\mathbf{Leg}(\mathbb{R}^3)) & \longrightarrow & \pi_0(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) & \longrightarrow & 0
 \end{array}$$

Notice that $\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)$ has the homotopy type of $\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3) \times \mathbb{S}^1 \times \mathbb{Z}$. Hence, $\pi_0(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) \cong \pi_0(\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z}$, where the integer is the rotation number. Moreover, $\pi_1(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) \cong \pi_1(\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z}$ and the \mathbb{Z} factor is given by the rotation number of the loop. Finally, $\pi_k(\widehat{\mathfrak{F}\mathbf{Leg}}(\mathbb{R}^3)) \cong \pi_k(\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3))$ for all $k \geq 2$.

The homotopy groups of \mathcal{F} are easily computable. Just observe that there is a fibration $\mathbf{Maps}(\mathbb{S}^1, \mathbb{S}^2) \rightarrow \mathbb{S}^2$ defined via the *evaluation map*, with fiber over $p \in \mathbb{S}^2$ given by $\Omega_p(\mathbb{S}^2)$. As every element $[f] \in \pi_n(\mathbb{S}^2)$ can be lifted to an element $[f_n] \in \pi_n(\mathbf{Maps}(\mathbb{S}^1, \mathbb{S}^2))$, defined as

$$f_n(p)(t) = f(p), t \in \mathbb{S}^1, p \in \mathbb{S}^n,$$

all the diagonal maps in the associated exact sequence to the fibration $\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2) \rightarrow \mathbb{S}^2$ are zero. This implies that there are short exact sequences $\pi_n(\Omega_p(\mathbb{S}^2)) \rightarrow \pi_n(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)) \rightarrow \pi_n(\mathbb{S}^2)$ for $n \geq 1$. In particular, since \mathbb{S}^2 is simply connected, we obtain that

$$\pi_0(\mathcal{F}) \cong \pi_1(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)) \cong \pi_1(\Omega_p(\mathbb{S}^2)) \cong \pi_2(\mathbb{S}^2) \cong \mathbb{Z}.$$

Moreover, these sequences are right split and, thus, split for $n > 2$ since the groups involved are abelian. So, we have

$$\begin{aligned} \pi_1(\mathcal{F}) &\cong \pi_2(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)) \cong \pi_2(\mathbb{S}^2) \oplus \pi_2(\Omega_p(\mathbb{S}^2)) \cong \pi_2(\mathbb{S}^2) \oplus \pi_3(\mathbb{S}^2) \cong \mathbb{Z} \oplus \mathbb{Z} \\ \text{and } \pi_2(\mathcal{F}) &\cong \pi_3(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)) \cong \pi_3(\mathbb{S}^2) \oplus \pi_3(\Omega_p(\mathbb{S}^2)) \cong \pi_3(\mathbb{S}^2) \oplus \pi_4(\mathbb{S}^2) \cong \mathbb{Z} \oplus \mathbb{Z}_2. \end{aligned}$$

Lemma 2.2.4 $\pi_0(\widehat{\mathfrak{Leg}}(\mathbb{R}^3)) \cong \pi_0(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \oplus \mathbb{Z}.$

Proof. It is sufficient to show that every element in $\pi_1(\widehat{\mathfrak{Leg}}(\mathbb{R}^3)) \cong \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \pi_1(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^1)) \cong \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z}$ can be lifted to an element in $\pi_1(\widehat{\mathfrak{Leg}}(\mathbb{R}^3))$.

Take a loop $(\gamma^\theta, F_1^\theta)$ in $\widehat{\mathfrak{Leg}}(\mathbb{R}^3)$. Let $F_0 = (\gamma^\theta)' : \mathbb{S}^1 \times \mathbb{S}^1(\theta, t) \rightarrow \mathbb{S}^2$ be the derivative $F_0(\theta, t) = (\gamma^\theta)'(t)$, we need to show that $F_0 = (\gamma^\theta)'$ is homotopic to the map $F_1 : \mathbb{S}^1 \times \mathbb{S}^1 \rightarrow \mathbb{S}^2$. Observe that the homotopy classes of maps from $\mathbb{S}^1 \times \mathbb{S}^1$ to \mathbb{S}^2 are classified by the degree and $\deg(F_1) = 0$, so we just need to show that $\deg(F_0) = 0$ to complete the proof.

Indeed, the map

$$G_\varepsilon(\theta, t) = \begin{cases} (\gamma^\theta)'(t) & \text{if } \varepsilon = 0, \\ \frac{\gamma^\theta(t+\varepsilon) - \gamma^\theta(t)}{\|\gamma^\theta(t+\varepsilon) - \gamma^\theta(t)\|} & \text{if } 0 < \varepsilon < 1, \\ -(\gamma^\theta)'(t) & \text{if } \varepsilon = 1, \end{cases} \quad (2.1)$$

is well-defined, because $\gamma^\theta, \theta \in \mathbb{S}^1$, is an embedding. Thus, $F_0 = (\gamma^\theta)'$ is homotopic to $-F_0 = -(\gamma^\theta)'$ and $\deg(F_0) = \deg(\gamma^\theta)' = 0$.

2.3 Classification of formal Legendrian embeddings in \mathbb{R}^3 .

We have checked that $\pi_0(\widehat{\mathfrak{Leg}}(\mathbb{R}^3)) \cong \pi_0(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \oplus \mathbb{Z}$. The first \mathbb{Z} corresponds to the *rotation number* and we will show that the second one corresponds to the *Thurston–Bennequin invariant*.

Let us refine the definition of formal Legendrian embedding to extend the definition of the *Thurston–Bennequin invariant* to the formal case.

Definition 2.3.1 A formal extended Legendrian embedding in \mathbb{R}^3 is a pair (γ, G_s) , satisfying:

- (i) $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is a embedding.
- (ii) $G_s : \mathbb{S}^1 \rightarrow \text{SO}(3)$, is a smooth family in the parameter $s \in [0, 1]$, such that $G_0 = \text{Id}$ and $G_1(\gamma') \in \xi_{\gamma(t)}$.

We denote $\widehat{\mathfrak{Leg}}(\mathbb{R}^3)$ for the space of formal extended Legendrian embeddings in \mathbb{R}^3 equipped with the C^∞ -topology in the first factor and the C^∞ -topology in the second one.

Remark 2.3.2 *The natural fibration $t : \mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3) \rightarrow \mathfrak{F}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3), (\gamma, G_s) \mapsto (\gamma, G_s(\gamma'))$, has contractible fibers. Note that the fiber is the space of paths in $\mathrm{SO}(2)$ starting at the Identity map and this space is contractible. Thus, this map is a weak homotopy equivalence.*

Given $(\gamma, G_s) \in \mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)$ we have a well-defined *formal contact framing* $\mathcal{F}_{\mathrm{FCont}}$ of the normal bundle $\nu(G_1(\gamma'))$ given by the Legendrian condition $G_1(\gamma') \subseteq \xi_{\gamma(t)}$. Then, $G_1^{-1}(\mathcal{F}_{\mathrm{FCont}})$ defines a framing of the normal bundle ν of γ . On the other hand, we have a *topological framing* $\mathcal{F}_{\mathrm{Top}}$ of ν given by a Seifert surface of γ .

Definition 2.3.3 *Let $(\gamma, G_s) \in \mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)$. The Thurston–Bennequin invariant is $\mathrm{tb}(\gamma, G_s) = \mathrm{tw}_\nu(G_1^{-1}(\mathcal{F}_{\mathrm{FCont}}), \mathcal{F}_{\mathrm{Top}})$, ie the twisting of $G_1^{-1}(\mathcal{F}_{\mathrm{FCont}})$ with respect to $\mathcal{F}_{\mathrm{Top}}$.*

The Thurston–Bennequin invariant is defined over $\widehat{\mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3) = \{(\gamma, G) | \gamma \in \mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3), G \in \mathfrak{Maps}(\mathbb{S}^1, \mathrm{SO}(3)), G(\gamma') \in \xi_{\gamma(t)}\}$. Furthermore, since the unique oriented \mathbb{S}^1 -bundle over \mathbb{S}^1 is the trivial one, $\pi_0(\widehat{\mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)) \cong \pi_0(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \oplus \mathbb{Z}$. The first \mathbb{Z} is just the *rotation number* and the second one corresponds to the *Thurston–Bennequin invariant*.

Now we can state the main result of this Section, which is folklore.

Theorem 2.3. *Formal Legendrian embeddings are classified by their parametrized knot type, rotation number and Thurston–Bennequin invariant.*

The proof of this result follows directly using the fibration $\hat{F} : \mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3) \rightarrow \widehat{\mathfrak{F}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)$ and the fact that the map $t : \mathfrak{F}\mathfrak{E}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3) \rightarrow \mathfrak{F}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)$ is a weak homotopy equivalence. Note also that the fibration \hat{F} has connected fiber, because its π_0 is given by $\pi_2(\mathrm{SO}(3)) = 0$. This completes the proof.

However to get a more geometric picture, we will express the isomorphism $\pi_0(t) \circ \pi_0(\hat{F})^{-1}$ in more concrete terms. It can be shown that it coincides with the one in Lemma 2.2.4 by comparing the maps involved. Clearly the isomorphism preserves the rotation invariant, ie the rotation number (γ, G_s) is sent to the rotation invariant of $(\gamma, G_s(\gamma'))$. To understand the rest of the isomorphism we fix a base point $(\gamma, F = \gamma')$ in $\widehat{\mathfrak{F}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)$ with $\mathrm{Rot}(\gamma, \gamma') = 0$, ie we declare the base point to be a Legendrian embedding with zero rotation. Now, given an element of the fiber, ie (γ, F^s) with $F^0 = F^1 = \gamma'$, we claim that for $[(\gamma, 0, k)] \in \pi_0(\mathfrak{F}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3))$, the isomorphism $\pi_0(t)$ is given by $\mathrm{tb}(\pi_0(t)^{-1}(\gamma, 0, k)) = \mathrm{tb}(\gamma, \gamma') - 2k$. In other words, it depends on the choice of base point. This is obvious if we check that given a double stabilization of the Legendrian knot, the value of the degree invariant in the fiber increases by 1 and it is a simple computation to check that the tb decreases by 2.

2.4 Computation of the Fundamental group of formal Legendrian Embeddings in \mathbb{R}^3 .

As a consequence of Lemma 2.2.4 we have that the following sequence is exact:

$$\begin{array}{ccccccc}
 \cdots & \longrightarrow & \pi_2(\mathfrak{F}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)) & \longrightarrow & \pi_2(\widehat{\mathfrak{F}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)) & & \\
 & & & \searrow & & & \\
 \pi_1(\mathcal{F}) & \longrightarrow & \pi_1(\mathfrak{F}\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)) & \longrightarrow & \pi_1(\widehat{\mathfrak{F}\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)) & \longrightarrow & 0
 \end{array}$$

Take a 2-sphere (γ^z, F_1^z) in $\widehat{\mathfrak{FLeg}}(\mathbb{R}^3)$, the diagonal map $d : \pi_2(\widehat{\mathfrak{FLeg}}(\mathbb{R}^3)) \rightarrow \pi_1(\mathcal{F})$ measures the obstruction to lifting (γ^z, F_1^z) to a 2-sphere (γ^z, F_s^z) in $\mathfrak{FLeg}(\mathbb{R}^3)$, ie the obstruction to find a homotopy between the derivative map $F_0 = (\gamma^z)' : \mathbb{S}^2 \times \mathbb{S}^1 \rightarrow \mathbb{S}^2$, $(z, t) \mapsto F_0(z, t) = (\gamma^z)'(t)$, and the map $F_1 : \mathbb{S}^2 \times \mathbb{S}^1 \rightarrow \mathbb{S}^2$. Note that by the Legendrian condition F_1 is nullhomotopic, since it is not surjective. The first obstruction to find this homotopy is just the degree of $F_0(z, 0) = (\gamma^z)'(0)$ and corresponds to the first \mathbb{Z} factor of $\pi_1(\mathcal{F}) \cong \pi_2(\mathbb{S}^2) \oplus \pi_3(\mathbb{S}^2) \cong \mathbb{Z} \oplus \mathbb{Z}$. In particular, since $(\gamma^z)'(0)$ is homotopic to $-(\gamma^z)'(0)$ this obstruction vanishes (see equation (2.1)).

Theorem 2.4. *The sequence*

$$0 \longrightarrow \mathbb{Z} \oplus \mathbb{Z}_m \longrightarrow \pi_1(\mathfrak{FLeg}(\mathbb{R}^3)) \longrightarrow \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \longrightarrow 0$$

is exact, where $m \geq 0$. In particular, if we fix the connected component $\mathfrak{Emb}'(\mathbb{S}^1, \mathbb{R}^3) \subseteq \mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ of the parametrized unknot or of the parametrized (p, q) torus knot we have that $m = 0$ and so

$$0 \longrightarrow \mathbb{Z} \oplus \mathbb{Z} \longrightarrow \pi_1(\mathfrak{FLeg}(\mathbb{R}^3)) \longrightarrow \pi_1(\mathfrak{Emb}'(\mathbb{S}^1, \mathbb{R}^3)) \oplus \mathbb{Z} \longrightarrow 0$$

is exact.

Proof. The Z_m denotes the subgroup of $Z = \pi_3(\mathbb{S}^2)$ which comes from the embedding part in the diagonal map above. Thus, we only need to check the particular cases mentioned above. For the connected component $\mathfrak{Emb}_0(\mathbb{S}^1, \mathbb{R}^3)$ of the parametrized unknot the result follows from Theorem 2.1, since $\pi_2(\widehat{\mathfrak{FLeg}}_0(\mathbb{R}^3)) = \pi_2(\mathfrak{Emb}_0(\mathbb{S}^1, \mathbb{R}^3)) = \pi_2(\mathrm{SO}(3)) = 0$, where $\widehat{\mathfrak{FLeg}}_0(\mathbb{R}^3)$ stands for a formal Legendrian connected component of the smooth unknot.

On the other hand, fix the connected component $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)$ of the parametrized (p, q) torus knot and consider the commutative diagram

$$\begin{array}{ccc} \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3) & \hookrightarrow & \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3) \\ \downarrow & & \downarrow \\ \mathfrak{Imm}(\mathbb{S}^1, \mathbb{R}^3) & \hookrightarrow & \mathfrak{Imm}(\mathbb{S}^1, \mathbb{S}^3) \end{array}$$

defined by the natural inclusions, where $\mathfrak{Imm}(N, M)$ denotes the space of immersions of a manifold N into a manifold M equipped with the C^r -topology, $r \geq 5$.

By the Smale–Hirsch Theorem for immersions (see [95] or [39]) we have that $\mathfrak{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ has the homotopy type of $\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)$ and $\mathfrak{Imm}(\mathbb{S}^1, \mathbb{S}^3)$ has the homotopy type of $\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^3) \times \mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)$. Moreover, the map induced by the inclusion $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3) \hookrightarrow \mathfrak{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ at π_2 -level sends the homotopy class $[\gamma^z] \in \pi_2(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3))$ to $[(\gamma^z)'] \in \pi_2(\mathfrak{Imm}(\mathbb{S}^1, \mathbb{R}^3)) \cong \pi_2(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2))$; ie coincides with the diagonal map $d : \pi_2(\widehat{\mathfrak{FLeg}}(\mathbb{R}^3)) \rightarrow \pi_1(\mathcal{F}) \cong \pi_2(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2))$.

Consider the induced commutative diagram at π_2 -level

$$\begin{array}{ccc} \pi_2(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)) & \longrightarrow & \pi_2(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)) \\ \downarrow & & \downarrow \\ \pi_2(\mathfrak{Imm}(\mathbb{S}^1, \mathbb{R}^3)) & \longrightarrow & \pi_2(\mathfrak{Imm}(\mathbb{S}^1, \mathbb{S}^3)) \end{array}$$

since $\pi_2(\mathbf{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3))$ is trivial (see Theorem 2.2) it is sufficient to show that the homomorphism $\pi_2(\mathcal{Imm}(\mathbb{S}^1, \mathbb{R}^3)) \rightarrow \pi_2(\mathcal{Imm}(\mathbb{S}^1, \mathbb{S}^3))$ is injective to conclude the proof. But this is clear, by using the h -principle for immersions, since the degree of the induced map $\mathbb{S}^2 \times \mathbb{S}^1$ to \mathbb{S}^3 is zero and the induced map for the derivative from $\mathbb{S}^2 \times \mathbb{S}^1$ to \mathbb{S}^2 is sent to itself by the inclusion.

2.5 Main application: formal non-triviality of Kálmán's loop.

2.5.1 Kálmán's loop.

Kálmán has constructed a series of examples of loops of Legendrian positive torus knots non-contractible in the space $\mathcal{Leg}(\mathbb{R}^3)$, though contractible in $\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ [71]. Let us prove that Kálmán's examples are non trivial even as loops of formal Legendrian embeddings; that is, in the space $\mathfrak{Leg}(\mathbb{R}^3)$. We will prove that they are not contractible for any choice of parametrization. Since the space of Legendrians can be seen inside the space of Formal Legendrians, thus these loops are not contractible as loops of unparametrized oriented Legendrian knots. Consider a




Figure 2.1: The loop in front projection ($p = 3, q = 7$). Note that we have to cycle $2p = 6$ times.

Legendrian positive (p, q) torus knot, a loop is described in Figure 2.1. The loop takes the p strands of the knot to the cyclic rotation of them. This geometrically corresponds to a $2\pi/p$ rotation along the core of the defining torus. Let us consider $2p$ concatenations of this loop. Thus, it is generated by two full rotations along the core of the torus.

2.5.1.1 Simplified position.

First, we will deform through formal loops the initial loop into a formal loop in a “simplified” position.

Step 1. Consider the contactomorphism $f(x, y, z) = (x/r, y/r, z/r^2)$ in the standard $(\mathbb{R}^3, \text{Ker}(dz - ydx))$. By using it, we assume that the defining torus for the loop has arbitrarily small meridional radius. Therefore, the knot is C^1 -close to the core β of the torus. We are not using the standard notion of C^1 -closeness, but a weaker one. We mean that a sequence of immersions $\hat{\gamma}_k$ is C^1 -close to an immersion $\tilde{\gamma}$ if for any $\varepsilon > 0$ and for any point $t \in \mathbb{S}^1$: for every $k \in \mathbb{Z}$ large enough there exists a point $\tau(t) \in \mathbb{S}^1$ such that $|\hat{\gamma}_k(t) - \tilde{\gamma}(\tau(t))| \leq \varepsilon$ and $|\frac{\hat{\gamma}'_k(t)}{\|\hat{\gamma}'_k(t)\|} - \frac{\tilde{\gamma}'(\tau(t))}{\|\tilde{\gamma}'(\tau(t))\|}| \leq \varepsilon$. This is a way of somehow formalising the notion of closeness for points and tangent vectors when regarded the knots as submanifolds (non-parametrized).

Moreover, by further shrinking, the knot and the core β , they may be assumed to be arbitrarily close to $\text{Ker } dz$; ie C^1 -close to their Lagrangian projections (i.e. the plane $\{z = 0\}$).




Figure 2.2: C^1 -approximation of the knot to the core, shown in the front projection.

Step 2. Denote γ^θ the initial loop of Legendrian embeddings. Understood as a formal loop, it is written as $(\gamma^\theta, F_s^\theta)$, where $F_s^\theta = (\gamma^\theta)'$. Let us construct a 1-parametric family of formal loops $(\gamma^{\theta,u}, F_s^{\theta,u})$, $u \in [0, 1]$, defined as follows

- (i) $\gamma^{\theta,u} = \gamma^\theta$,
- (ii) $F_s^{\theta,u} = (1-s)(\gamma^\theta)' + s((1-u)(\gamma^\theta)' + u\partial_y)$.

This is a family of formal loops, since $(\gamma^\theta)'$ is never a negative multiple of ∂_y .² To check that, note that β is C^1 -close to γ^θ .

Step 3. We consider the family of rotations $\{r_v \in \text{SO}(3)\}_{v \in [0,1]}$ taking the quadrant XY to the quadrant $-ZX$, see Figure 2.3. Construct a family of formal loops $(\gamma^{\theta,u}, F_s^{\theta,u})$, $u \in [1, 2]$, as follows

- (i) $\gamma^{\theta,u} = r_{u-1} \cdot \gamma^\theta$,
- (ii) $F_s^{\theta,u} = (1-s)(\gamma^{\theta,u})' + s\partial_y$.

Again, this is a family of formal loops because $(\gamma^{\theta,u})'$ is never a negative multiple of ∂_y . Note that we are using that γ^θ is C^1 -close to β .

Step 4. Finally, we turn over the left lobe of the unknot core $r_1 \cdot \beta$ by an isotopy defined as follows. Take polar coordinates (r, φ) in the plane YZ and $\varepsilon > 0$ small enough. We define the isotopy as:

$$f_u(x, r, \varphi) = \begin{cases} (x, r, \varphi + u \cdot \chi(x) \cdot \pi) & \text{if } x \leq \varepsilon, \\ (x, r, \varphi) & \text{if } \varepsilon \leq x \end{cases}$$

where $0 \leq u \leq 1$ and $\chi : \mathbb{R} \rightarrow \mathbb{R}$ is a non-decreasing smooth function satisfying

² We are assuming an orientation of the Legendrians. The argument with the opposite orientation runs in the same way by changing ∂_y by $-\partial_y$.

- $\chi(x) = 1$ for all $x \leq 0$,
- $\chi(x) = 0$ for all $x \geq \varepsilon$.

We apply the isotopy to $r_1 \cdot \gamma^\theta$, see Figure 2.3. Again, the derivative is never tangent to $-\partial_y$ and thus we can interpolate to ∂_y .

We have proven that our initial loop of Legendrian embeddings is homotopic to the loop of formal Legendrian embeddings $(\tilde{\gamma}^\theta, \tilde{F}_s^\theta)$ defined as follows:

- (i) $\tilde{\gamma}^\theta$ is the loop of parametrized (p, q) torus knots supported in the torus associated to the unknot contained in the plane XZ . The loop is obtained by a rotation of 4π radians of the standard (p, q) -embedding in the direction of the parallel of the supporting torus.
- (ii) $\tilde{F}_s^\theta = (1 - s)(\tilde{\gamma}^\theta)' + s\partial_y$.




Figure 2.3: Construction of the path of loops. We represent the moves of the core β .

2.5.2 Set of parametrizations of the family of loops.

As an outcome of the previous discussion, we may assume that our formal Legendrian parametrized (p, q) torus knot can be written as (γ, F_s) , where

- $\gamma(t) = \begin{pmatrix} (\cos(2\pi pt) + 2) \cos(2\pi qt) \\ \sin(2\pi pt) \\ (\cos(2\pi pt) + 2) \sin(2\pi qt) \end{pmatrix},$
- $F_s(t) = s\partial_y + (1 - s)(\gamma)'(t).$

One particular parametrization of the loop can be written as $(\gamma^\theta, F_s^\theta)$, where

- $\gamma^\theta(t) = \begin{pmatrix} \cos(4\pi\theta) & 0 & -\sin(4\pi\theta) \\ 0 & 1 & 0 \\ \sin(4\pi\theta) & 0 & \cos(4\pi\theta) \end{pmatrix} \begin{pmatrix} (\cos(2\pi pt) + 2) \cos(2\pi qt) \\ \sin(2\pi pt) \\ (\cos(2\pi pt) + 2) \sin(2\pi qt) \end{pmatrix},$
- $F_s^\theta(t) = s\partial_y + (1 - s)(\gamma^\theta)'(t).$

We will show that any possible parametrization of the loop gives raise to a non-trivial loop of parametrized formal Legendrian embeddings. Up to homotopy, the possible parametrizations of the formal Legendrian loop are given by:

- $\gamma^{\theta,k}(t) = \gamma^\theta(t + k\theta)$,
- $F_s^{\theta,k}(t) = F_s^\theta(t + k\theta) = s\partial_y + (1-s)(\gamma^{\theta,k})'(t)$,

where $k \in \mathbb{Z}$. This is because $\pi_1(\text{Diff}^+(\mathbb{S}^1)) = \pi_1(\text{SO}(2)) = \mathbb{Z}$.³

We will prove the following statement.

Proposition 2.5.1 *The loop of formal Legendrian embeddings $(\gamma^{\theta,k}, F_s^{\theta,k})$ is non trivial for any $k \in \mathbb{Z}$.*

This proves that the loop is non trivial as a loop of non parametrized formal Legendrian knots.

2.5.3 Proof of Proposition 2.5.1.

It follows from the previous discussion that the loop $\gamma^{\theta,k}(t)$ of smooth embeddings lies in $\text{SO}(4) \subset \mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$, ie $\gamma^{\theta,k}(t) = A_{\theta,k}\gamma(t)$, where $A_{\theta,k} \in \text{SO}(4)$. More specifically, on $\mathbb{S}^3(\sqrt{2})$, the $\sqrt{2}$ radius sphere in \mathbb{C}^2 , we have

$$\gamma^{\theta,0}(t) = \begin{pmatrix} 1 & 0 \\ 0 & e^{4\pi i \theta} \end{pmatrix} \begin{pmatrix} e^{2\pi i p t} \\ e^{2\pi i q t} \end{pmatrix}.$$

Thus, in these coordinates the other parametrizations are given by

$$\gamma^{\theta,k}(t) = \gamma^{\theta,0}(t + k\theta) = \begin{pmatrix} e^{2\pi i p k \theta} & 0 \\ 0 & e^{2\pi i (2+qk)\theta} \end{pmatrix} \begin{pmatrix} e^{2\pi i p t} \\ e^{2\pi i q t} \end{pmatrix}.$$

By Theorem 2.2 the parametrized loop $\gamma^{\theta,k}$ is trivial in $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$ if and only if $2 + k(p+q)$ is even. From now on we will assume that this is the case. Thus, there is a family $\{\tilde{A}_{(r,\theta),k}\}_{(r,\theta) \in \mathbb{D}}$ such that $\tilde{A}_{(1,\theta),k} = A_{\theta,k}$. Since $\pi_2(\text{SO}(4)) = 0$, the disk $\tilde{A}_{(r,\theta),k}$ is unique up to homotopy fixing the boundary and the same holds for the disk $\tilde{A}_{(r,\theta),k} \cdot \gamma$ in $\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)$.

By Theorem 2.4, we have the following exact sequence:

$$0 \longrightarrow \pi_1(\mathcal{F}) \xrightarrow{m_1} \pi_1(\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3)) \xrightarrow{m_2} \pi_1(\widehat{\mathfrak{L}\mathfrak{eg}}(\mathbb{R}^3)) \longrightarrow 0 \quad (2.2)$$

Thus, in order to prove that our loop $\ell_k = [(\gamma^{\theta,k}, F_s^{\theta,k})] \in \pi_1(\mathfrak{L}\mathfrak{eg}(\mathbb{R}^3))$ is non trivial we distinguish two cases:

2.5.3.1 Case 1. $k \neq 0$.

We claim that $m_2(\ell_k) \neq 0$, ie $[\gamma^{\theta,k}] \neq 0 \in \pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3))$.⁴

³ Remember that the Legendrian knots are oriented.

⁴ Observe that $\text{Rot}_L(\gamma^{\theta,k}, F_s^{\theta,k}) = 0$.

Recall from Corollary 2.1.1 that $\pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{R}^3)) \cong \pi_1(\mathcal{Stereop}_{p,q})$ and that we have an exact sequence

$$0 \longrightarrow G_{p,q} \longrightarrow \pi_1(\mathcal{Stereop}_{p,q}) \longrightarrow \pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3)) \longrightarrow 0$$

Thus, we must show that $[(A_{\theta,k}\gamma, \infty)] \in \pi_1(\mathcal{Stereop}_{p,q})$ is non trivial. It is the family of loops that is trivial by hypothesis in $\pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3))$ does not admit a capping disk whose evaluation map avoids $\infty \in \mathbb{S}^3$. We check it by composing with the 1-parametric family of loops $\tilde{A}_{(r,\theta),k}^{-1} \in \text{SO}(4)$, we obtain a 1-parametric family of loops $(\tilde{A}_{(r,\theta),k}^{-1}A_{\theta,k}\gamma, \tilde{A}_{(r,\theta),k}^{-1}\infty)$, $r \in [0, 1]$. For $r = 0$ we obtain the initial loop and for $r = 1$ we obtain the loop $(\gamma, A_{\theta,k}^{-1}\infty)$. Thus, these two loops represent the same element of $\pi_1(\mathcal{Stereop}_{p,q})$. Moreover, $[(\gamma, A_{\theta,k}^{-1}\infty)]$ can be lifted to $G_{p,q}$, since it lies on the fiber defined by the element $[\gamma] \in \pi_1(\mathfrak{Emb}_{p,q}(\mathbb{S}^1, \mathbb{S}^3))$. So, we are reduced to check whether $[A_{\theta,k}^{-1}\infty] \in G_{p,q}$ is trivial. The knot group of the (p, q) torus knot is $G_{p,q} = \langle a, b : a^p = b^q \rangle$. Thus, $[A_{\theta,k}^{-1}\infty] = b^{pk} \neq 0$, since b is a non torsion element of $G_{p,q}$.




Figure 2.4: Visualization of the loop $\gamma^{\theta,k}$ for a $(5, 2)$ torus knot ($k \neq 0$).

2.5.3.2 Case 2. $k = 0$.

Since $m_2(\ell_0) \in \pi_1(\widehat{\mathfrak{FLeg}}(\mathbb{R}^3))$ is zero, there exists $A \in \pi_1(\mathcal{F})$ such that $m_1(A) = \ell_0$. We are going to geometrically check that $A \neq 0$ and therefore, by the injectivity of m_1 provided by the sequence (2.2), the non triviality of ℓ_0 follows.

Write $(\gamma^\theta, F_s^\theta) = (\gamma^{\theta,0}, F_s^{\theta,0})$. Note that in $(\mathbb{R}^3(x, y, z), \text{Ker}(dz - ydx))$ the parametrized loop is written as:

- $\gamma^\theta(t) = B_\theta \gamma(t) = \begin{pmatrix} \cos(4\pi\theta) & 0 & -\sin(4\pi\theta) \\ 0 & 1 & 0 \\ \sin(4\pi\theta) & 0 & \cos(4\pi\theta) \end{pmatrix} \begin{pmatrix} (\cos(2\pi pt) + 2)\cos(2\pi qt) \\ \sin(2\pi pt) \\ (\cos(2\pi pt) + 2)\sin(2\pi qt) \end{pmatrix},$
- $F_s^\theta(t) = s\partial_y + (1-s)(\gamma^\theta)'(t).$

Moreover, γ^θ is bounded by $\tilde{\gamma}^{(r,\theta)}(t) = \tilde{B}_{(r,\theta)}\gamma(t)$, where $\tilde{B}_{(r,\theta)} \in \text{SO}(3)$, such that $\tilde{B}_{(1,\theta)} = B_\theta$ and is homotopic to $\tilde{A}_{(r,\theta)}$ inside $\text{SO}(4)$. Thus, we have a disk $\mathcal{D}(r, \theta) = \{(\tilde{B}_{r,\theta}\gamma, \partial_y)\}$ in $\widehat{\mathfrak{FLeg}}(\mathbb{R}^3)$ that bounds $m_2(\ell_0)$. We try to lift it to a disk in $\mathfrak{FLeg}(\mathbb{R}^3)$ that bounds ℓ_0 . There is no homotopical obstruction to lifting it for the punctured disk $\mathcal{D}(r, \theta)|_{r>0}$. Therefore, the homotopy obstruction is represented by an element of $\pi_1(\mathcal{F})$. So we obtain a loop $(\gamma_0, \tilde{F}_s^\theta)$ over the fiber of (γ_0, ∂_y) where $\gamma_0 = B_\theta\gamma$. It follows by construction that $A = [(\gamma_0, \tilde{F}_s^\theta)]$. Let us perform the computation.




Figure 2.5: Visualization of the loop $\gamma^{\theta,0}$ for a $(5,2)$ torus knot ($k=0$).

$$\tilde{F}_s^\theta = \begin{cases} \tilde{B}_{2s,\theta}(\gamma') & \text{if } 0 \leq s \leq \frac{1}{2}, \\ F_{2s-1}^\theta = (2s-1)\partial_y + (2-2s)B_\theta(\gamma') & \text{if } \frac{1}{2} \leq s \leq 1. \end{cases} \quad (2.3)$$

Observe that $\tilde{F}_0^\theta = \gamma'_0$ and $\tilde{F}_1^\theta = \partial_y$. Thus, we can understand \tilde{F}_s^θ as a map $\tilde{F} : \mathbb{S}^2 \rightarrow \mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)$. It follows that $A = [\tilde{F}] \in \pi_1(\mathcal{F}) \cong \pi_2(\mathfrak{Maps}(\mathbb{S}^1, \mathbb{S}^2)) \cong \pi_2(\mathbb{S}^2) \oplus \pi_2(\Omega_p(\mathbb{S}^2))$, that is the fundamental group of the fiber. We already computed this group in Theorem 2.4. Moreover, we also showed that the morphism to $\pi_1(\mathfrak{Leg}(\mathbb{R}^3))$ induced by the inclusion is injective.

To conclude the proof we must check that $[\tilde{F}] \neq 0$. In order to see this, we will verify that $\deg(\tilde{F}_s^\theta(0) : \mathbb{S}^2 \rightarrow \mathbb{S}^2)$ is nonzero. This degree is the first coordinate of $[\tilde{F}] \in \mathbb{Z} \oplus \mathbb{Z} \cong \pi_1(\mathcal{F}) \cong \pi_2(\mathbb{S}^2) \oplus \pi_2(\Omega_p(\mathbb{S}^2))$.

We give an explicit description of $\{\tilde{B}_{r,\theta}\}$ in $\text{SO}(3)$. Identify $\text{SO}(3)$ with \mathbb{RP}^3 in the usual way. We understand \mathbb{RP}^3 as the 3-ball of radius π with its boundary points identified via the antipodal map. Then a point $p = (p_1, p_2, p_3)$ in the described 3-ball corresponds in $\text{SO}(3)$ to the rotation of angle $\sqrt{p_1^2 + p_2^2 + p_3^2}$ in $\mathbb{R}^3(x, y, z)$ around the axis described by its position vector p . In these coordinates, see the left drawing in Figure 2.6, the loop B_θ is given by

$$B_\theta = \begin{cases} (0, 4\pi\theta, 0) & \text{if } 0 \leq \theta \leq \frac{1}{4}, \\ (0, -2\pi + 4\pi\theta, 0) & \text{if } \frac{1}{4} \leq \theta \leq \frac{3}{4}, \\ (0, -4\pi + 4\pi\theta, 0) & \text{if } \frac{3}{4} \leq \theta \leq 1. \end{cases}$$

Define the disk $\{\tilde{B}_{r,\theta}\}$ as the intersection of the plane $\{z=0\}$ with the π -radius ball. It produces an $\mathbb{RP}^2 = \mathbb{RP}^3 \cap \{z=0\} \subseteq \mathbb{RP}^3$. We have $\hat{B} = \{B_\theta : \theta \in \mathbb{S}^1\} \subseteq \mathbb{RP}^2$. We obtain $\mathbb{RP}^2 \setminus \hat{B}$ is an embedded 2-disk. See Figure 2.6.

We have that $\gamma'(0) = 2\pi p\partial_y + 6\pi q\partial_z$. Substituting in equation (2.3) for $t=0$, we get

$$\tilde{F}_s^\theta(0) = \begin{cases} \tilde{B}_{2s,\theta}(2\pi p\partial_y + 6\pi q\partial_z) & \text{if } 0 \leq s \leq \frac{1}{2}, \\ F_{2s-1}^\theta = (2s-1)\partial_y + (2-2s)B_\theta(2\pi p\partial_y + 6\pi q\partial_z) & \text{if } \frac{1}{2} \leq s \leq 1. \end{cases}$$

Moreover, the maps, $u \in [0, 1]$,

$$G(s, \theta, u) = \begin{cases} \tilde{B}_{2s,\theta}((1-u)(2\pi p\partial_y + 6\pi q\partial_z) + u\partial_z) & \text{if } 0 \leq s \leq \frac{1}{2}, \\ F_{2s-1}^\theta = (2s-1)\partial_y + (2-2s)B_\theta((1-u)(2\pi p\partial_y + 6\pi q\partial_z) + u\partial_z) & \text{if } \frac{1}{2} \leq s \leq 1 \end{cases}$$

are always non zero. Thus, the map $\tilde{F}_s^\theta(0) = G(s, \theta, 0)$ is homotopic to

$$G(s, \theta) = G(s, \theta, 1) = \begin{cases} \tilde{B}_{2s, \theta}(\partial_z) & \text{if } 0 \leq s \leq \frac{1}{2}, \\ (2s - 1)\partial_y + (2 - 2s)B_\theta(\partial_z) & \text{if } \frac{1}{2} \leq s \leq 1. \end{cases}$$

In order to compute the degree of $G(s, \theta)$, we write $G(s, \theta) = (g_x(s, \theta), g_y(s, \theta), g_z(s, \theta)) \in \mathbb{S}^2(\partial_x, \partial_y, \partial_z)$ and we check that $\#G^{-1}(-\partial_y) = 1$:

- If $\frac{1}{2} \leq s \leq 1$ then $G(s, \theta)$ is a linear combination with positive coefficients between ∂_y and $B_\theta(\partial_z) \in \mathbb{S}^1(\partial_x, \partial_z) \subseteq \mathbb{S}^2(\partial_x, \partial_y, \partial_z)$ thus $g_y(s, \theta) \geq 0$.
- If $0 \leq s \leq \frac{1}{2}$ then $G(s, \theta)$ is just a rotation around an axis in the XY -plane acting over ∂_z . The rotation of angle $\frac{\pi}{2}$ around the X -axis is the unique rotation that sends ∂_z to $-\partial_y$. Thus, $G^{-1}(-\partial_y) = \{(s_0, \theta_0)\}$ where (s_0, θ_0) is the only point that satisfies that $\tilde{B}_{2s_0, \theta_0}$ is the mentioned rotation.

The map G is a local diffeomorphism in a neighborhood of the point (s_0, θ_0) . Thus, $-\partial_y$ is a regular value for G and $|\deg(\tilde{F}_s^\theta(0))| = |\deg(G(s, \theta))| = 1 \neq 0$.




Figure 2.6: Explicit construction of the capping disk.

Chapter 3

Connected sum of parametric families of Legendrian embeddings.

3.1 Topological relations between $\pi_1(\mathfrak{Leg}(\mathbb{R}^3, \xi_{\text{std}}))$ and $\pi_1(\mathfrak{Leg}(\mathbb{S}^3, \xi_{\text{std}}))$.

Let $(\mathbb{R}^3, \xi_{\text{std}} = \text{Ker}(dz - ydx))$ be the standard contact structure on $\mathbb{R}^3(x, y, z)$ and $(\mathbb{S}^3, \xi_{\text{std}} = \frac{i}{2} \sum (z_j d\bar{z}_j - \bar{z}_j dz_j)$ be the standard contact structure on $\mathbb{S}^3 \subseteq \mathbb{C}^2(z_1, z_2)$. Throughout the section (M, ξ_{std}) denotes $(\mathbb{R}^3, \xi_{\text{std}})$ or $(\mathbb{S}^3, \xi_{\text{std}})$ unless other thing is said.

On $(\mathbb{R}^3, \xi_{\text{std}})$ fix the framings $\xi_{\text{std}} = \langle \partial_x + y\partial_z, \partial_y \rangle$ and $T\mathbb{R}^3 = \langle \partial_x + y\partial_z, \partial_y, \partial_z \rangle$; and on $(\mathbb{S}^3, \xi_{\text{std}})$, with quaternionic notation, the framings $\xi_{\text{std},p} = \langle jp, kp \rangle$ and $T_p\mathbb{S}^3 = \langle ip, jp, kp \rangle$. The choice of framings is unique up to homotopy. From now on we understand $F_s : \mathbb{S}^1 \rightarrow \mathbb{S}^2$ and $F_1 : \mathbb{S}^1 \rightarrow \mathbb{S}^1 = \mathbb{S}^2 \cap \xi_{\text{std}}$; i.e. we mean that after trivialising the tangent bundle $TM = T\mathbb{R}^3$ for $M = \mathbb{R}^3, \mathbb{S}^3$ we consider norm-1 vectors. Denote by LX the free loop space of a connected manifold X , then for any formal Legendrian embedding (γ, F_s) we have that $F_1 \in LS^1$ and F_s defines a path in LS^2 between γ' and F_1 . The *evaluation map* $LX \rightarrow X$ defines a fibration with fiber $\Omega_p(X)$. This fibration has a section which assigns to each point in X the constant loop based at that point. Moreover, if X is a Lie group then $LX \cong X \times \Omega_{Id}(X)$. Denote by $\mathfrak{Leg}(M, \xi_{\text{std}})$ the space of formal Legendrian embeddings in (M, ξ_{std}) . The aim is to understand the map induced in homotopy by the natural inclusion

$$\begin{aligned} j : \mathfrak{Leg}(M, \xi_{\text{std}}) &\hookrightarrow \mathfrak{Leg}(M, \xi_{\text{std}}) \\ \gamma &\longmapsto (\gamma, F_s \equiv \gamma'). \end{aligned} \quad (3.1)$$

In this thesis we are interested in studying the homomorphism $\pi_1(j)$.

In this subsection we study the topological relations between the fundamental groups of the spaces $\mathfrak{Leg}(\mathbb{R}^3, \xi_{\text{std}})$ and $\mathfrak{Leg}(\mathbb{S}^3, \xi_{\text{std}})$.

Denote by $\mathfrak{Leg}_{p,v}(M, \xi_{\text{std}})$, $p \in M$ and $v \in \mathbb{S}(\xi_{\text{std}})_p$, where $\mathbb{S}(\xi_{\text{std}})$ denotes the sphere bundle of ξ_{std} ; i. e. $\mathfrak{Leg}_{p,v}(M, \xi_{\text{std}})$ is the subspace of Legendrian embeddings γ such that $\gamma(0) = p$ and $\gamma'(0) = v$. Analogously, define the space $\mathfrak{Leg}_{p,v}(M, \xi_{\text{std}})$ as the space of formal Legendrian embeddings (γ, F_s) such that $\gamma(0) = p$ and $F_1(0) = v$. Consider the auxiliary space

$$\text{Stereo}\mathfrak{Leg}_{N,jN} = \{(\gamma, p) : p \notin \gamma(\mathbb{S}^1)\} \subseteq \mathfrak{Leg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}}) \times (\mathbb{S}^3 \setminus \{N\}).$$

Recall that the standard contact structure over \mathbb{S}^3 is defined as the complex tangencies of $T\mathbb{S}^3$, thus we have a natural inclusion

$$i_{\mathbb{S}^3} : \text{U}(2) \hookrightarrow \text{Cont}(\mathbb{S}^3, \xi_{\text{std}}). \quad (3.2)$$

This inclusion has a left inverse given by

$$\begin{aligned} \text{Cont}(\mathbb{S}^3, \xi_{\text{std}}) &\longrightarrow \text{U}(2) \\ \varphi &\longmapsto A_\varphi, \end{aligned} \quad (3.3)$$

where $A_\varphi = (\varphi(N), \frac{1}{\|d\varphi_N(jN)\|} d\varphi_N(jN))$. Thus, it defines an homotopy injection.

Observe that the inclusion $\mathcal{L}\mathbf{eg}_{0,(\partial_x)_0}(\mathbb{R}^3, \xi_{\text{std}}) \hookrightarrow \mathcal{S}tereo\mathcal{L}\mathbf{eg}_{N,jN}$ induces a weak homotopy equivalence.

There is a natural map $\mathcal{L}\mathbf{eg}(\mathbb{S}^3, \xi_{\text{std}}) \rightarrow \mathcal{U}(2), \gamma \mapsto A_\gamma = (\gamma(0), \gamma'(0))$. This is a homotopy equivalence:

$$\begin{aligned} \Phi : \mathcal{L}\mathbf{eg}(\mathbb{S}^3, \xi_{\text{std}}) &\longrightarrow \mathcal{U}(2) \times \mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}}) \\ \gamma &\longmapsto (A_\gamma, A_\gamma^{-1}\gamma), \end{aligned} \quad (3.4)$$

This homotopy equivalence was observed by E. Fernández (this is explained in detail in ([46])). On the other hand there is a Serre fibration

$$\begin{array}{ccc} \mathbb{S}^3 \setminus K & \hookrightarrow & \mathcal{S}tereo\mathcal{L}\mathbf{eg}_{N,jN} \\ & & \downarrow \\ & & \mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}}), \end{array}$$

where $K = \gamma(\mathbb{S}^1)$ is the embedded Legendrian associated to a base point $\gamma \in \mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}})$. Notice that given any compact family $\gamma^z \in \mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}})$ there is a way to lift it to a family $(\gamma^z, p^z) \in \mathcal{S}tereo\mathcal{L}\mathbf{eg}_{N,jN}$ taking p^z to be the image under the Reeb flow of $\gamma^z(0)$ at time $t = \varepsilon > 0$ for some small enough $\varepsilon > 0$. Moreover, it is a consequence of the *Sphere Theorem* that the complement of a knot is aspherical (see [66, Corollary 3.9]). We conclude that

Lemma 3.1.1 *The inclusion $\mathcal{L}\mathbf{eg}_{0,(\partial_x)_0}(\mathbb{R}^3, \xi_{\text{std}}) \hookrightarrow \mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}})$ induces a right split¹ exact sequence*

$$0 \rightarrow \pi_1(\mathbb{S}^3 \setminus K) \rightarrow \pi_1(\mathcal{L}\mathbf{eg}_{0,(\partial_x)_0}(\mathbb{R}^3, \xi_{\text{std}})) \rightarrow \pi_1(\mathcal{L}\mathbf{eg}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}})) \rightarrow 0.$$

From the discussion above, for general Legendrian embeddings we conclude the following

Corollary 3.1.2 *There is a right split exact sequence*

$$0 \rightarrow \pi_1(\mathbb{S}^3 \setminus K) \rightarrow \pi_1(\mathcal{L}\mathbf{eg}(\mathbb{R}^3, \xi_{\text{std}})) \rightarrow \pi_1(\mathcal{L}\mathbf{eg}(\mathbb{S}^3, \xi_{\text{std}})) \rightarrow 0.$$

It is interesting to relate the difference between the fundamental groups of the space of Legendrians in \mathbb{R}^3 and \mathbb{S}^3 with the difference between the fundamental groups of their smooth counterparts. To do this it is enough to redo the previous argument working with the space

$$\mathcal{S}tereo = \{(\gamma, q) : q \notin \gamma(\mathbb{S}^1)\} \subseteq \mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3) \times \mathbb{S}^3,$$

in the smooth case, and the space $\mathcal{S}tereo\mathcal{L}\mathbf{eg} = \mathcal{S}tereo \cap (\mathcal{L}\mathbf{eg}(\mathbb{S}^3, \xi_{\text{std}})) \times \mathbb{S}^3$ in the Legendrian case. It follows that there is a commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) & \longrightarrow & \pi_1(\mathcal{L}\mathbf{eg}(\mathbb{R}^3, \xi_{\text{std}})) & \longrightarrow & \pi_1(\mathcal{L}\mathbf{eg}(\mathbb{S}^3, \xi_{\text{std}})) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) / \pi_2(\mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3)) & \longrightarrow & \pi_1(\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3)) & \longrightarrow & \pi_1(\mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3)) \longrightarrow 0 \end{array}$$

where the horizontal lines are exact sequences. Note that the quotient by the $\pi_2(\mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3))$ factor comes from the boundary map coming from the exact sequence of the corresponding fibration.

¹ Since we are not working with abelian groups this does not imply that the sequence splits in general.

Remark 3.1.3 *Observe that in general we cannot assume that $\pi_1(\mathbb{S}^3 \setminus K) / \pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) = \pi_1(\mathbb{S}^3 \setminus K)$. For example, consider the path component of the smooth unknot in \mathbb{S}^3 . In this case we have that the knot complement is diffeomorphic to a solid torus and, moreover, the generator of $\pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) \cong \mathbb{Z}$ can be regarded as the \mathbb{S}^2 -family of oriented lines passing through $0 \in \mathbb{R}^3 \subseteq \mathbb{S}^3$. Finding a \mathbb{S}^2 -family of points in the complement of each line is equivalent to finding a vector field tangent to \mathbb{S}^2 without zeros which is not possible by the Poincaré–Hopf Theorem. In fact, any generic vector field tangent to \mathbb{S}^2 has $\chi(\mathbb{S}^2) = 2$ zeroes. Thus, the homomorphism*

$$\pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^2)) \cong \mathbb{Z} \rightarrow \pi_1(\mathbb{S}^3 \setminus K) \cong \mathbb{Z}$$

is just multiplication by 2. This fits with the fact that the space of parametrized unknots in \mathbb{R}^3 is homotopy equivalent to $\mathrm{SO}(3)$ and in \mathbb{S}^3 to $V_{4,2}$ (see [63]).

If $\pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3))$ is zero, e.g. for torus knots, we obtain

Corollary 3.1.4 *Let $\gamma^\theta \in \mathfrak{Leg}(\mathbb{R}^3, \xi_{\mathrm{std}})$ be a loop of Legendrian embeddings.*

Assume that $\pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3), \gamma^0) = 0$. If γ^θ is trivial as an element of $\pi_1(\mathfrak{Leg}(\mathbb{S}^3, \xi_{\mathrm{std}}))$ and as an element of $\pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3))$ then it is trivial as an element of $\pi_1(\mathfrak{Leg}(\mathbb{R}^3, \xi_{\mathrm{std}}))$.

Remark 3.1.5 *It follows from this corollary that those Kálmán’s examples which are non contractible as loops of smooth embeddings in \mathbb{R}^3 (recall that this loops are non contractible as loops of (formal) Legendrian embeddings in $(\mathbb{R}^3, \xi_{\mathrm{std}})$ [71, 44]) are, in fact, non contractible as loops of Legendrian embeddings in $(\mathbb{S}^3, \xi_{\mathrm{std}})$.*

3.1.1 The space of Darboux balls in a contact manifold.

Alexander’s trick allows to prove that the space $\mathfrak{Emb}^+(\mathbb{D}^n, \mathbb{R}^n)$ of orientation preserving embeddings $\mathbb{D}^n \rightarrow \mathbb{R}^n$ linearise, i.e. it is homotopy equivalent to $\mathrm{GL}^+(n, \mathbb{R})$. This, together with the *Isotopy Extension Theorem*, implies that on a closed oriented n -manifold N^n the space $\mathfrak{Emb}^+(\mathbb{D}^n, N^n)$ is homotopy equivalent to the total space of the oriented frame bundle $\mathrm{Fr}^+(N)$. Explicitly, there is a map of fibrations

$$\begin{array}{ccc} \mathfrak{Emb}_p^+(\mathbb{D}^n, N^n) & \longrightarrow & \mathrm{GL}^+(n, \mathbb{R}) \\ \downarrow & & \downarrow \\ \mathfrak{Emb}^+(\mathbb{D}^n, N^n) & \longrightarrow & \mathrm{Fr}^+(N^n) \\ \downarrow \mathrm{ev}_0 & & \downarrow \\ N^n & \xrightarrow{\mathrm{Id}} & N^n \end{array}$$

Where the maps between the fibers and between the bases are homotopy equivalences. Thus, the natural map $\mathfrak{Emb}^+(\mathbb{D}^n, N^n) \rightarrow \mathrm{Fr}^+(N^n)$ is a homotopy equivalence.

In the contact category there is also an Alexander trick. Indeed, in $(\mathbb{R}^{2n+1}, \xi_{\mathrm{std}} = \mathrm{Ker}(dz - \sum_i y_i dx_i))$ the dilation

$$\delta_t : \mathbb{R}^{2n+1} \rightarrow \mathbb{R}^{2n+1}, (\mathbf{x}, \mathbf{y}, z) \mapsto (t\mathbf{x}, t\mathbf{y}, t^2 z);$$

is a contactomorphism for any $t > 0$. This implies that the space $\mathfrak{Emb}((\mathbb{D}^{2n+1}, \xi_{\mathrm{std}}), (\mathbb{R}^{2n+1}, \xi_{\mathrm{std}}))$ of co-oriented embeddings (preserving the contact structure with co-orientation) of Darboux balls

$(\mathbb{D}^{2n+1}, \xi_{\text{std}})$ into $(\mathbb{R}^{2n+1}, \xi_{\text{std}})$ is homotopy equivalent to the space of contact framings, i.e. to $U(n)$. We refer the reader to [53, Section 2.6.2] for further details. In particular, the Isotopy Extension Theorems in Contact Topology implies, in the same way as in the smooth case, that

Lemma 3.1.6 *Let (N, ξ) be a closed co-oriented $(2n+1)$ -contact manifold.*

The space $\mathfrak{CEmb}((\mathbb{D}^{2n+1}, \xi_{\text{std}}), (N, \xi))$ is homotopy equivalent to the total space of the bundle of contact framings $\text{CFr}(N, \xi)$ over (N, ξ) , which has fiber $U(n)$; i.e. a Darboux ball is determined by the centre of the ball and the induced framing of ξ at that point.

In particular, for $(\mathbb{S}^3, \xi_{\text{std}})$ the space of parametrized Darboux balls is homotopy equivalent to $\mathbb{S}^3 \times U(1)$.

3.2 Connected sum at the π_0 -level.

We follow Etnyre and Honda [43] to define the connected sum of a pair of Legendrian embeddings in $(\mathbb{S}^3, \xi_{\text{std}})$. Let us explain first the smooth case. Consider two smooth embeddings $\gamma_1, \gamma_2 \in \mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)$ and complete $\gamma'_1(0)$ and $\gamma'_2(0)$ to oriented framings $\langle \gamma'_1(0), f_1^2, f_1^3 \rangle$ and $\langle \gamma'_2(0), f_2^2, f_2^3 \rangle$ of $T_{\gamma_1(0)}\mathbb{S}^3$ and $T_{\gamma_2(0)}\mathbb{S}^3$, respectively. This determines two orientation preserving embeddings $F_{\gamma_j} : \mathbb{D}_\varepsilon^3(x, y, z) \rightarrow \mathbb{S}^3$, $j \in \{1, 2\}$; such that

- $F_{\gamma_j}(x, 0, 0) = \gamma_j(x)$, $(x, 0, 0) \in \mathbb{D}_\varepsilon^3$, and
- $F_{\gamma_j}(\mathbb{D}_\varepsilon^3) \cap \gamma_j(\mathbb{S}^1) = \{F_{\gamma_j}(x, 0, 0) : (x, 0, 0) \in \mathbb{D}_\varepsilon^3\} = \gamma_j(-\varepsilon, \varepsilon)$

for $\varepsilon > 0$ small enough. Perform the ambient connected sum $\mathbb{S}^3 \# \mathbb{S}^3$ along these disks by using an orientation reversing diffeomorphism to glue the boundaries of the disks that takes $F_{\gamma_1}(\partial \mathbb{D}_\varepsilon^3) \cap \gamma_1(-\varepsilon, \varepsilon)$ to $F_{\gamma_2}(\partial \mathbb{D}_\varepsilon^3) \cap \gamma_2(-\varepsilon, \varepsilon)$ coherently with the orientations of the embeddings, i.e. the image of $\gamma_1(\pm \varepsilon)$ is $\gamma_2(\mp \varepsilon)$. Notice that the framings $\langle f_j^2, f_j^3 \rangle$ naturally extends to a trivialization of the normal bundle of $\gamma_j|_{(-\varepsilon, \varepsilon)}$, $j \in \{1, 2\}$. Moreover, the image of the framing $\langle f_1^2, f_1^3 \rangle$ can be assumed to be $\langle f_2^2, -f_2^3 \rangle$. Clearly, the two embeddings $\gamma_1|_{\mathbb{S}^1 \setminus (-\varepsilon, \varepsilon)}$ and $\gamma_2|_{\mathbb{S}^1 \setminus (-\varepsilon, \varepsilon)}$ produce an embedding $\gamma_1 \# \gamma_2 : \mathbb{S}^1 \rightarrow \mathbb{S}^3 = \mathbb{S}^3 \# \mathbb{S}^3$. We define the *connected sum* of γ_1 and γ_2 as the embedding $\gamma_1 \# \gamma_2$.

In the Legendrian setting the construction is similar. Consider two Legendrian embeddings $\gamma_1, \gamma_2 \in \mathfrak{Leg}(\mathbb{S}^3, \xi_{\text{std}})$. There is a canonical framing of $\gamma_j^* \xi_{\text{std}}$, $j \in \{1, 2\}$, associated to the Legendrian condition, namely $\langle \gamma'_j, i\gamma'_j \rangle$ (the j here is just a subindex, it does not stand for quaternion j). This framing together with $i\gamma_j$ defines a trivialization of $\gamma_j^* T\mathbb{S}^3$. In particular, there is a natural contact framing over $T_{\gamma_j(0)}\mathbb{S}^3$ which defines a co-oriented contact embedding of a Darboux ball $F_{\gamma_j} : (\mathbb{D}_\varepsilon^3(x, y, z), \xi_{\text{std}}) \rightarrow (\mathbb{S}^3, \xi_{\text{std}})$, $j \in \{1, 2\}$, such that

- $F_{\gamma_j}(x, 0, 0) = \gamma_j(x)$, $(x, 0, 0) \in \mathbb{D}_\varepsilon^3$, and
- $F_{\gamma_j}(\mathbb{D}_\varepsilon^3) \cap \gamma_j(\mathbb{S}^1) = \{F_{\gamma_j}(x, 0, 0) : (x, 0, 0) \in \mathbb{D}_\varepsilon^3\} = \gamma_j(-\varepsilon, \varepsilon)$

for some $\varepsilon > 0$ small enough. Perform the contact connected sum along these Darboux balls by using an orientation reversing diffeomorphism of $\partial \mathbb{D}_\varepsilon^3$ taking the characteristic foliation to itself to obtain the connected sum Legendrian embedding $\gamma_1 \# \gamma_2$ on the contact manifold $(\mathbb{S}^3, \xi_{\text{std}}) \cong (\mathbb{S}^3, \xi_{\text{std}}) \#_{\gamma_1, \gamma_2} (\mathbb{S}^3, \xi_{\text{std}})$. Equivalently, the connected sum can be defined by working with *long Legendrian embeddings*, denote by $N \in \mathbb{S}^3$ the north pole of the sphere, by a contact isotopy we may assume that the intersection of γ_1 with the north hemisphere coincides with the Legendrian great circle $\langle N, jN \rangle^2$, i.e. with the Legendrian embedding $\gamma_{N, jN}(t) = (\cos t, \sin t) \in \mathbb{S}^3 \subseteq \mathbb{C}^2$, and

² We use quaternionic notation. Recall that $\xi_{\text{std}, p} = \langle jp, kp \rangle$.

that the intersection of γ_2 with the south hemisphere also coincides with $\langle N, jN \rangle$. It is clear that in this case we can glue both embeddings together to obtain the connected sum $\gamma_1 \# \gamma_2$. This point of view also works in the smooth case.

3.3 Connected sum at the π_1 -level.

Let us explain first the smooth case. Consider two loops $(\gamma_j^\theta, F_j^\theta)$, $\theta \in \mathbb{S}^1$ and $j \in \{1, 2\}$; where

- $\gamma_j^\theta \in \mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3)$,
- $F_j^\theta = \langle (\gamma_j^\theta)'(0), f_j^{\theta,2}, f_j^{\theta,3} \rangle$ is a framing of $T_{\gamma_j^\theta(0)}\mathbb{S}^3$.

Let us define the *connected sum loop* of γ_1^θ and γ_2^θ with framings F_1^θ and F_2^θ , denoted by $\gamma_1^\theta \#_{F_1^\theta, F_2^\theta} \gamma_2^\theta$, which depends on the choice of framings.

Associated with each loop we have a framed embedding, namely a pair of $s_j : \mathbb{S}^1 \rightarrow \mathbb{S}^3 \times \mathbb{S}^1, \theta \mapsto (\gamma_j^\theta(0), \theta)$, $j \in \{1, 2\}$ and the framing F_j^θ of the normal bundle of s_j . Use this framed embeddings to glue two copies of $\mathbb{S}^3 \times \mathbb{S}^1$ and obtain a 4-manifold $S = \mathbb{S}^3 \times \mathbb{S}^1 \#_{s_1, s_2} \mathbb{S}^3 \times \mathbb{S}^1$, i.e. S is obtained by performing at each slice $\mathbb{S}^3 \times \{\theta\}$ the connected sum in the π_0 -sense of the embeddings γ_1^θ and γ_2^θ . In particular, there is an \mathbb{S}^3 -bundle $p : S \rightarrow \mathbb{S}^1$, where the fiber over $\theta \in \mathbb{S}^1$ is \mathbb{S}_θ^3 together with the connected sum embedding $\gamma_1^\theta \# \gamma_2^\theta$. In order to obtain a loop of embeddings in \mathbb{S}^3 we need to trivialize this bundle. Note that this can always be done because $\pi_0(\text{Diff}^+(\mathbb{S}^3)) = 0$ but not in a canonical way since $\pi_1(\text{Diff}^+(\mathbb{S}^3)) \cong \mathbb{Z}_2$. However, there is a choice of framings for which a choice of trivialization can be made and we will restrict ourselves to the case where the loops of framings given by $(\gamma_j^\theta(0), F_j^\theta)$ are contractible (recall from the discussion in 3.1.1 that $\pi_1(\text{Fr}^+(\mathbb{S}^3)) \cong \pi_1(\mathbf{Emb}^+(\mathbb{D}^3, \mathbb{S}^3)) \cong \pi_1(\text{SO}(3)) = \mathbb{Z}_2$) and we are asking here that our loops are the trivial element in this \mathbb{Z}_2 . Therefore, we can understand the 4-manifold S as the boundary of an \mathbb{S}^3 -bundle over \mathbb{D}^2 and, in particular, there is a unique trivialization (up to homotopy) of the bundle S that extends over the disk. More precisely, if the loop $(\gamma_j^\theta(0), F_j^\theta)$, $j \in \{1, 2\}$, is trivial we can find, for each $j \in \{1, 2\}$, a disk $(p_j^{r,\theta}, \hat{F}_j^{r,\theta}) \in \text{Fr}^+(\mathbb{S}^3)$ such that $(p_j^{1,\theta}, \hat{F}_j^{1,\theta}) = (\gamma_j^\theta(0), F_j^\theta)$. By using these framings we can construct the desired \mathbb{S}^3 -bundle over \mathbb{D}^2 just by performing the \mathbb{D}^2 -family of connected sums determined by them. This process is unique up to homotopy because $\pi_2(\text{Fr}^+(\mathbb{S}^3)) \cong \pi_2(\mathbf{Emb}^+(\mathbb{D}^3, \mathbb{S}^3)) \cong \pi_2(\mathbb{S}^3) \oplus \pi_2(\text{SO}(3)) = 0$.

In conclusion, the connected sum of loops is canonically defined whenever we use trivial loops of framings. In this case the homotopy class of the loop $\gamma_1^\theta \#_{F_1^\theta, F_2^\theta} \gamma_2^\theta$ only depends on the homotopy class of $(\gamma_j^\theta, F_j^\theta)$, $j \in \{1, 2\}$. This is the content of the next

Lemma 3.3.1 *Let $(\gamma_j^\theta, F_j^\theta), (\hat{\gamma}_j^\theta, \hat{F}_j^\theta)$, $j \in \{1, 2\}$, be any pair of loops of smooth embeddings equipped with a trivial loop of framings at the base points. Assume that $(\gamma_j^\theta, F_j^\theta)$ is homotopic to $(\hat{\gamma}_j^\theta, \hat{F}_j^\theta)$, $j \in \{1, 2\}$, then, the loop $\gamma_1^\theta \#_{F_1^\theta, F_2^\theta} \gamma_2^\theta$ is homotopic to $\hat{\gamma}_1^\theta \#_{\hat{F}_1^\theta, \hat{F}_2^\theta} \hat{\gamma}_2^\theta$*

Proof. Let $(\gamma_1^{\theta,s}, F_1^{\theta,s})$, $s \in [0, 1]$, be any homotopy between $(\gamma_1^{\theta,0}, F_1^{\theta,0}) = (\gamma_1^\theta, F_1^\theta)$ and $(\gamma_1^{\theta,1}, F_1^{\theta,1}) = (\hat{\gamma}_1^\theta, \hat{F}_1^\theta)$. It is enough to check that $\gamma_1^{\theta,0} \#_{F_1^{\theta,0}, F_2^\theta} \gamma_2^\theta$ is homotopic to $\gamma_1^{\theta,1} \#_{F_1^{\theta,1}, F_2^\theta} \gamma_2^\theta$. Perform the connected sum of the loops $(\gamma_1^{\theta,s}, F_1^{\theta,s})$ and $(\gamma_2^\theta, F_2^\theta)$, $s \in [0, 1]$, to build an \mathbb{S}^3 -bundle $X \rightarrow \mathbb{S}^1 \times [0, 1]$ over the cylinder $\mathbb{S}^1 \times [0, 1]$. Since the loops of framings $(\gamma_1^{\theta,0}(0), F_1^{\theta,0})$ and $(\gamma_1^{\theta,1}(0), F_1^{\theta,1})$ are trivial we can extend this bundle to a \mathbb{S}^3 -bundle $\hat{X} \rightarrow \mathbb{S}^2$. Now the statement follows from the fact that $\pi_2(\mathbf{Emb}^+(\mathbb{D}^3, \mathbb{S}^3)) = 0$ and thus the bundle \hat{X} is the boundary of an \mathbb{S}^3 -bundle over the 3-disk \mathbb{D}^3 and hence canonically trivializable. \square

Consider now two loops of Legendrian embeddings $\gamma_1^\theta, \gamma_2^\theta \in \mathcal{LEG}(\mathbb{S}^3, \xi_{\text{std}})$, $\theta \in \mathbb{S}^1$. Recall that a Legendrian embedding γ is naturally a framed embedding by means of the framing $F_\gamma = \langle \gamma', i\gamma', i\gamma \rangle$. Thus, the choice of framings is canonical and we can denote the connected loop as $\gamma_1^\theta \# \gamma_2^\theta$. The construction is analogous to the smooth case. Use the natural contact framings to build a contact fiber bundle $p : S \rightarrow \mathbb{S}^1$ with fiber $p^{-1}(\theta) = (\mathbb{S}_\theta^3, \xi_\theta, \gamma_1^\theta \# \gamma_2^\theta) = (\mathbb{S}^3, \xi_{\text{std}}) \#_{\gamma_1^\theta, \gamma_2^\theta} (\mathbb{S}^3, \xi_{\text{std}})$. As in the smooth case we want to identify, in a canonical way, all the fibers with the standard contact structure over the 3–sphere. The bundle can be trivialized because $\pi_0(\text{Cont}(\mathbb{S}^3, \xi_{\text{std}})) = \pi_0(\text{U}(2)) = 0$ but the trivialization is not unique: $\pi_1(\text{Cont}(\mathbb{S}^3, \xi_{\text{std}})) = \pi_1(\text{U}(2)) \cong \mathbb{Z}$. As in the smooth case we restrict ourselves to the case when the contact framings associated to the loops are trivial. In this case we can understand the contact fiber bundle $p : S \rightarrow \mathbb{S}^1$ as the boundary of a contact bundle over \mathbb{D}^2 and, thus, trivialize it in a unique way. Moreover, the procedure is canonical because $\pi_2(\mathcal{CEmb}((\mathbb{D}^3, \xi_{\text{std}}), (\mathbb{S}^3, \xi_{\text{std}}))) = \pi_2(\mathbb{S}^3) \oplus \pi_2(\text{U}(1)) = 0$.

Let us explain geometrically under which conditions the loop of contact framings associated with a loop of Legendrian embeddings is trivial. Consider the natural map

$$\begin{aligned} F : \mathcal{LEG}(\mathbb{S}^3, \xi_{\text{std}}) &\longrightarrow \text{CFr}(\mathbb{S}^3, \xi_{\text{std}}) \\ \gamma &\mapsto F(\gamma) = \langle \gamma'(0), i\gamma'(0), i\gamma(0) \rangle. \end{aligned} \quad (3.5)$$

Recall that the space of Darboux balls in $(\mathbb{S}^3, \xi_{\text{std}})$ is homotopy equivalent to $\mathbb{S}^3 \times \text{U}(1)$ and the homotopy equivalence is given by associating to each contact framing $\langle e_1, e_2, e_3 \rangle$ over a point p in \mathbb{S}^3 the pair $(p, e_1) \in \mathbb{S}^3 \times \text{U}(1)$ (see Lemma 3.1.6). The induced map at π_1 –level is given by

$$\begin{aligned} \pi_1(F) : \pi_1(\mathcal{LEG}(\mathbb{S}^3, \xi_{\text{std}})) &\longrightarrow \pi_1(\text{CFr}(\mathbb{S}^3, \xi_{\text{std}})) \cong \mathbb{Z} \\ \gamma^\theta &\mapsto \text{Rot}_{\pi_1}(\gamma^\theta). \end{aligned} \quad (3.6)$$

Lemma 3.3.2 *Let $\gamma_j^\theta \in \mathcal{LEG}(\mathbb{S}^3, \xi_{\text{std}})$, $j \in \{1, 2\}$, be any loop of Legendrian embeddings with rotation number zero. Then, the Legendrian connected sum loop, denoted by $\gamma_1^\theta \# \gamma_2^\theta$, is a well-defined loop of Legendrian embeddings in $(\mathbb{S}^3, \xi_{\text{std}})$. Moreover, if $\hat{\gamma}_j^\theta$, $j \in \{1, 2\}$, is homotopic to γ_j^θ then $\hat{\gamma}_1^\theta \# \hat{\gamma}_2^\theta$ is homotopic to $\gamma_1^\theta \# \gamma_2^\theta$.*

Proof. The existence part has been proven in the discussion above. The statement about the homotopy invariance follows in the same way as in the smooth case (Lemma 3.3.1) because $\pi_2(\mathcal{CEmb}((\mathbb{D}^3, \xi_{\text{std}}), (\mathbb{S}^3, \xi_{\text{std}}))) = 0$.

Remark 3.3.3 *The discussion above (the smooth and Legendrian cases) can be generalized for multi-parametric families and for higher dimensional spheres.*

Let us explain this operation from the point of view of *long Legendrian embeddings*. In order to do this consider the homotopy equivalence

$$\begin{aligned} \Phi : \mathcal{LEG}(\mathbb{S}^3, \xi_{\text{std}}) &\longrightarrow \text{U}(2) \times \mathcal{LEG}_{N,jN}(\mathbb{S}^3, \xi_{\text{std}}) \\ \gamma &\mapsto (A_\gamma, A_\gamma^{-1}\gamma), \end{aligned}$$

The rotation number of a loop γ^θ is just the homotopy class of $A_{\gamma^\theta} \in \pi_1(\text{U}(2)) \cong \mathbb{Z}$. I.e. Lemma 3.3.2 just says that the connected sum operation for loops is well-defined for loops of long Legendrian embeddings. In fact, to perform the connected sum of two loops γ_1^θ and γ_2^θ of long Legendrian embeddings we can proceed as follows: let $\mathbb{S}_N^3 = \{(z_1, z_2) = (x_1 + iy_1, x_2 + iy_2) \in \mathbb{S}^3 : x_1 \geq 0\}$ be the north hemisphere of \mathbb{S}^3 and \mathbb{S}_S^3 the south hemisphere, by using a contact isotopy we may assume that:

- $\gamma_1^\theta(t) = (\cos t, \sin t)$ for $t \in [-\frac{\pi}{2}, \frac{\pi}{2}] \subseteq \mathbb{S}^1$,
- $\gamma_2^\theta(t) = (-\cos t, -\sin t)$ for $t \in [-\frac{\pi}{2}, \frac{\pi}{2}] \subseteq \mathbb{S}^1$,

- $\gamma_1^\theta(\frac{\pi}{2}, \frac{3\pi}{2}) \cap \mathbb{S}_N^3 = \emptyset$,
- $\gamma_2^\theta(\frac{\pi}{2}, \frac{3\pi}{2}) \cap \mathbb{S}_S^3 = \emptyset$.

Since $(\gamma_1^\theta(\pm\frac{\pi}{2}), (\gamma_1^\theta)'(\pm\frac{\pi}{2})) = (\gamma_2^\theta(\mp\frac{\pi}{2}), (\gamma_2^\theta)'(\mp\frac{\pi}{2}))$ this two loops can be glued to produce the loop of Legendrian embeddings

$$\gamma_1^\theta \widehat{\#} \gamma_2^\theta(t) = \begin{cases} \gamma_1^\theta(t + \frac{\pi}{2}) & \text{if } 0 \leq t \leq \pi, \\ \gamma_2^\theta(t - \frac{\pi}{2}) & \text{if } \pi \leq t \leq 2\pi. \end{cases} \quad (3.7)$$

This definition of connected sum (in the rotation zero case) is the same, up to homotopy, to the previous one. Therefore, from now on we denote this operation also with the symbol $\#$ instead of $\widehat{\#}$.

The aforementioned homotopy equivalence allows us to define the connected sum of loops of Legendrians with any rotation number:

Definition 3.3.4 Let $\gamma_j^\theta \in \mathfrak{L}\mathfrak{eg}(\mathbb{S}^3, \xi_{\text{std}})$, $j \in \{1, 2\}$, be any loop of Legendrian embeddings. The connected sum loop of γ_1^θ and γ_2^θ is the loop

$$\gamma_1^\theta \widetilde{\#} \gamma_2^\theta = A_{\gamma_1^\theta} A_{\gamma_2^\theta} ((A_{\gamma_1^\theta})^{-1} \gamma_j^\theta \# (A_{\gamma_2^\theta})^{-1} \gamma_j^\theta).$$

Remark 3.3.5 It follows from Lemma 3.3.2 that in the case of rotation number zero the loops $\gamma_1^\theta \# \gamma_2^\theta$ and $\gamma_1^\theta \widetilde{\#} \gamma_2^\theta$ are homotopic. Thus, we will denote both of them by $\gamma_1^\theta \# \gamma_2^\theta$.

From this point of view we can check the commutativity, up to homotopy, of the connected sum for loops of Legendrian embeddings:

Lemma 3.3.6 Let $\gamma_j^\theta \in \mathfrak{L}\mathfrak{eg}(\mathbb{S}^3, \xi_{\text{std}})$, $j \in \{1, 2\}$, be any loop of Legendrian embeddings. The loops $\gamma_1^\theta \# \gamma_2^\theta$ and $\gamma_2^\theta \# \gamma_1^\theta$ are homotopic.

Proof. It is enough to check the statement in the case that $\text{Rot}_{\pi_1}(\gamma_1^\theta) = \text{Rot}_{\pi_1}(\gamma_2^\theta) = 0$. Consider contact isotopy given by the matrices $A_s = \begin{pmatrix} \cos \pi s & -\sin \pi s \\ \sin \pi s & \cos \pi s \end{pmatrix} \in \mathcal{U}(2)$, $s \in [0, 1]$. Clearly, the loop $A_1(\gamma_1^\theta \# \gamma_2^\theta)$ is homotopic to the loop $\gamma_2^\theta \# \gamma_1^\theta$, see Equation (3.7).

3.3.1 Connected sum of loops of Legendrian embeddings in $(\mathbb{R}^3, \xi_{\text{std}})$

Let us recall the connected sum operation at π_0 -level explained in [43]. Let γ_1 and γ_2 be two Legendrian embeddings in $(\mathbb{R}^3, \xi_{\text{std}}) \subseteq (\mathbb{S}^3, \xi_{\text{std}})$. Perform the connected sum of both as Legendrian embeddings in the 3-sphere. To obtain a Legendrian embedding in $(\mathbb{R}^3, \xi_{\text{std}})$ we just need to eliminate a point $\infty \in \mathbb{S}^3 \setminus \gamma_1 \# \gamma_2(\mathbb{S}^1)$. Since the complement of a knot in \mathbb{S}^3 is path-connected the construction is independent, up to Legendrian isotopy, of the choice of infinity point in the 3-sphere.

Consider now two loops, γ_1^θ and γ_2^θ , of Legendrian embeddings in $(\mathbb{R}^3, \xi_{\text{std}})$. Understand both loops as Legendrian embeddings in $(\mathbb{S}^3, \xi_{\text{std}})$ and perform their connected sum to obtain a loop $\gamma_1^\theta \# \gamma_2^\theta$ of Legendrians in the 3-sphere. To go back to $(\mathbb{R}^3, \xi_{\text{std}})$ we need to choose a loop $p_\theta \in \mathbb{S}^3$ of points in the complement of each embedding to eliminate. In other words, we need to choose a section of the fiber bundle $\{(p, \theta) \in \mathbb{S}^3 \times \mathbb{S}^1 : p \in \mathbb{S}^3 \setminus \gamma_1^\theta \# \gamma_2^\theta(\mathbb{S}^1)\} \rightarrow \mathbb{S}^1$, $(p, \theta) \mapsto \theta$. This choice is not unique: the obstruction to homotope two sections is measured by the knot group of $\gamma_1^\theta \# \gamma_2^\theta(\mathbb{S}^1)$ (see Corollary 3.1.2).

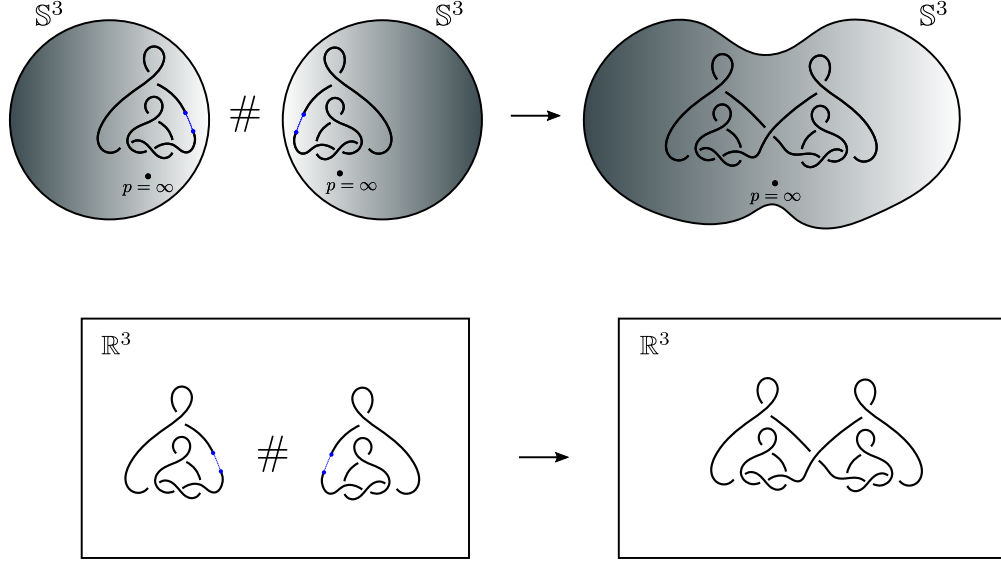


Figure 3.1: Connected sum at the π_0 -level for two legendrian trefoils in \mathbb{S}^3 . Note that by eliminating a point $p = \infty \in \mathbb{S}^3$ which does not belong to any of the knots we can regard the constructions, via the stereographic projection, as a connected sum of knots in \mathbb{R}^3 . This is depicted in the second row, where suitable projections have been taken.

We denote the connected sum loop associated to the choice p_θ by $\gamma_1^\theta \#_{p_\theta} \gamma_2^\theta \in \mathcal{Leg}(\mathbb{R}^3, \xi_{\text{std}})$.

A geometric representation of the π_1 -connected sum: the fly and the elephant.

For smooth embedding spaces satisfying $\pi_2(\mathbf{Emb}(\mathbb{S}^1, \mathbb{R}^3)) = 0$, we provide a geometric representation of π_1 -connected sums of loops which are smoothly trivial as loops in \mathbb{R}^3 (and thus in \mathbb{S}^3). Note that, since $\pi_2(\mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3)) = 0$, then the obstruction to homotoping two sections of the fiber bundle

$$\{(p, \theta) \in \mathbb{S}^3 \times \mathbb{S}^1 : p \in \mathbb{S}^3 \setminus \gamma_1^\theta \# \gamma_2^\theta(\mathbb{S}^1)\} \rightarrow \mathbb{S}^1, (p, \theta) \mapsto \theta$$

is measured by $\pi_1(\mathbb{S}^3 \setminus K) / \pi_2(\mathbf{Emb}(\mathbb{S}^1, \mathbb{S}^3)) = \pi_1(\mathbb{S}^3 \setminus K)$.

Consider two Legendrian embeddings $\gamma_1, \gamma_2 \in \mathcal{Leg}(\mathbb{R}^3)$ and two smoothly trivial loops of legendrians $\gamma_1^\theta, \gamma_2^\theta \in \pi_1(\mathcal{Leg}(\mathbb{R}^3))$ based on them, respectively, with rotation number $\text{Rot}_{\pi_1}(\gamma_1^\theta) = \text{Rot}_{\pi_1}(\gamma_2^\theta) = 0$. We then produce a loop of legendrians in the connected component of $\gamma_1 \# \gamma_2$ described by the following construction (which we call the fly and the elephant construction). Note that, as explained above, we will think the construction in \mathbb{S}^3 (where the connected-sum operation is defined) and we will come back to \mathbb{R}^3 afterwards.

Step 1

For small enough $\varepsilon > 0$, consider the small arcs $\gamma_1(-\varepsilon, \varepsilon)$ and $\gamma_2(-\varepsilon, \varepsilon)$ and perform the π_0 -connected sum $\gamma_1 \# \gamma_2$ of both Legendrian embeddings.

Step 2

Shrink the component of γ_2 in this connected sum embedding until it becomes arbitrarily small. This can be schematically understood as in Figure 3.7, where we have the connected sum of a Legendrian trefoil with a Legendrian p, q -torus knot (the latter can be understood as a shrunk embedding and thus represented as a small box).

Now regard the embedding $\gamma_1 \# \gamma_2$ (where the component of γ_2 has been shrunk) as γ_1 with a tiny box (encoding the topology of γ_2) attached to it (more precisely, taking the place of the segment $\gamma_1(-\varepsilon, \varepsilon)$). Perform now the loop γ_1^θ to this connected sum (where the box just moves rigidly with the embedding). Note that since $\text{Rot}_{\pi_1}(\gamma_1^\theta) = 0$, this box can be thought (precomposing with a

loop of matrices in $U(1)$) as constant (not moving) all along the loop at the homotopy level. See Figure 3.9 for visualizing Step 2 with a particular example.

Step 3.

Once the loop in Step 2 has been performed and, thus, the embedding $\gamma_1 \# \gamma_2$ is in its original position (at $\theta = 0$), we enlarge the component of γ_2 until it reaches its original size. We now shrink the component of γ_1 in the embedding and perform the analogous construction from Step 2 but replacing the roles of the γ_1 - and γ_2 -components. In other words, we perform now the loop γ_2^θ to this connected sum (where the box encoding the topology of γ_1 just moves rigidly with the embedding). Finally, we enlarge γ_2 until it reaches its original size. This completes the construction.

Remark 3.3.7 *The terminology fly and elephant construction has been chosen since when each component in the connected sum has been shrunk and moves rigidly along the loop performed by the other component, this resembles a fly posed in a big moving elephant.*

Note that this construction is a representation of the π_1 -connected sum for two loops of Legendrians in \mathbb{S}^3 where the loop of points p_θ in the complement of each embedding is constant. By means of the following diagram

$$\begin{array}{ccccccc}
0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) & \longrightarrow & \pi_1(\mathfrak{Leg}(\mathbb{R}^3, \xi_{\text{std}})) & \longrightarrow & \pi_1(\mathfrak{Leg}(\mathbb{S}^3, \xi_{\text{std}})) \longrightarrow 0 \\
& & \downarrow & & \downarrow & & \downarrow \\
0 & \longrightarrow & \pi_1(\mathbb{S}^3 \setminus K) / \pi_2(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) & \longrightarrow & \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3)) & \longrightarrow & \pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{S}^3)) \longrightarrow 0
\end{array}$$

this is the case since both loops are smoothly trivial in \mathbb{R}^3 and \mathbb{S}^3 . We conclude that the obstruction (in $\pi_1(\mathfrak{Emb}(\mathbb{S}^1, \mathbb{R}^3))$) to going back to \mathbb{R}^3 is zero.

3.4 Invariants for one-parameter families of Legendrians.

3.4.1 Chekanov–Eliashberg Differential Graded Algebra

The presentation in this Subsection is based on [71], from where we learnt about invariants for one-parameter families of Legendrians.

Y. Chekanov proved [26] that the homology $H(K)$ associated to a generic Legendrian knot $K \in \widehat{\mathfrak{Leg}}(\mathbb{R}^3, \xi_{\text{std}})$ is invariant under Legendrian isotopies. We first introduce the differential graded algebra (DGA) associated to a Legendrian knot. For our purposes it suffices to work with \mathbb{Z}_2 -graduated algebra.

Fix an oriented Legendrian K . For simplicity we assume that $\text{Rot}(K) = 0$ which is the case discussed in this chapter, which is assumed as an hypothesis in the rest of the Section, unless otherwise stated. We define its associated DGA \mathcal{A}_K over \mathbb{Z}_2 as the algebra whose generators are the crossings $C = \{c_1, \dots, c_k\}$ of its Lagrangian projection K_L . Let us define the grading. Following [71], we say that the capping path of a crossing $c \in K_L$ is the (oriented) path in K_L starting at the undercrossing and whose endpoint is the uppercrossing of c . Moreover, assume that all the crossings take place at right angles (in the Lagrangian projection) and, therefore, the rotation number of the capping path r_c takes values $(2k + 1)/4, k \in \mathbb{Z}$. Define the grading of c as

$$|c| = -2r_c - \frac{1}{2} \in \mathbb{Z}$$

and extend it in the usual way to the whole algebra \mathcal{A}_K . Observe that the condition on the rotation number of the Legendrian implies that the grading is independent of the choice of orientation and therefore it is well defined. Define the Reeb sign associated to a quadrant at a crossing by the following convention: fix the ordered basis given by the tangent vector of the Legendrian at the upper crossing and the tangent vector at the under crossing. We say that the Reeb sign is positive if the orientation of this basis coincides with the standard one in \mathbb{R}^2 and it is negative otherwise.

Fix crossings p, c_1, \dots, c_n ; denote by $\mathcal{P}(p; c_1, \dots, c_n)$ the set of (non parametrized) $(n+1)$ -polygons P in \mathbb{R}^2 satisfying the following properties:

- (i) P is immersed everywhere except at the ordered set of points $\mathcal{Q} = \{p, c_1, \dots, c_n\}$.
- (ii) $\partial P \subset K_L$.
- (iii) The non immersed curve ∂P fails to be immersed at the sequential set of points \mathcal{Q} . Moreover, it follows a positive (resp. negative) quadrant for p (resp. for c_j).

Do note that c_j and $c_{j'}$, for $j \neq j'$, may be the same crossing. Moreover, observe that we consider polygons, instead of their boundaries because of two facts. First, we want to avoid two different smooth branches of the curve ∂P that go through a crossing without turning. Second, we want make sure that $\infty \in \mathbb{S}^2 \simeq \mathbb{R}^2 \cup \{\infty\} \supseteq \mathbb{R}^2$ is not contained in P .

Now, define the differential of a crossing p via the formula:

$$\partial(p) = \sum_{n \in \mathbb{N}, (p, c_1, \dots, c_n) \in C^n} \# \mathcal{P}(p; c_1, \dots, c_n) c_1 \cdots c_n$$

and extend it to the rest of the words by the (graded) Leibniz rule.




Figure 3.2: Example of a Legendrian connected sum of two trefoil knots with the described choice of signs for the crossings. In yellow, an example of one of the polygons contributing to the differential of the uppermost crossing in the diagram.

Remark 3.4.1 Define the height of a crossing c , which is a strict positive number $h(c) > 0$ to be the difference of the z coordinates of the two points on c when lifting the Lagrangian projection to \mathbb{R}^3 , ie the length of the Reeb chord joining them.

If $P \in \mathcal{P}(p; c_1, \dots, c_n)$, fix a parametrization of P namely $\tilde{P} : \mathbb{D}^2 \rightarrow \mathbb{R}^2$, then by Stokes' Theorem,

$$h(p) - \sum_{i=1}^n h(c_i) = \int_{\mathbb{D}^2} \tilde{P}^*(dx \wedge dy) > 0.$$

This formula does not depend on the parametrization.

Observe that whenever we fix a crossing p all the words in $\partial(p)$ are made of crossings with strictly smaller height than p .

Remark 3.4.2 The differential ∂ is well defined since this sum has a finite number of non-vanishing elements, because of the previous property. Moreover, it has degree -1 and its square is zero. The algebra \mathcal{A}_K equipped with the boundary operator ∂ is known as Chekanov–Eliashberg DGA of the Legendrian knot K . This DGA gives rise to the Chekanov–Eliashberg Legendrian Contact homology $H_*(K) = \text{Ker}(\delta)/\text{Im}(\delta)$ of a Legendrian K which depends just on the (Legendrian) isotopy class of the Legendrian [26].

Let K_t , $t \in [0, 1]$ be a generic path of Legendrians and consider the Lagrangian projection of the path $\pi_L(K_t), t \in [0, 1]$. Then, four essentially different Legendrian Reidemeister moves ($\mathcal{R} - \text{II}$, $\mathcal{R} - \text{II}^{-1}$, $\mathcal{R} - \text{III}_a$, $\mathcal{R} - \text{III}_b$) take place in this projection at finitely many times $\{t_0, \dots, t_k\}$ (see Figure 3.3), and no other bifurcation takes place for different values of the parameter t .

Denote by $(\mathcal{A}_K, \partial)$ the differential algebra associated to the Legendrian before one of the Reidemeister moves takes place and by $(\mathcal{A}'_K, \partial')$ the one associated to the embedding just after such a Reidemeister move. Then, the diagram is only locally affected where the Reidemeister move has taken place and, therefore, there is an obvious mapping $p \mapsto p'$ for all generators $p \in \mathcal{A}$ of the algebra not affected by the move. For the ones affected by the Reidemeister move we have to name the old and the new corners in order to setup a correspondence in each of the eight Reidemeister moves. There are four (see Figure 3.3).

The notation is as follows. In any of the four $\mathcal{R} - \text{II}$ moves the pair of points that appear in the intersecting branches have opposite signs. So, we just label x to the positive one and y to the other one. For the four $\mathcal{R} - \text{III}$ moves we denote by z the crossing that remains fixed under the movement and by x and y the two crossings of the moving branch with respect to the branches that define the corner z . Moreover, x is chosen to be the negative corner and y the positive one.

We will carefully discuss the holonomy associated to each of the four drawn Reidemeister moves.




Figure 3.3: Possible different Reidemeister moves taking place in the Lagrangian projection of a generic path of Legendrian embeddings.

Theorem 3.1 (Chekanov, [26]). Let K_t , $t \in [0, 1]$ be a generic path of Legendrians. Denote by $\{t_0, \dots, t_k\}$ the times where some Legendrian Reidemeister move takes place in the Lagrangian projection of the path. The following algebra homomorphisms, called holonomy maps, $g : \mathcal{A}_{K_{t_i-\varepsilon}} \rightarrow \mathcal{A}_{K_{t_i+\varepsilon}}$ (described below for generators depending on the different Reidemeister move taking place at t_i) induce isomorphisms at the homology-level.

Name the generators of \mathcal{A}_K as follows ordered by their heights:

$$h(x_q) \geq h(x_{q-1}) \geq \dots \geq h(x) > h(y) \geq h(y_1) \geq h(y_2) \geq \dots h(y_\ell)$$

and let $\partial(x) = y + w$.

- $\mathcal{R} - II$

The holonomy map for the $\mathcal{R} - II$ move maps trivially every y_j ; i.e. $\varphi(y'_j) = y_j$ and is defined recursively for the other generators x_1, \dots, x_q starting by defining it for x_1 . Without loss of generality express $\partial(x_1) = \sum Y_1 y Y_2 y \cdots Y_m y X$, where $Y_1, \dots, Y_m \in T(y_1, \dots, y_\ell)$ are monomials and where $X \in T(x, y, y_1, \dots, y_\ell)$ is a monomial where every y is preceded by an x . Then, the holonomy map acts on x'_1 as follows:

$$\begin{aligned} \varphi(x'_1) = & x_1 + \sum Y_1 x Y_2 y Y_3 y Y_4 y \cdots Y_m y X \\ & + Y_1 w Y_2 x Y_3 y Y_4 y \cdots Y_m y X \\ & + Y_1 w Y_2 w Y_3 x Y_4 y \cdots Y_m y X \\ & + Y_1 w Y_2 w Y_3 w Y_4 x \cdots Y_m y X \\ & + \cdots \\ & + Y_1 w Y_2 w Y_3 w Y_4 w \cdots Y_m x X \end{aligned}$$

and the holonomy map acts recursively on each $x_j, j \geq 1$ as follows:

Write $\partial(x_j) = \sum Y_1 y Y_2 y Y_3 y Y_4 y \cdots Y_k y X$, where $Y_1, \dots, Y_k \in T(y_1, \dots, y_\ell, x_1, \dots, x_{j-1})$ and where in the factor X every y is preceded by an x . Then,

$$\begin{aligned} \varphi(x'_j) = & x_j + \sum Z_1 x Y_2 y Y_3 y Y_4 y \cdots Y_m y X \\ & + Z_1 w Z_2 x Y_3 y Y_4 y \cdots Y_m y X \\ & + Z_1 w Z_2 w Z_3 x Y_4 y \cdots Y_m y X \\ & + Z_1 w Z_2 w Z_3 w Z_4 x \cdots Y_m y X \\ & + \cdots \\ & + Z_1 w Z_2 w Z_3 w Z_4 w \cdots Y_m x X \end{aligned}$$

where each Z_i is obtained from Y_i just by substituting each x_1, \dots, x_{j-1} by the corresponding $\varphi(x_1), \dots, \varphi(x_{j-1})$ obtained in previous steps of the iterative process.

- $\mathcal{R} - II^{-1}$

$x \mapsto 0, y \mapsto w'$ and $p \mapsto p'$ for any other $p \in \mathcal{A}_k$.

- $\mathcal{R} - III_a$.

$p \mapsto p'$ for any $p \in \mathcal{A}_k$.

- $\mathcal{R} - III_b$.

$x \mapsto x' + z'y'$ and $p \mapsto p'$ for any other generator $p \in \mathcal{A}_K$, where p corresponds to the same crossing p after the move.

3.4.2 Invariants for one-parameter families of Legendrians.

Theorem 3.2 ([71]). Let $K \in \widehat{\mathfrak{Leg}}(\mathbb{R}^3, \xi_{\text{std}})$ be a generic Legendrian, then the continuation map on the Chekanov-Eliashberg Legendrian Homology (see Theorem 3.1) defines a multiplicative group homomorphism

$$\pi_1(\widehat{\mathfrak{Leg}}(\mathbb{R}^3, \xi_{\text{std}}), K) \rightarrow \text{Aut}(H_*(K)).$$

We call this homomorphism defined for loops of Legendrians the monodromy invariant of the loop.

The following proposition is useful in order to compute the monodromy of a loop in practice:

Proposition 3.4.3 ([71], Prop. 3.5.) *Let $p \in \mathcal{A}_K$ be a generator such that $\partial(p) = 0$, then the holonomy $\varphi : \mathcal{A}_K \rightarrow \mathcal{A}'_K$ corresponding to a $\mathcal{R} - II$ move acts as the identity on p ; i.e. $\varphi(p) = p$.*

3.5 Legendrian (p, q) -tangles and associated words.

A diagrammatic realization of a (p, q) -torus knot in the Lagrangian projection consists on the following construction. Take the usual Legendrian (p, q) -braid and close it by joining the i -th strand at the left side of the braid with the i -th strand at the right side of the braid via a closed unknotted curve with a kink (for every $1 \leq i \leq q$). See Figure 3.4 (A). The enclosed 0-area condition must be satisfied when performing this diagrammatic realization but for practical reasons we will not take this into consideration in our Figures.

Definition 3.5.1 *We define a Legendrian (p, q) -tangle as a tangle consisting on opening a Legendrian (p, q) -torus knot at a point Q in the curve joining the points in the q -th position both at left and right side of the Legendrian (p, q) -braid in the knot.*

See Figure 3.4 for an example of a $(4, 3)$ -torus knot (A) and its corresponding $(4, 3)$ -tangle (B).

Remark 3.5.2 *Any choice of point Q gives raise, up to Legendrian homotopy, to the same tangle diagram and, therefore, we can speak about the (p, q) -tangle with no ambiguity. Observe that if you do not assume that the point Q is in the top most strand in the Lagrangian diagram this is not true by similar reasons to the ones explained in Subsection 3.3.1.*




Figure 3.4

Proposition 3.5.3 *The Legendrian connected sum of a (p, q) -torus knot and a (p', q') -torus knot admits the diagrammatic realization in the Lagrangian projection consisting on the concatenation of the (p, q) -tangle and the (p', q') -tangle together with an unknotted curve with a kink joining them (see third step of Figure 3.5).*

Proof. The first arrow in Figure 3.5 corresponds to the canonical Legendrian connected sum of the Legendrian (p, q) -torus knot and the Legendrian (p', q') -torus knot. From this diagrammatic realization we can achieve the desired realization by performing the II^{-1} -Reidemeister move described by the second arrow. The area of the 2-gon in dark grey can be made smaller than the sum of the areas of adjacent regions (in light grey) at its vertices (for example, by taking the (p', q') -torus knot very small). Since this is the condition in order for this Reidemeister II^{-1} -move to be valid in the Legendrian category (see [71], Theorem 4.1), the result follows. \square




Figure 3.5: Legendrian canonical connected sum of a (p, q) -torus knot and a (p', q') -torus knot followed by a legit Reidemeister II^{-1} move. The area of the involved 2-gon (in dark grey) can be made smaller than the sum of the areas of adjacent regions (in light grey) at its vertices.

Given a (p, q) -tangle $\mathcal{T}_{p,q}$ with endpoints p_0 and p_1 , consider the crossings $b_1, \dots, b_n \in C$, where C denotes the set of crossings in the Legendrian (p, q) -braid inside the tangle, i.e. crossings of degree 0 in the DGA. Denote by \hat{C} the set of ordered subsets of C . For any $\{b_{j_1}, \dots, b_{j_m}\} \in \hat{C}$ denote by $\mathcal{M}_{p,q}(b_{j_1}, \dots, b_{j_m})$ the set of smooth curves $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ up to reparametrization satisfying the following properties:

- (i) $\gamma(0) = p_0$ and $\gamma(1) = p_1$.
- (ii) γ is immersed everywhere except at the points $\gamma(t_1) = b_{j_1}, \dots, \gamma(t_m) = b_{j_m}$; $0 < t_1 < \dots < t_m < 1$.
- (iii) $\gamma([0, 1]) \subset \mathcal{T}_{p,q}$.
- (iv) Near a point b_j , $j \in \{j_1, \dots, j_n\}$, where the curve is not immersed the trace of the curve follows a negative quadrant.

Definition 3.5.4 Define the word $W_{p,q}$ associated to a (p, q) -tangle as the formal sum:

$$W_{p,q} = \sum_{\{b_{j_1}, \dots, b_{j_m}\} \in \hat{C}} \# \mathcal{M}_{p,q}(b_{j_1}, \dots, b_{j_m}) b_{j_1} \cdots b_{j_m}$$




Figure 3.6: b_5 is one of the monomials in the word $W_{4,3}$ associated to a $(4, 3)$ -torus tangle; i.e. $W_{4,3} = b_5 + \dots$.

Let us elaborate a bit on the differential graded algebras of a particular kind of torus knots. Consider the positive Legendrian torus knots $K_{n,2}$. T. Kálmán describes the differentials of the 0-index generators b_i (corresponding to the crossings in the braid from the diagrammatic realization explained above).

Definition 3.5.5 The braid β corresponding to $K_{n,2}$ has an associated matrix called the **path matrix** defined as:

$$B_\beta := \begin{pmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{pmatrix} = \begin{pmatrix} b_1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} b_2 & 1 \\ 1 & 0 \end{pmatrix} \cdots \begin{pmatrix} b_n & 1 \\ 1 & 0 \end{pmatrix}$$

Remark 3.5.6 In fact, Kálmán defines such a matrix for any braid but since we will focus on $(n,2)$ -torus knots we will not introduce it in such generality.

Lemma 3.5.7 (Theorem 6.7 in [71]) For the legendrian torus knot $K_{n,2}$ with a_1, a_2 the index-1 generators and b_1, \dots, b_n index-0 generators, we have:

$$\partial(b_i) = 0, \quad \partial(a_1) = 1 + B_{1,1}, \quad \partial(a_2) = 1 + B_{2,2} + B_{2,1}B_{1,2}$$

where $B_{i,j}$ are the entries of the path matrix associated to $K_{n,2}$.

Example 3.5.8 (Example 3.3 in [72]) For the positive right-handed legendrian trefoil $K_{2,3}$ we have

$$B_\beta = \begin{pmatrix} b_1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} b_2 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} b_3 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} b_1 + b_3 + b_1b_2b_3 & 1 + b_1b_2 \\ 1 + b_2b_3 & b_2 \end{pmatrix}.$$

We obtain then that $\partial(a_1) = 1 + b_1 + b_3 + b_1b_2b_3$ and $\partial(a_2) = 1 + 1 + b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2$.

Definition 3.5.9 We call the **length** of a polynomial P , and we denote it by $\ell(P)$, to the number of words in P .

Lemma 3.5.10 Let $K_{n,2}$ be a torus knot with index-1 generators a_1, a_2 with a diagrammatic representation as described in the beginning of Subsection 3.5. Then,

$$\ell(B_{1,1}) = F_{n+1}, \quad \ell(B_{1,2}) = \ell(B_{2,1}) = F_n, \quad \ell(B_{2,2}) = F_{n-1},$$

where F_i denotes the i -th **Fibonacci number**.

Proof. The proof follows by induction.

- **Base case: $n=3$.** The claim is true for $K_{2,3}$ (see example 3.5.8) since

$$\ell(B_{1,1}) = F_4 = 3, \quad \ell(B_{1,2}) = \ell(B_{2,1}) = F_3 = 2, \quad \ell(B_{2,2}) = F_1 = 1.$$

- **Inductive step: $n=k+1$.** We assume the claim true for $n = k$ and we will prove it for $n = k+1$.

Use the following notation:

$$B_\beta^{2,k} = \begin{pmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{pmatrix}, \quad B_\beta^{2,k+1} = \begin{pmatrix} \tilde{B}_{1,1} & \tilde{B}_{1,2} \\ \tilde{B}_{2,1} & \tilde{B}_{2,2} \end{pmatrix}.$$

We then have

$$B_\beta^{2,k+1} = B_\beta^{2,k} \cdot \begin{pmatrix} b_{k+1} & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{pmatrix} \cdot \begin{pmatrix} b_{k+1} & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} B_{1,1} \cdot b_{k+1} + B_{1,2} & B_{1,1} \\ B_{2,1} \cdot b_{k+1} + B_{2,2} & B_{2,1} \end{pmatrix}.$$

Therefore, it follows that

$$- \ell(\tilde{B}_{1,1}) = F_{k+1} + F_k = F_{k+2},$$

- $\ell(\tilde{B}_{1,2}) = F_{k+1}$,
- $\ell(\tilde{B}_{2,1}) = F_k + F_{k-1} = F_{k+1}$,
- $\ell(\tilde{B}_{2,2}) = F_k$.

This proves the inductive step, thus yielding the claim.

We now introduce a class of knots that will play a central role along the chapter. These are the so called **even ∂ -class**.

Definition 3.5.11 *We say that a Legendrian knot is of **even ∂ -class** if it admits a diagrammatic realization in the Lagrangian projection where the differential of every index-1 crossing a_i is an expression $\partial(a_i)$ containing an even number of words.*

The following proposition shows that there are infinitely many Legendrian torus knots of even ∂ -class and thus the study of this class of knots is meaningful.

Proposition 3.5.12 *Legendrian positive $(n, 2)$ torus knots are of even ∂ -class if and only if $n \not\equiv 1 \pmod{3}$.*

Proof. By Lemma 3.5.7 and Lemma 3.5.10 we know that $\ell(\partial(a_1))$ has the same parity as the number $A_1 := F_{n+1} - 1$ while $\ell(\partial(a_2))$ has the same parity as the number $A_2 := F_{n-1} + F_n^2 - 1$.

The i -th Fibonacci number F_i is even if and only if $i \equiv 0 \pmod{3}$. Therefore, if $n \equiv 1 \pmod{3}$ or $n \equiv 2 \pmod{3}$, both A_1 and A_2 are even numbers, while if $n \equiv 0 \pmod{3}$ then this is not the case, thus yielding the claim.

3.6 New examples of loops with non-trivial monodromy.

Definition 3.6.1 *Consider a certain knot diagram with set of crossings C . Then for any formal expression*

$$\phi = \sum_{k \in \mathbb{N}, (b_1, \dots, b_k) \in C^k} \lambda(b_1, \dots, b_k) b_1 \cdots b_k$$

Denote by $\max_{b_j}(\phi)$ the maximum number of b_j letters appearing in a single word of the expression ϕ (among all the words involved in ϕ).

Denote by $h_{b_j}^r(\phi)$ the number of words in ϕ containing exactly r b_j -letters. We define the function $\tau_{b_j}(\phi)$ as follows:

$$\tau_{b_j}(\phi) = h_{b_j}^{\max_{b_j}(\phi)}(\phi)$$

Lemma 3.6.2 *Let K denote the connected-sum knot $K_{3,2} \# K_{p,q}$ where $K_{p,q}$ is a legendrian torus knot of even ∂ -class. Then for any element $\phi \in \text{Im}(\partial) \subset I = \langle \partial(a_1), \partial(a_2), \dots, \partial(\tilde{a}_1), \dots, \partial(\tilde{a}_k) \rangle$ (where \tilde{a}_i are the index-1 generators in the p, q -tangle and a_1, a_2 the other two index-1 generators, see Figure 3.7), the following property holds:*

- *The integer $\tau_{b_3}(\phi)$ is even.*

Proof. If there were no $K_{p,q}$ -tangle involved we recover the trefoil knot, where we have

$$\partial(a_1) = 1 + b_1 + b_3 + b_1 b_2 b_3$$

$$\partial(a_2) = 1 + 1 + b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2.^3$$




Figure 3.7: Connected sum of a (p, q) -torus knot with a trefoil and corresponding labelling of some of the crossings.

Since we do have an additional (p, q) -tangle involved in the knot between the crossing a_2 and the corresponding path joining it with b_1, b_2, b_3 , then there exists an expression P only containing terms in the (p, q) -tangle such that:

$$\partial(a_1) = 1 + b_1 + b_3 + b_1b_2b_3,$$

$$\partial(a_2) = 1 + P(1 + b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2).$$

Also, it is clear that b_1, b_2, b_3 do not appear in the expressions $\partial(\tilde{a}_i)$ for any of the \tilde{a}_i .

We have that

$$\tau_{b_3}(\partial(a_1)) = h_{b_3}^1(\partial(a_1)) = 2 \in 2\mathbb{Z},$$

$$\tau_{b_3}(\partial(a_2)) = h_{b_3}^1(\partial(a_2)) = 2 \cdot \#\{\text{words in } P\} \in 2\mathbb{Z}.$$

Therefore the Property holds for both a_i generators and it also holds trivially for the other \tilde{a}_i generators. Since $K_{p,q}$ is of even ∂ -class, all the expressions $\partial\tilde{a}_i$ contain an even number of words and, thus, it is clear then that the Property holds for any expression in the ideal I . Since $\text{Im}(\partial) \subset I$, the claim follows.

Lemma 3.6.3 *Let $K_{p,q}^\theta$ be a loop of Legendrian knots based on $K_{p,q}$, a Legendrian (p, q) -torus knot, and let \tilde{K}^θ be an arbitrary loop of Legendrian knots based on \tilde{K} . Denote by FE^θ the connected sum loop of both loops. The restriction of the holonomy morphism to the crossings of degree 0 in the $K_{p,q}$ block corresponding to the first part of the loop (when $K_{p,q}$ knot plays the role of the fly) coincide with the identity map.*

Proof. We must check that when a branch of the loop passes over (or below) the $K_{p,q}$ -block, the 0-index crossings of the (p, q) -block are mapped trivially by the holonomy morphism. In particular, we will show this fact for each of the elementary moves appearing in the homotopy.

For $R - III_a$ this follows automatically since it maps all generators trivially. Note that since the moving branch in any $R - III_b$ move corresponds to the elephant knot, the crossings of the

³ We write $1 + 1 + \dots$ even if we are working in \mathbb{Z}_2 -coefficients since when considering the additional (p, q) -tangle in the next case, one of the monogones will produce a different word.




Figure 3.8: Branch passing under and over a $K_{p,q}$ -block, respectively.

(p, q) -block can only play the role of the crossings y or z in diagram (B) of Figure 3.3 and, thus, are mapped trivially as well. For $R - II$ moves this fact follows from an application of Proposition 3.4.3 together with the fact that the differential of all 0-index crossings in a (p, q) -block is always zero. Finally, holonomies of $R - II^{-1}$ are the identity for all points except for two points (x and y crossings in Diagram (A) of Figure 3.3). But, in our case, these two points do not correspond to crossings in the (p, q) -block when a branch passes over (or below) the block. Therefore all the cases are covered and the claim follows. \square

Remark 3.6.4 *Another approach to prove the previous Lemma is based on observing that the action of any crossing in the fly knot is arbitrarily small with respect to the action of any crossing of the elephant knot. (We are thankful to Tobias Ekholm for this remark).*

Theorem 3.3. *Let $K_{2,3}$ be an embedded legendrian trefoil and $K_{p,q}$ a (p, q) -legendrian torus knot of even ∂ -class (recall Definition 3.5.11). Consider Kálmán's loop on the trefoil $K_{2,3}^\theta$ and another loop $K_{p,q}^\theta$ based on $K_{p,q}$. Let $(K_{2,3} \# K_{p,q})^\theta$ be the 1-parametric connected sum loop of both loops. Then the H_0 -restricted monodromy associated to $(K_{2,3} \# K_{p,q})_{H_0(K_{2,3} \# K_{p,q})}^\theta$ is not the identity and, therefore, this loop is not contractible in the space $\widehat{\mathfrak{Lcg}}(\mathbb{R}^3, \xi_{\text{std}})$.*

Proof. We will focus on the monodromy map restricted to the 0-index generators of the trefoil block. By Lemma 3.6.3 these generators are mapped trivially when the $(2, 3)$ -block plays the role of the tiny knot (half of the loop) and, thus, it suffices to study their holonomies for half the loop (when the block they form part of plays the role of the big knot). Take a diagrammatic realization for this half of the loop as in Figure 3.9.

The first arrow does not involve any Reidemeister move and, therefore, all crossings are mapped trivially via the holonomies. The second sequence of moves (first arrow in Figure 3.9) involves a bunch of II, II^{-1}, III_a and III_b -moves but neither of them affect the 0-index generators in the $(2, 3)$ -tangle, the reason being that II^{-1} and III -moves do not act on generators remotely and II -moves do not affect these generators by Proposition 3.4.3. Subsequent II and III_b -moves (third and fourth arrows, respectively) do not affect these generators for the same reason. The fifth arrow corresponds to a III_b -move that only affects b_1 mapping it to b_1 as it is shown in the diagram. The sixth arrow does not involve any Reidemeister move and thus maps trivially any crossing.

The seventh arrow involves a bunch of Reidemeister moves that do not affect the 0-index crossings of the $(2, 3)$ -tangle because the only move type that could potentially affect them is the II -type move but this is not the case by Proposition 3.4.3.

The eighth arrow corresponds to a II^{-1} -move whose holonomies restricted to 0-index crossings coincide with the holonomies of this move in Kálmán's original loop with an extra factor $W_{(p,q)}$ in b_1 . In order to check this, just note that one of the crossings disappearing via this II^{-1} move will map to zero whereas b_1 will be mapped to the differential of the other crossing before the move.

Therefore, the 0-index crossings are mapped as follows after the 8th arrow in Figure 3.9:

$$b_1 \mapsto (1 + b_3 A) W_{(p,q)}$$

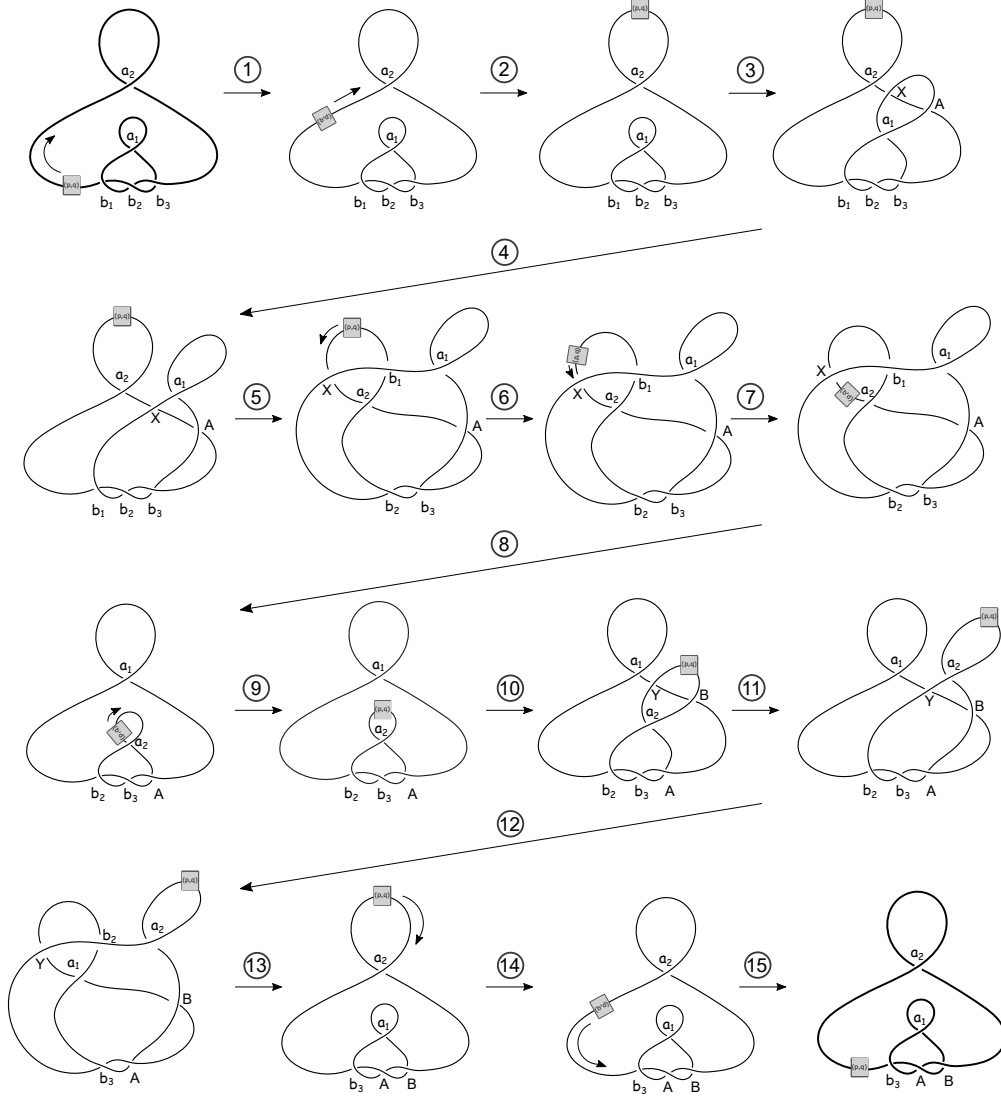


Figure 3.9: Half the loop consists on the concatenation of three times the depicted loop, which we denote by $\Omega_{3,2}^{\#(p,q)}$ (the trefoil plays the role of the big knot whereas the (p, q) –knot plays the role of the tiny knot).

$$b_2 \mapsto b_2$$

$$b_3 \mapsto b_3$$

The ninth arrow maps every crossing trivially since it does not involve any Reidemeister move.

There are a bunch of Reidemeister moves taking place in the tenth arrow but they do not affect 0-index generators of the $(2, 3)$ –tangle for the same reason as in previous arrows. The eleventh arrow corresponds to a III_b –move not affecting 0–index crossings of the $(2, 3)$ –tangle. The twelfth arrow maps b_2 to b_2 as shown in Figure 3.9 and any other 0–index crossing of the $(2, 3)$ –tangle is mapped trivially.

Arrow thirteenth corresponds to a II^{-1} –move that maps b_2 to $1 + AB$ and the rest of 0–index crossings are mapped trivially. Finally, arrows 14^{th} and 15^{th} do not affect 0–index crossings of the $(2, 3)$ –tangle.

Note that at the end of the loop b_3 takes the original position of b_1 , crossing A the one of b_2 and crossing B the one of b_3 . Therefore, the monodromy of $\Omega_{3,2}^{\#(p,q)}$ restricted to the 0-index generators of the $(2,3)$ -tangle is described as follows:

$$\begin{aligned} b_1 &\mapsto W_{(p,q)} + b_1 b_2 W_{(p,q)} \\ b_2 &\mapsto 1 + b_2 b_3 \\ b_3 &\mapsto b_1 \end{aligned}$$

Nonetheless, the loop consists on the concatenation of 3-times $\Omega_{3,2}^{\#(p,q)}$. We check how this monodromy map acts on generator b_3 . First, b_3 is mapped to b_1 in $\Omega_{3,2}^{\#(p,q)}$. It is mapped to $W_{(p,q)} + b_1 b_2 W_{(p,q)}$ after $\left(\Omega_{3,2}^{\#(p,q)}\right)^{\times 2}$. Finally, it is mapped to $W_{(p,q)} + (W_{(p,q)} + b_1 b_2 W_{(p,q)})(1 + b_2 b_3)W_{(p,q)}$ after $\left(\Omega_{3,2}^{\#(p,q)}\right)^{\times 3}$. In other words, the monodromy acts as follows on b_3 :

$$\mu(b_3) = W_{(p,q)} + W_{(p,q)}^2 + W_{(p,q)} b_2 b_3 W_{(p,q)} + b_1 b_2 W_{(p,q)}^2 + b_1 b_2 W_{(p,q)} b_2 b_3 W_{(p,q)}.$$

We need to show that $\mu(b_3) - b_3$ is not the zero class in $H_0(K_{2,3} \# K_{p,q})$.

$$\mu(b_3) - b_3 = W_{(p,q)} + W_{(p,q)}^2 + W_{(p,q)} b_2 b_3 W_{(p,q)} + b_1 b_2 W_{(p,q)}^2 + b_1 b_2 W_{(p,q)} b_2 b_3 W_{(p,q)} - b_3$$

does not satisfy the Property in Lemma 3.6.2 since $\tau_{b_3}(\mu(b_3) - b_3)$ is not an even number independently of the parity of the number of words in $W_{(p,q)}$. Therefore we conclude that $\mu(b_3) - b_3 \notin \text{Im}(\partial)$ and $[\mu(b_3) - b_3] \in H_0(K_{2,3} \# K_{p,q})$ is not the zero class in homology. This completes the proof.

As a corollary we can construct infinitely many new families of loops of Legendrians with non-trivial monodromy. Denote by $K_{p,q}^\theta$ Kálmán's loop for any (p, q) torus knot and by $K_{p,q}^\theta$ its inverse loop (i.e. with reverse orientation in the parameter). Denote by $K_{p,q}^{\text{const}}$ the constant loop based at $K_{p,q}$. Then, as an application of Theorem 3.3 we get:

Theorem 3.4 (Infinitely many new examples of loops with non-trivial monodromy).

The monodromy at the level of H_0 associated to the following families of loops is not the identity and, thus, all these families of loops are not contractible in the space $\widehat{\mathfrak{Lcg}}(\mathbb{R}^3, \xi_{\text{std}})$:

For $m \geq 1$, the connected sum loop of Kálmán's loop based at a trefoil and m loops K_{p_i, q_i}^θ where for $1 \leq i \leq m$ is one of the following:

- a) Kálmán's loop based at the torus knot K_{p_i, q_i} ,
- b) Kálmán's inverse loop based at the torus knot K_{p_i, q_i} , or
- c) The constant loop $K_{p_i, q_i}^{\text{const}}$ based at the torus knot K_{p_i, q_i} .




Figure 3.10: Connected sum of a trefoil with m torus knots K_{p_i, q_i}^θ , $1 \leq i \leq m$. For each $m \geq 1$, the loops described in Theorem 3.4 are based in a knot with this structure. First, the (p_i, q_i) -tangle blocks play the role of the fly while the trefoil plays the role of the elephant loop. Then, the roles are reversed.

Part II

Horizontal and transverse embedding spaces

Chapter 4

h –Principle for horizontal and transverse curves.

4.1 Introduction

In this section we recall some of the basic theory of singularity and rigidity for horizontal curves (Subsection 4.1.1). For further details we refer the reader to [61, 81].

4.1.1 Regularity of horizontal curves

We now recall how the phenomenon of singularity for horizontal curves shows up.

Given a (M, \mathcal{D}) distribution and a point $p \in M$, we write $\mathfrak{Maps}_p([0, 1]; M, \mathcal{D})$ for the space of horizontal maps of the interval $[0, 1]$ into (M, \mathcal{D}) with initial point $\gamma(0) = p$; we endow it with the C^∞ -topology.

Definition 4.1.1 *The **endpoint map** is defined as the evaluation map at $1 \in [0, 1]$:*

$$\begin{aligned} \mathbf{ep} : \mathfrak{Maps}_p([0, 1]; M, \mathcal{D}) &\longrightarrow M \\ \gamma &\longmapsto \gamma(1) \end{aligned}$$

This map is smooth. If it were submersive, its fibres would be smooth Frechet manifolds consisting of horizontal paths with given endpoints. The issue is that this is not always the case, leading to the conclusion that the fibres may develop singularities in which the tangent space is not well-defined. These singularities are thus horizontal curves that present issues in order to be deformed.

Definition 4.1.2 *A curve $\gamma \in \mathfrak{Maps}_p([0, 1]; M, \mathcal{D})$ is **regular** if the endpoint map \mathbf{ep} is submersive at γ . Otherwise, a curve is said to be **singular**.*

Equivalently, regularity means that, given any vector $v \in T_{\gamma(1)}M$, there exists a variation $(\gamma_s)_{s \in (-\varepsilon, \varepsilon)}$ such that

$$d_\gamma \mathbf{ep} \left(\frac{d}{ds} \gamma_s \right) = \frac{d}{ds} (\mathbf{ep}(\gamma_s)) = v.$$

We denote by Var_γ the space of infinitesimal variations of γ , endowed with the C^∞ -topology. A more down-to-earth description, when \mathcal{D} is a connection, is that Var_γ corresponds to the space of infinitesimal variations of the projection of γ (to the base of the bundle). This approach will be used repeatedly in upcoming sections.

For the purposes of this thesis, we are interested both in horizontal paths and horizontal loops. Then:

Definition 4.1.3 A curve $\gamma \in \mathcal{L}(M, \mathcal{D})$ is **regular** if it is regular as a path (using the quotient map $[0, 1] \rightarrow \mathbb{S}^1$ given by a choice of basepoint).

It is not difficult to see that being regular does not depend on the auxiliary choice of basepoint.

Remark 4.1.4 This is a well known fact. One could just elaborate on the argument in [68, Proposition 4, Corollary 5]: the variations generating infinitesimal directions in the source can be “localised” by an appropriate use of bump functions in the domain. Therefore, a finite basis of variations can be found and, thus, they can be modelled by a finite dimensional Euclidean vector space. One can just now choose another basepoint and the claim readily follows. See [30] for a similar approach.

4.2 ε -horizontality and ε -transversality

In this section we introduce ε -horizontal curves. These are curves that form an angle of at most ε with the distribution and thus serve as approximations of horizontal curves. They provide a convenient starting point for the h -principle arguments that will appear later in the chapter.

In Subsection 4.2.1 we introduce some additional notation regarding horizontal curves. ε -horizontal curves appear in Subsection 4.2.2. The main result is Proposition 4.2.4: the space of ε -horizontal curves is weakly equivalent to the space of formal horizontal curves. We then introduce analogues of this idea in the transverse setting. This is done in Subsections 4.2.3 and 4.2.4.

We assume that the reader is familiar with the h -principle language. The standard references on the topic are [39, 59].

4.2.1 Horizontal curves

Fix a manifold and a distribution (M, \mathcal{D}) . We already introduced the spaces of immersed horizontal loops $\mathfrak{Imm}(M, \mathcal{D})$ and embedded horizontal loops $\mathfrak{Emb}(M, \mathcal{D})$. The phenomenon of rigidity forced us to look instead into $\mathfrak{Imm}^r(M, \mathcal{D})$ and $\mathfrak{Emb}^r(M, \mathcal{D})$, the subspaces of regular curves. We want to compare these to the formal analogues $\mathfrak{Imm}^f(M, \mathcal{D})$ and $\mathfrak{Emb}^f(M, \mathcal{D})$. This comparison relates geometrically-defined spaces to spaces that are topological in nature.

Proofs in h -principle are local in nature. That is to say, in order to prove our theorems, we will reduce them to analogous statements for horizontal paths, relative boundary. This motivates us to introduce the following notation. Given a 1-dimensional manifold I , we write

$$\mathfrak{Imm}^r(I; M, \mathcal{D}) \longrightarrow \mathfrak{Imm}(I; M, \mathcal{D}) \longrightarrow \mathfrak{Imm}^f(I; M, \mathcal{D})$$

for the spaces of regular horizontal immersions, horizontal immersions, and formal horizontal immersions of I into (M, \mathcal{D}) . Similarly, we write

$$\mathfrak{Emb}^r(I; M, \mathcal{D}) \longrightarrow \mathfrak{Emb}(I; M, \mathcal{D}) \longrightarrow \mathfrak{Emb}^f(I; M, \mathcal{D})$$

in the case of embeddings. All spaces are endowed with the weak Whitney topology.

4.2.2 ε -horizontal curves

Being horizontal is a closed differential relation. These are typically more difficult to handle than open relations; dealing with them often requires some input from PDE theory or the use of a trick that transforms the problem into one involving an open relation. In this thesis we follow the second route, manipulating horizontal curves through their projections to the space of controls (Section 4.3).

We now introduce ε -horizontality. ε -horizontal curves can also be manipulated using their projections, with the added advantage of being described by an open relation. Fix a riemannian metric g in M . We can measure the (unsigned) angle \angle , in terms of the metric g , between any two linear subspaces at a given $T_p M$.

Definition 4.2.1 Fix a constant $0 < \varepsilon < \pi/2$. The space of ε -horizontal embeddings is defined as:

$$\mathfrak{Emb}^\varepsilon(M, \mathcal{D}) := \{ \gamma \in \mathfrak{Emb}(M) \mid \angle(\gamma', \mathcal{D}) < \varepsilon \}.$$

Its formal analogue, the space of **formal ε -horizontal embeddings**, is defined as:

$$\mathfrak{Emb}^{f,\varepsilon}(M, \mathcal{D}) := \{ (\gamma, (F_s)_{s \in [0,1]}) : \gamma \in \mathfrak{Emb}(M), \quad F_s \in \text{Mon}_{\mathbb{S}^1}(T\mathbb{S}^1, \gamma^* TM), \\ F_0 = \gamma', \quad \angle(F_1, \gamma^* \mathcal{D}) < \varepsilon \}.$$

4.2.2.1 Some flexibility statements

It is a classic result due to M. Gromov that the h -principle holds in the ε -horizontal setting:

Lemma 4.2.2 Consider (M, \mathcal{D}) with $\text{rank}(\mathcal{D}) \geq 2$. Then, the inclusion $\mathfrak{Emb}^\varepsilon(M, \mathcal{D}) \rightarrow \mathfrak{Emb}^{f,\varepsilon}(M, \mathcal{D})$ is a weak homotopy equivalence.

Proof. This follows from convex integration for open and ample relations [39, Theorem 18.4.1]. The relation is clearly open. Ampleness follows from the fact that principal subspaces are in correspondence with tangent fibres $T_p M$, and the relation in each is an open conical set (as depicted in Figure 4.1) that is path-connected and ample, because \mathcal{D} has at least rank 2.




Figure 4.1: Schematic depiction of the principal subspaces associated to ε -horizontality. The figure shows the rank 1 case, in which the relation is conical and open but has two components. When the rank is at least 2, the relation is path-connected and thus ample.

Furthermore:

Lemma 4.2.3 *The inclusion $\mathfrak{Emb}^f(M, \mathcal{D}) \rightarrow \mathfrak{Emb}^{f, \varepsilon}(M, \mathcal{D})$ is a homotopy equivalence. In particular, $\mathfrak{Emb}^\varepsilon(M, \mathcal{D})$ and $\mathfrak{Emb}^f(M, \mathcal{D})$ are weakly homotopy equivalent.*

Proof. Just note that the fiberwise orthogonal riemannian projection of F_1 onto \mathcal{D} provides a homotopy inverse.

These results that we have just stated are also relative in the parameter, relative in the domain, and satisfy C^0 -closeness. More precisely:

Proposition 4.2.4 *Let K be a compact manifold. Let (M, \mathcal{D}) be a manifold endowed with a distribution of rank greater or equal to 2. Suppose we are given a map $(\gamma, F_s) : K \rightarrow \mathfrak{Emb}^f([0, 1], M, \mathcal{D})$ satisfying the boundary conditions:*

- $(\gamma, F_s)(k)|_{\mathcal{O}_p(\{0, 1\})}$ is a ε -horizontal embedding for all $k \in K$,
- $(\gamma, F_s)(k) \in \mathfrak{Emb}^\varepsilon([0, 1], M, \mathcal{D})$ for $k \in \mathcal{O}_p(\partial K)$.

Then, (γ, F_s) extends to a homotopy $(\tilde{\gamma}, \tilde{F}_s) : K \times [0, 1] \rightarrow \mathfrak{Emb}^f([0, 1], M, \mathcal{D})$ that:

- *restricts to (γ, F_s) at time $s = 0$,*
- *maps into $\mathfrak{Emb}^\varepsilon([0, 1], M, \mathcal{D})$ at time $s = 1$*
- *is relative in the parameter (i.e. relative to $k \in \mathcal{O}_p(\partial K)$),*
- *is relative in the domain of the curves (i.e relative to $t \in \mathcal{O}_p(\{0, 1\})$),*
- *has underlying curves $\tilde{\gamma}(k, s)$ that are C^0 -close to $\gamma(k)$ for all k and s .*

4.2.2.2 The punchline

We can summarise the previous statements using the following commutative diagram:

$$\begin{array}{ccc} \mathfrak{Emb}^r(M, \mathcal{D}) \subset \mathfrak{Emb}(M, \mathcal{D}) & \longrightarrow & \mathfrak{Emb}^f(M, \mathcal{D}) \\ \downarrow & & \downarrow \cong \\ \mathfrak{Emb}^\varepsilon(M, \mathcal{D}) & \xrightarrow{\cong} & \mathfrak{Emb}^{f, \varepsilon}(M, \mathcal{D}) \end{array}$$

It follows that, in order to prove our main Theorem 1.2, it is sufficient to understand the inclusion $\mathfrak{Emb}^r(M, \mathcal{D}) \hookrightarrow \mathfrak{Emb}^\varepsilon(M, \mathcal{D})$. This simplification (passing from formal to ε) is commonly used in the h -principle literature, see for instance [84, 24].

4.2.2.3 The case of immersions

One can define, analogously, the space of immersed ε -horizontal loops:

$$\mathfrak{Imm}^\varepsilon(M, \mathcal{D}) := \{\gamma \in \mathfrak{Imm}(M) : \angle(\gamma', \mathcal{D}) < \varepsilon\}.$$

From the arguments above it follows that:

Lemma 4.2.5 *Let \mathcal{D} be a distribution with $\text{rank}(\mathcal{D}) \geq 2$. The map $\mathfrak{Imm}^\varepsilon(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}^f(M, \mathcal{D})$ is a weak homotopy equivalence.*

This h -principle is also relative in the parameter, relative in the domain, and C^0 -close. We leave it to the reader to spell out the analogue of Proposition 4.2.4. Theorem 1.1 reduces then to the study of the inclusion $\mathfrak{Imm}^r(M, \mathcal{D}) \hookrightarrow \mathfrak{Imm}^\varepsilon(M, \mathcal{D})$.

4.2.3 Transverse curves

We have already introduced the spaces of transverse immersed loops $\mathfrak{Imm}_\mathcal{T}(M, \mathcal{D})$, transverse embedded loops $\mathfrak{Emb}_\mathcal{T}(M, \mathcal{D})$, formally transverse immersions $\mathfrak{Imm}_\mathcal{T}^f(M, \mathcal{D})$, and formally transverse embeddings $\mathfrak{Emb}_\mathcal{T}^f(M, \mathcal{D})$. It was stated in the introduction that

$$\mathfrak{Imm}_\mathcal{T}(M, \mathcal{D}) \longrightarrow \mathfrak{Imm}_\mathcal{T}^f(M, \mathcal{D}), \quad \mathfrak{Emb}_\mathcal{T}(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}_\mathcal{T}^f(M, \mathcal{D})$$

are weak equivalences whenever the corank of \mathcal{D} is larger than 1, thanks to convex integration.

Assumption 4.2.6 *Whenever we work with transverse curves, we do so under the assumption that the corank of \mathcal{D} is one.*

4.2.3.1 Coorientations

Suppose that \mathcal{D} is coorientable and fix a coorientation. We do not need this assumption for our results. However, we will make use of it as follows: due the relative nature of the arguments, we will reduce our theorems to h -principles in which the target manifold is the Euclidean space. In this local picture, the distribution is parallelisable and co-parallelisable. Furthermore, formal transverse curves induce a preferred coorientation. This will allow us to define a suitable replacement of ε -horizontality in the transverse setting.

Definition 4.2.7 *A curve $\gamma : \mathbb{S}^1 \rightarrow M$ is **positively transverse** if γ' defines the positive coorientation in TM/\mathcal{D} .*

If \mathcal{D} is cooriented, $\mathfrak{Imm}_\mathcal{T}(M, \mathcal{D})$, $\mathfrak{Emb}_\mathcal{T}(M, \mathcal{D})$, $\mathfrak{Imm}_\mathcal{T}^f(M, \mathcal{D})$, and $\mathfrak{Emb}_\mathcal{T}^f(M, \mathcal{D})$ split into two different path components (the positively transverse and the negatively transverse). In order not to overload notation, we will follow the convention that if \mathcal{D} is cooriented, we focus on the positive component.

We can now fix a riemannian metric g on M and define the **oriented angle** $\angle(v(p), \mathcal{D}_p)$ between a vector $v \in T_p M$ and the corank-1 distribution \mathcal{D} . Its absolute value agrees with the (unsigned) angle $\angle(v, \mathcal{D})$ and its sign is positive if v is positively transverse.

4.2.3.2 Immersions

If I is a 1-manifold, we write

$$\mathfrak{Imm}_\mathcal{T}(I; M, \mathcal{D}) \rightarrow \mathfrak{Imm}_\mathcal{T}^f(I; M, \mathcal{D}), \quad \mathfrak{Emb}_\mathcal{T}(I; M, \mathcal{D}) \rightarrow \mathfrak{Emb}_\mathcal{T}^f(I; M, \mathcal{D})$$

for the spaces of transverse immersions, formal transverse immersions, transverse embeddings and formal transverse embeddings of I into (M, \mathcal{D}) . Once again, if \mathcal{D} is cooriented, these denote only the positively transverse component.




Figure 4.2: On the left we depict a vector field ν , providing the global choice of coorientation, and a positively transverse vector $v \in \mathcal{T}_p M$; the two lie on the same side of \mathcal{D} . On the right we depict a negatively transverse vector.

4.2.4 ε -transverse curves

Working under coorientability assumptions allows us to introduce the notion of ε -transversality.

Definition 4.2.8 Let \mathcal{D} be a cooriented distribution of corank-1. Fix a positive number $\pi/2 > \varepsilon > 0$. The space of ε -**transverse embeddings** is defined as:

$$\mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}(M, \mathcal{D}) = \{\gamma \in \mathfrak{Emb}(M) : \angle(\gamma', \mathcal{D}) > -\varepsilon\}.$$

We can also consider its formal analogue, the space of **formal ε -transverse embeddings**:

$$\mathfrak{Emb}_{\mathcal{T}}^{f, \varepsilon}(M, \mathcal{D}) = \left\{ (\gamma, (F_s)_{s \in [0,1]}) : \begin{array}{l} \gamma \in \mathfrak{Emb}(M), \quad F_s \in \text{Mon}_{\mathbb{S}^1}(TS^1, \gamma^* TM), \\ F_0 = \gamma', \quad \angle(F_1, \gamma^* \mathcal{D}) > -\varepsilon \end{array} \right\}.$$

4.2.4.1 Flexibility statements

The following is an analogue of Proposition 4.2.4, with a milder condition on the rank. The proof is analogous, using convex integration and projection to \mathcal{D} :

Proposition 4.2.9 Let \mathcal{D} be a cooriented corank-1 distribution with $\text{rank}(\mathcal{D}) \geq 1$. Then, the following inclusions are weak equivalences:

$$\mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}_{\mathcal{T}}^{f, \varepsilon}(M, \mathcal{D}), \quad \mathfrak{Emb}^f(M, \mathcal{D}) \longrightarrow \mathfrak{Emb}^{f, \varepsilon}(M, \mathcal{D}).$$

A statement that is relative in the parameter, relative in the domain, and C^0 -close also holds:

Proposition 4.2.10 Let K be a compact manifold. Let (M, \mathcal{D}) be a manifold endowed with a cooriented corank-1 distribution of rank at least 1. Suppose we are given a map $(\gamma, F_s) : K \rightarrow \mathfrak{Emb}_{\mathcal{T}}^f([0, 1]; M, \mathcal{D})$ satisfying:

- $(\gamma, F_s)(k)|_{\mathcal{O}_p(\{0,1\})}$ is a ε -transverse embedding for all $k \in K$,
- $(\gamma, F_s)(k) \in \mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}([0, 1]; M, \mathcal{D})$ for $k \in \mathcal{O}_p(\partial K)$.

Then, (γ, F_s) extends to a homotopy $(\tilde{\gamma}, \tilde{F}_s) : K \times [0, 1] \rightarrow \mathfrak{Emb}_{\mathcal{T}}^f([0, 1]; M, \mathcal{D})$ that:

- restricts to (γ, F_s) at time $s = 0$,
- maps into $\mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}([0, 1]; M, \mathcal{D})$ at time $s = 1$




Figure 4.3: The fiberwise relation defining ε -transversality defines a relation that is ample.

- is relative in the parameter (i.e. relative to $k \in \mathcal{O}p(\partial K)$),
- is relative in the domain of the curves (i.e. relative to $t \in \mathcal{O}p(\{0, 1\})$),
- has underlying curves $\tilde{\gamma}(k, s)$ that are C^0 -close to $\gamma(k)$ for all k and s .

4.2.4.2 The punchline

We obtain the following commutative diagram:

$$\begin{array}{ccc} \mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D}) & \longrightarrow & \mathfrak{Emb}_{\mathcal{T}^f}(M, \mathcal{D}) \\ \downarrow & & \downarrow \cong \\ \mathfrak{Emb}_{\mathcal{T}^\varepsilon}(M, \mathcal{D}) & \xrightarrow{\cong} & \mathfrak{Emb}_{\mathcal{T}^{f,\varepsilon}}(M, \mathcal{D}), \end{array}$$

telling us that we should focus on the inclusion $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D}) \hookrightarrow \mathfrak{Emb}_{\mathcal{T}^\varepsilon}(M, \mathcal{D})$. We will do so to prove Theorem 1.5.

4.2.4.3 The immersion case

We can also define the space of ε -transverse immersions $\mathfrak{Imm}_{\mathcal{T}^\varepsilon}(M, \mathcal{D})$ and deduce that:

Lemma 4.2.11 *Let \mathcal{D} be a cooriented, corank-1 distribution of rank at least 1. The map $\mathfrak{Imm}_{\mathcal{T}^\varepsilon}(M, \mathcal{D}) \rightarrow \mathfrak{Imm}_{\mathcal{T}^f}(M, \mathcal{D})$ is a weak homotopy equivalence.*

This h -principle is also relative in the parameter, relative in the domain, and C^0 -close. We will henceforth focus on the inclusion $\mathfrak{Imm}_{\mathcal{T}}(M, \mathcal{D}) \hookrightarrow \mathfrak{Imm}_{\mathcal{T}^\varepsilon}(M, \mathcal{D})$ in order to prove Theorem 1.19.

4.2.4.4 Almost transversality

To wrap up this section, consider the following definition:

Definition 4.2.12 *Let \mathcal{D} be a cooriented, corank-1 distribution. Let I be a 1-dimensional manifold. The space of **almost transverse embeddings** is:*

$$\mathfrak{Emb}_{\mathcal{AT}}(I; M, \mathcal{D}) := \{ \gamma \in \mathfrak{Emb}(I, M) \mid \angle(\gamma', \mathcal{D}) \geq 0 \}.$$

We write $\mathfrak{Emb}_{\mathcal{AT}}(M, \mathcal{D})$ in the particular case of loops.

This may be regarded as the closure of the space of positively transverse embeddings $\mathfrak{Emb}_{\mathcal{T}}(I, M, \mathcal{D})$. We only introduce it because it will allow us to translate flexibility statements about horizontal curves to the transverse setting.

4.3 Graphical models

One of the two standard projections used in Contact Topology to study legendrian (i.e. horizontal) knots in standard contact $(\mathbb{R}^3, \xi_{\text{std}} = \text{Ker}(dy - zdx))$ is the so-called Lagrangian projection:

$$\begin{aligned} \pi_L : \mathbb{R}^3 &\longrightarrow \mathbb{R}^2 \\ (x, y, z) &\mapsto (x, z). \end{aligned}$$

It projects ξ_{std} , at each $p \in \mathbb{R}^3$, isomorphically onto the tangent space $T_{\pi_L(p)}\mathbb{R}^2$ of the base. This projects a legendrian knot to an immersed planar curve. From the projected curve one can recover the missing y -coordinate by integration:

$$z(t) = z(t_0) + \int_{t_0}^t x(s)y'(s)ds.$$

Indeed, the integral on the right-hand side, when evaluated over the whole curve, computes the area it bounds due to Stokes' theorem. This turns the problem of manipulating Legendrian knots into a problem about planar curves satisfying an area constraint.

In this Section we introduce *graphical models*. These are opens in Euclidean space, endowed with a bracket-generating distribution that is graphical over some of the coordinates; we call these the *base*. Projecting to the base and manipulating curves there is analogous to using the Lagrangian projection. This line of reasoning is also classic in Geometric Control Theory: the tangent space of the base, upon choosing a framing, corresponds to the space of controls.

Graphical models are introduced in Subsection 4.3.1. The related notion of ODE model appears in Subsection 4.3.2. In Subsection 4.3.3 we explain how to cover any (M, \mathcal{D}) by graphical models. These local models will be used in Section 4.4 to manipulate horizontal and transverse curves.

4.3.1 Graphical models

Fix a rank q , an ambient dimension n , and a step m . We now introduce the main definition of this section. It may remind the reader of the ideas used to construct balls in Carnot-Caratheodory geometry [51, 60, 81]:

Definition 4.3.1 *A graphical model is a tuple consisting of:*

- a radius $r > 0$,
- a constant-growth, bracket-generating, rank- q distribution \mathcal{D} defined over the ball $B_r \subset \mathbb{R}^n$,
- the projection $\pi : B_r \subset \mathbb{R}^n \rightarrow \mathbb{R}^q$ to the so-called **base**,
- a framing $\{X_1, \dots, X_n\}$ of TB_r .

The Lie flag of \mathcal{D} will be denoted by

$$\mathcal{D} = \mathcal{D}_1 \subset \mathcal{D}_2 \subset \cdots \subset \mathcal{D}_m = TB_r,$$

and we write $q_i = \text{rank}(\mathcal{D}_i)$.

This tuple must satisfy the following conditions:

- $\{X_1, \dots, X_{q_i}\}$ is a framing of \mathcal{D}_i .
- Given $j = q_{i-1} + 1, \dots, q_i$, there is a formal bracket expression A_j and a collection of indices $l_a^j = 1, \dots, q$ satisfying:

$$X_j = A_j(X_{l_1^j}, \dots, X_{l_i^j}),$$

a. $d\pi(X_j) = \partial_j$ for all $j = 1, \dots, q$. In particular, $d\pi$ is a fibrewise isomorphism between \mathcal{D} and $T\mathbb{R}^q$.

b. $X_j(0) = \partial_j$ for all $j = 1, \dots, n$.

The first two conditions are unnamed because they simply state that the given framing is compatible with the Lie flag. Condition (a) says that the distribution is a connection over \mathbb{R}^q and that its framing is the lift of the standard coordinate framing. Condition (b) controls $\{X_1, \dots, X_n\}$, saying that they agree with the standard coordinate directions at the origin. This will allow us to describe, quantitatively, how paths in the base lift to horizontal curves. As one may expect, we will be able to estimate this up to an error of size $O(r)$.

4.3.2 ODE models

Since \mathcal{D} is a connection over \mathbb{R}^q , any curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}^q$ can be lifted to a horizontal curve of \mathcal{D} , uniquely once a lift of $\gamma(0)$ has been chosen. Conversely, any horizontal curve is recovered uniquely from its projection by lifting (using the appropriate initial point). This is a consequence of the fundamental theorem of ODEs.

The caveat is that \mathcal{D} is only defined over B_r , so the claimed lift may escape the model and therefore not exist for all times; this is the usual issue one encounters with non-complete flows. Still, the lift

$$\text{lift}(\gamma) : U \rightarrow (B_r, \mathcal{D})$$

is uniquely defined over some maximal open interval $U \subset \mathbb{R}$ that contains zero. In order to discuss this a bit further, we introduce:

Definition 4.3.2 Consider a curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}^q$ mapping to the base of a graphical model. Thanks to π , \mathbb{R}^n can be seen as a fibration F over \mathbb{R}^q , allowing us to consider the tautological map

$$\Psi : \gamma^*F \cong \mathbb{R} \times \mathbb{R}^{n-q} \rightarrow \mathbb{R}^n = \mathbb{R}^q \times \mathbb{R}^{n-q}$$

that is transverse to \mathcal{D} (whenever the latter is defined).

The **ODE model** associated to γ (and to the graphical model) consists of:

- an open subset $D \subset \gamma^*F \cap \Psi^{-1}(B_r)$,
- the tautological map $\Psi : D \rightarrow B_r$,
- the line field $\Psi^*\mathcal{D}$, whose domain of definition is D .

That is, (D, Ψ) parametrises the region of B_r that lies over γ . Do note that Ψ is an immersion/embedding only if γ itself is immersed/embedded. Our discussion above states that lifting γ amounts to choosing a basepoint in D , solving the ODE given by $\Psi^*\mathcal{D}$, and pushing forward with Ψ .




Figure 4.4: A graphical model and an ODE model. The distribution is seen as a connection, where $T\mathbb{R}^n = \mathcal{D} \oplus \text{Ker } d\pi$. A curve γ is shown in the base, with its ODE model lying above. \mathcal{D} restricts to the ODE model as a line field. A lift of γ is shown in magenta.

The reason behind introducing ODE models is that they allow us to state the following trivial lemma:

Lemma 4.3.3 *Fix a graphical model and consider the following objects:*

- A curve $\gamma : [0, 1] \rightarrow \mathbb{R}^q$ mapping to the base.
- A (defined for all time) lift $\text{lift}(\gamma) : [0, 1] \rightarrow B_r$ of γ .
- The ODE model $(D, \Psi, \Psi^*\mathcal{D})$ of γ .
- The unique integral curve $\nu : [0, 1] \rightarrow (D, \Psi^*\mathcal{D})$, graphical over $[0, 1]$, such that $\Psi \circ \nu = \text{lift}(\gamma)$.

Then, there is a constant $\delta > 0$ and coordinates $\phi : [0, 1] \times B_\delta \rightarrow D$ such that:

- $\phi(t, 0) = \nu(t)$,
- $\phi^*\Psi^*\mathcal{D}$ is spanned by ∂_1 (the first coordinate direction).

This is a consequence of the flowbox theorem, so we will call ϕ **flowbox coordinates**. This allows us to see horizontal/transverse curves as graphs of functions and see horizontality/transversality in terms of their slope. See Figure 4.5.

4.3.2.1 Size of ODE models

Using the properties of a graphical model, we can estimate how large the constant δ appearing in Lemma 4.3.3 can be.

Write $\pi_v : D \rightarrow \mathbb{R}^{n-q}$ for the standard projection to the vertical. According to the definition of graphical model, \mathcal{D} and the foliation by planes parallel to \mathbb{R}^q differ by $O(r)$. In particular, the




Figure 4.5: An ODE model in flowbox coordinates. The distribution is of corank-1 and its coorientation is given by “going up”. It contains a ε -horizontal curve γ with three differentiated regions. The left region is negatively transverse. In right region the curve is positively transverse. In the middle region γ is parallel to the x -axis and is thus horizontal.

slope of \mathcal{D} is bounded by $O(r)$. The same holds for $\Psi^*\mathcal{D}$ in D . This implies a bound for the vertical displacement:

$$|\pi_v \circ \text{lift}(\gamma)(1) - \pi_v \circ \text{lift}(\gamma)(0)| < \text{len}(\gamma)O(r) \quad (4.1)$$

where $\text{len}(\gamma)$ is the length of γ . It is valid as long as it is smaller than the distance d of $\gamma(0)$ to the boundary of the model. Choosing $\text{len}(\gamma)$ sufficiently small, we can choose δ to be of the magnitude of d .

The punchline is that, in order to manipulate horizontal curves effectively on a manifold (M, \mathcal{D}) , it will be necessary to cover it with graphical models whose radii are very small, as this will allow us to estimate vertical displacement of curves in an effective manner.

4.3.3 Adapted charts

Fix a manifold M and a bracket-generating distribution \mathcal{D} . We now prove that (M, \mathcal{D}) can be covered by graphical models. For notational ease, let us introduce a definition first. Given a point $p \in M$, an **adapted chart** is a graphical model $(B_r \subset \mathbb{R}^n, \mathcal{D}_U)$ together with a chart

$$\phi : (B_r, \mathcal{D}_U) \longrightarrow (M, \mathcal{D}),$$

such that $\phi^*\mathcal{D} = \mathcal{D}_U$ and $\phi(0) = p$.

Proposition 4.3.4 *Let (M, \mathcal{D}) be a manifold endowed with a bracket-generating distribution. Then, any point $p \in M$ admits an adapted chart.*

Proof. We argue at a fixed but arbitrary point $p \in M$. Fix a basis $\{Y_1, \dots, Y_n\}$ of $T_p M$ such that $\{Y_1, \dots, Y_q\}$ spans \mathcal{D}_p . We construct local coordinates around p applying the exponential map. In these new coordinates we have that Y_i is ∂_i . Condition (a) in the definition of graphical model follows.

In the new coordinates, we have local projection $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^q$. This allows us to define a framing $F_{\mathcal{D}} := \{X_1, \dots, X_q\}$ of \mathcal{D} by lifting the coordinate vector fields ∂_i of \mathbb{R}^q . Note that these X_i coincide with the Y_i at the point p . Due to the bracket-generating condition, all vector fields around p can be written as linear combinations of Lie brackets involving vector fields in \mathcal{D} . Since all such vector fields are themselves sums of elements in $F_{\mathcal{D}}$, it can be deduced that TM is spanned by bracket expressions involving only $F_{\mathcal{D}}$. This allows us to extend $F_{\mathcal{D}}$ to a frame $\{X_1, \dots, X_n\}$ such that:

- $\{X_1, \dots, X_{q_i}\}$ spans \mathcal{D}_i .
- $X_j = A_j(X_{i_1}, \dots, X_{i_l})$, with $i_a \leq q$ for every $a = 1, \dots, l$ and A_j some bracket-expression.

By construction, the elements in $F_{\mathcal{D}}$ commute with one another upon projecting to \mathbb{R}^q . I.e. their Lie brackets are purely vertical, meaning that the vector fields $\{X_{q+1}, \dots, X_n\}$ are tangent to the fibres of π . This implies that, by applying a linear transformation fibrewise, we can produce new coordinates in which $X_j(p)$ is ∂_j . Condition (b) follows.

4.3.3.1 Covering by nice adapted charts

It is crucial for our arguments to be able to produce coverings by graphical models that are arbitrarily fine, and whose behaviour is controlled regardless of how fine we need them to be. That is the content of the following corollary:

Lemma 4.3.5 *Let (M, \mathcal{D}) be a compact manifold endowed with a bracket-generating distribution. Then, there are constants $C, r > 0$ such that:*

- i. Any point $p \in M$ admits an adapted chart of radius r .*
- ii. The bound $|X_j(x) - \partial_j| < C|x|$ holds for all such adapted charts and all elements X_j in the corresponding framing.*

The measuring in both properties is done using the Euclidean distance given by the adapted chart.

Proof. The construction in Proposition 4.3.4 is parametric on p . This is certainly true for the exponential map, which yields the resulting local coordinates, the projection π , and thus the frame $F_{\mathcal{D}}$. This is not necessarily the case, globally, for the choice of bracket expressions A_j , but it is still true if we argue on opens of some sufficiently fine cover $\{U_i\}$ of M . The fibrewise linear transformation chosen at the end of the argument is unique and thus parametric.

It follows that the statement holds for constants $C_i, r_i > 0$ over each U_i . We extract a finite covering to conclude the argument.

Do note that, upon zooming-in at p , the distribution \mathcal{D} locally converges to a Carnot group called the nilpotentisation [81, Chapter 4]. That is, to a nilpotent Lie group endowed with a bracket-generating distribution that is left-invariant and invariant under suitable weighted scaling. This implies that, by taking a sufficiently fine cover of (M, \mathcal{D}) , one can produce graphical models that are as close as required to a Carnot group. This is an improvement on Lemma 4.3.5, but it is not needed for our arguments.

4.4 Microflexibility of curves

The results in this thesis follow an overall strategy that is standard in h -principle. Namely: we first perform a series of simplifications that are meant to reduce the proof to a problem that is localised in a small ball. We call this *reduction*. Reduction arguments can be technical but often follow some standard heuristics and patterns. Once we have passed to a localised setting, the second step begins. This is the core of the proof and requires some input that is specific to the geometric setup at hand.

This step is called *extension* because it often amounts to extending a solution from the boundary of small ball to its interior¹.

In this section we prove several lemmas dealing with deformations of horizontal and transverse curves; they are meant to be used in the reduction step. These deformations often take place along stratified subsets of positive codimension, and can therefore be understood as microflexibility phenomena for horizontal and transverse curves². Proofs boil down to patching up local constructions happening in graphical models (Section 4.3).

In Subsection 5.4.4 we review Thurston's jiggling, which we use to triangulate our manifolds and thus argue one simplex at a time. Local statements taking place in graphical models are presented in Subsection 4.4.2. We globalise these constructions in Subsections 4.4.3 (for horizontal curves) and 4.4.4 (for transverse curves).

4.4.1 Triangulations

In order to reduce our arguments to a euclidean situation, we fix a (sufficiently fine) triangulation and we work on neighbourhoods of the simplices. For our purposes it is important that these triangulations are well-behaved. This can be achieved using the Thurston jiggling Lemma. We state it for the case of line fields (which is all we need):

Lemma 4.4.1 (Thurston, [101]) *Let N be a smooth n -manifold equipped with a line field ξ . Fix a metric g . Then, there exists a sequence of triangulations \mathcal{T}_b satisfying the following properties:*

- i. Each simplex $\Delta \in \mathcal{T}_b$ is transverse to ξ .*
- i'. Each n -simplex is homeomorphic to a flowbox of ξ .*
- ii. The radius of each simplex in \mathcal{T}_b is bounded above by $1/b$.*
- iii. The number of simultaneously incident simplices in \mathcal{T}_b is bounded above by a constant independent of b .*

Conditions (i) and (i') say that ξ is almost constant (upon taking it fine enough) in the coordinates provided by Δ . Condition (ii) says that the triangulations are becoming finer as b increases (and indeed all of them can be assumed to be refinements of some given triangulation). Condition (iii) states that the combinatorics of the triangulation remain controlled upon refinement (which is needed to prove Condition (i')).

We will also need a version for manifolds with the boundary:

Corollary 4.4.2 *In Lemma 4.4.1, suppose that N has boundary. Then we can furthermore assume that:*

- \mathcal{T}_b extends a triangulation of the boundary ∂N .*
- Conditions (i), (i'), (ii), (iii) hold for all simplices of \mathcal{T}_b not fully contained in ∂N .*
- The pair $(\partial N, \mathcal{T}_b|_{\partial N})$ satisfies the conclusions of the Lemma.*

¹ The proof of Theorem 1.5 follows these general lines. The proof of Theorem 1.2 presents some subtleties that force us to do something slightly different; see Remark 4.7.2

² Do note that, due to rigidity, horizontal curves are not microflexible in general.

4.4.2 Local arguments

We now present a series of statements dealing with families of curves mapping into graphical models. In order to streamline notation, let us denote the target graphical model by (V, \mathcal{D}) . Its projection to the base \mathbb{R}^q is denoted by π and the projection to the vertical by π_{\downarrow} . We also write K for the smooth compact manifold serving as the parameter space of the families.

4.4.2.1 Horizontalisation in graphical models

We will often construct horizontal curves as lifts of curves in \mathbb{R}^1 . The following is a quantitative statement about the existence of lifts:

Lemma 4.4.3 *Consider a family $\gamma : K \rightarrow \mathfrak{Emb}^\varepsilon([0, 1]; \mathbb{D}_{r_1}, \mathcal{D})$. Then, there is a unique family $\nu : K \rightarrow \mathfrak{Emb}([0, \delta]; \mathbb{D}_r, \mathcal{D})$ satisfying*

- $\pi \circ \nu = \pi \circ \gamma$.
- $\nu(k)(0) = \gamma(k)(0)$.

Furthermore: Let l be an upper bound for the velocity $\|(\pi \circ \gamma(k))'\|$. Then we can assume

$$\delta > \frac{r - r_1}{l(\varepsilon + O(r))}.$$

Proof. The uniqueness of ν is immediate from the discussion in Subsection 4.3.2, since ν is obtained from $\pi \circ \gamma$ by lifting with a given initial value. The bound on δ follows from the fact that the slope of \mathcal{D} with the horizontal is at most $\varepsilon + O(r)$, so the difference between γ and ν , which is purely vertical is bounded by $\delta l(\varepsilon + O(r))$. For ν to remain within the r -ball, this quantity must be smaller than $r - r_1$, yielding the claim.

Do note that the coefficient in the expression $O(r)$ can be bounded above in terms of the derivatives of \mathcal{D} . In Lemma 4.3.5 we observed that this coefficient can be bounded globally over a compact manifold.

4.4.2.2 Stability of horizontalisation

Given a family of horizontal curves, we may want to produce a nearby horizontal family by manipulating its projection to the base. The following lemma says that this is indeed possible:

Lemma 4.4.4 *Consider families*

$$\gamma : K \rightarrow \mathfrak{Emb}([0, 1]; \mathbb{D}_r, \mathcal{D}), \quad \alpha : K \rightarrow \mathfrak{Imm}([0, 1]; \mathbb{R}^q)$$

such that $\pi \circ \gamma$ and α are C^0 -close and their lengths are close. Then, there is a lift $\nu : K \rightarrow \mathfrak{Imm}([0, 1]; \mathbb{D}_r, \mathcal{D})$ of α that is C^0 -close to γ .

Proof. The family ν is obtained by lifting α , as in Subsection 4.3.2. The conclusion forces us to choose an initial value that is close to $\gamma(k)(0)$. The hypothesis on α (closeness in C^0 and length) imply (Lemma 4.4.8) that the ODE behind the lifting process is close to the ODE associated to $\pi \circ \gamma$. This implies that the lifting exists over $[0, 1]$ and is close to γ . \square

In concrete instances we will be able to argue that the resulting family ν also consists of embedded curves. This will follow from the specific properties of the family α under consideration.

4.4.2.3 Interpolation statements

We sometimes consider deformations of ε -horizontal curves in which the projection to the base remains fixed and the vertical component changes. This is explained in the following lemma, whose proof we leave to the reader:

Lemma 4.4.5 *Fix a family of curves $\alpha : K \rightarrow \mathfrak{Imm}([0, 1]; \mathbb{R}^q)$. Then, there exists a constant $\delta > 0$ such that any two families of curves*

$$\gamma, \nu : K \rightarrow \mathfrak{Imm}^\varepsilon([0, 1]; \mathbb{D}_r, \mathcal{D})$$

lifting α and satisfying $|\gamma(k) - \nu(k)|_{C^0} < \delta$ are homotopic through a family of ε -horizontal curves also lifting α .

Furthermore, this homotopy may be assumed to be relative to $\mathcal{O}p(\partial(K \times I))$ if the families already agree there.

The analogue for transverse curves reads:

Lemma 4.4.6 *Suppose \mathcal{D} is of corank-1 and cooriented. Fix a family of curves $\alpha : K \rightarrow \mathfrak{Imm}([0, 1]; \mathbb{R}^q)$. Then, there exists a constant $\delta > 0$ such that any two families of curves*

$$\gamma, \nu : K \rightarrow \mathfrak{Imm}_{\mathcal{T}}([0, 1]; \mathbb{D}_r, \mathcal{D})$$

lifting α and satisfying $|\gamma(k) - \nu(k)|_{C^0} < \delta$ are homotopic through a family of transverse curves also lifting α .

Furthermore, this homotopy may be assumed to be relative to $\mathcal{O}p(\partial(K \times I))$ if the families already agree there.

4.4.2.4 Deforming ε -horizontal curves

The following is an analogue of Lemma 4.4.4 in the ε -horizontal setting.

Lemma 4.4.7 *Consider families*

$$\gamma : K \rightarrow \mathfrak{Emb}^\varepsilon([0, 1]; \mathbb{D}_r, \mathcal{D}) \quad \alpha : K \rightarrow \mathfrak{Imm}([0, 1]; \mathbb{R}^q)$$

such that $\pi \circ \gamma$ and α are C^0 -close and their lengths are close. Then, there is a lift $\nu : K \rightarrow \mathfrak{Imm}^\varepsilon([0, 1]; \mathbb{D}_r, \mathcal{D})$ of α that is C^0 -close to γ .

Proof. Constructing ν amounts to choosing its vertical component $\pi_\perp \circ \nu$. Naively, we could set $\pi_\perp \circ \nu = \pi_\perp \circ \gamma$, but there is no reason why this would preserve ε -horizontality. The strategy to be pursued instead is to mimic the proof of Lemma 4.4.4. Namely, we want to see γ as the solution of an ODE (that makes an angle of at most ε with respect to \mathcal{D}) and produce ν as a solution of a similar ODE.

We define families of vector fields

$$X, Y, Z, W : K \times I \rightarrow T(TV)$$

satisfying the condition:

- X is tangent to \mathcal{D} and satisfies $d\pi(X(k, t)) = d\pi(\gamma(k)'(t))$ over all points lying over $\pi \circ \gamma(k)(t)$.

- Y is vertical and satisfies $Y(k, t) = X(k, t) - \gamma(k)'(t)$. It follows that, along γ , we have the inequality:

$$\frac{|Y(k, t)|}{|\gamma(k)'(t)|} = \angle(\gamma(k)'(t), \mathcal{D}) < \sin(\varepsilon).$$

- W is tangent to \mathcal{D} and satisfies $d\pi(W(k, t)) = \alpha(k)'(t)$ over all points projecting to $\alpha(k)(t)$
- Z is vertical and given by the expression $Z(k, t) = Y(k, t) \frac{|\alpha(k)'(t)|}{|X(k, t)|}$.

According to these definitions, $\gamma(k)$ is an integral line of $X(k, t) + Y(k, t)$. We define $\nu(k)(t)$ to be the integral line of $W(k, t) + Z(k, t)$ with initial condition $\gamma(k)(0)$.

By construction, $\pi \circ \nu = \alpha$ and therefore ν is immersed. Furthermore, due to our definition Z , ν is ε -horizontal as long as α is sufficiently C^1 -close to $\pi \circ \gamma$. Lastly, C^0 -closeness of ν and γ follows from the closeness of α and $\pi \circ \gamma$ in C^0 and length.

4.4.3 Horizontalisation

We now present semi-local analogues of Lemma 4.4.3. Since the Lemma only provides short-time existence of horizontal curves, generalisations must also present this feature. The reader should think of the upcoming statements as analogues of the holonomic approximation theorem [39, Theorem 3.1.1]. However, they involve no wiggling.

The general setup is the following: We fix a pair (M, \mathcal{D}) . The distribution need not be bracket-generating. Our families of curves are parametrised by a compact manifold K and have the unit interval $I = [0, 1]$ as their domain. The product $K \times I$ contains a stratified subset A such that all its strata are transverse to the I -factor.

4.4.3.1 Horizontalisation along the skeleton

The following result shows that any family of ε -horizontal curves can be made horizontal on a neighbourhood of .

Lemma 4.4.8 *Given a family $\gamma : K \rightarrow \mathfrak{Emb}^\varepsilon(I; M, \mathcal{D})$, there exists a family*

$$\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}^\varepsilon(I; M, \mathcal{D})$$

such that:

- i. $\tilde{\gamma}(k, 0) = \gamma(k)$.
- ii. $\tilde{\gamma}(k)(t) = \gamma(k)(t)$ if $(k, t) \in (K \times I) \setminus \mathcal{Op}A$
- iii. $\tilde{\gamma}(k, s)$ is C^0 -close to $\gamma(k)$, for all s .
- iii'. the length of $\tilde{\gamma}(k, s)$ is close to the length of $\gamma(k)$, for all s .
- iv. $\tilde{\gamma}(k, 1)$ is horizontal close to A .

Proof. The proof is inductive on the strata of A , starting from the smallest one. At a given step, working with a stratum B , we will achieve Property (iv) over B , preserving it as well along smaller strata. The other properties will follow as long as our perturbations are small and localised close to A .

Let U be a neighbourhood of the smaller strata in which $\tilde{\gamma}(k, 1)$ is already horizontal. We can then consider a closed submanifold $B' \subset B$ such that $\{B', U\}$ cover B . We can triangulate B' using Lemma 4.4.1, turning it into a stratified set itself, so that each simplex Δ is mapped by γ to some adapted chart (V, ϕ) . We then proceed inductively from the smaller simplices. A crucial observation is that simplices along ∂B are contained in U and therefore no further changes are required there.

For the inductive step consider an l -simplex Δ . The inductive hypothesis is that there is a family of curves β that has been obtained from γ by a homotopy satisfying Properties (i) to (iii') and that is already horizontal over all smaller simplices (and U). Since B is transverse to the I -direction, Δ has a neighbourhood parametrised as

$$\Phi : \mathbb{D}^l \times \mathbb{D}^{k-l} \times (-\delta, \delta) \longrightarrow K \times I.$$

The map Φ preserves the foliation in the direction of the last component and $\Phi|_{\mathbb{D}^l \times 0}$ is an arbitrarily small extension of Δ to a smooth disc. We write $\eta = \phi \circ \beta \circ \Phi$ for the restriction of the family to this neighbourhood, mapping now into the graphical model V . From the induction hypothesis it follows that η is horizontal over $\mathcal{O}p(\partial \mathbb{D}^l) \times \mathbb{D}^{k-l} \times (-\delta, \delta)$.

We have thus reduced the claim to the situation in which our stratified set is just a disc, and we have to work relatively to the boundary of $\mathbb{D}^l \times \mathbb{D}^{k-l} \times (\delta, \delta)$. There is a unique family of horizontal curves ν such that ν and η share the same projection to the base of V and such that $\nu(k)(0) = \eta(k)(0)$ for all $k \in \mathbb{D}^l \times \mathbb{D}^{k-l}$. We can argue that this family exists for all time if our triangulation was fine enough. Alternatively, we just observe that there is some $\delta' > 0$ such that ν is defined over $\mathbb{D}^l \times \mathbb{D}^{k-l} \times (-\delta', \delta')$ and η lives within an ODE model associated to it (Lemma 4.3.3).

We now deform η , relatively to the boundary of the model, to a family that agrees with ν over $\mathbb{D}^l \times \mathbb{D}^{k-l} \times (-\delta, \delta)$. We can do so keeping the projection to the base the same (Lemma 4.4.3). This, together with a sufficiently small choice of δ' , guarantees Properties (iii) and (iii'). This concludes the inductive argument to handle a stratum B and thus the inductive argument across all strata.

It is immediate from the proof that the statement also holds relatively to regions of A in which the curves are already horizontal.

Corollary 4.4.9 *Assume that \mathcal{D} is cooriented of corank-1 and that γ is positively transverse. Then, the conclusions of Lemma 4.4.8 hold and additionally $\tilde{\gamma}$ can be assumed to be almost transverse.*

Proof. The horizontalisation process described in the proof of Lemma 4.4.8 was based on passing locally to some ODE model. In such a model it is immediate that introducing zero slope (making the curves horizontal) can be done while preserving non-negative slope everywhere (being almost transverse).

4.4.3.2 Direction adjustment

Lemma 4.4.8 explained to us how to perturb a family of ε -horizontal curves so that it becomes horizontal along A . The next lemma states that one can prescribe the behaviour along A , as long as A is contractible and $\text{rank}(\mathcal{D}) \geq 2$.

Lemma 4.4.10 *Suppose that A is a k -disc and \mathcal{D} has rank at least 2. Fix a family $\gamma : K \rightarrow \mathfrak{Emb}^\varepsilon(I; M, \mathcal{D})$ and a family of horizontal curves ν , defined only on a neighbourhood of A . Assume that $\nu|_A = \gamma|_A$.*

Then, there exists a family $\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}^\varepsilon(I; M, \mathcal{D})$ such that the conclusions of Lemma 4.4.8 hold and, additionally:

v . $\tilde{\gamma}(-, 1)$ agrees with ν in $\mathcal{O}p(A)$.

Proof. Since the argument takes place on a neighbourhood of A and is relative to its boundary, we may as well assume that $K = \mathcal{O}p(\mathbb{D}^k)$ and $A = \mathbb{D}^k \times \{1/2\} \subset K \times [0, 1]$.

Since A is contractible and \mathcal{D} has rank at least 2, we can find a tangential rotation

$$(v_\theta)_{\theta \in [0,1]} : \mathbb{D}^k \longrightarrow TM$$

such that:

- $v_\theta(k) \in T_{\gamma(k)(1/2)}M$ is ε -horizontal.
- $v_0(x) = \gamma(k)'(1/2)$.
- $v_1(x) = \nu(x)'(1/2)$.

That is, v is a lift of $\gamma|_A$ providing a tangential rotation of its velocity vector to the velocity vector of ν .

A further simplification enters the proof now: γ may be assumed to take values in a graphical model V . Otherwise we triangulate A in a sufficiently fine manner and argue inductively on neighbourhoods of its simplices. There is then a homotopy of linear maps

$$(\Phi_\theta)_{\theta \in [0,1]} : \mathbb{D}^k \longrightarrow \text{GL}(\mathbb{R}^q, \mathbb{R}^q)$$

that satisfies $\Phi_0(k) = \text{Id}$ and $\Phi_\theta(k)(d\pi(\gamma(k)'(1/2))) = d\pi(v_\theta(k))$. It exists due to the homotopy lifting property. It provides us with a rotation of \mathbb{R}^q extending the tangential rotation $d\pi \circ v_\theta$.

Let χ be a cut-off function that is 1 on a neighbourhood of A and zero away from it. Consider the homotopy of curves given by

$$\alpha_\theta(k)(t) := \Phi_{\chi(k,t)\theta}(k)(\pi \circ \gamma(k)(t)).$$

We claim that $\alpha_\theta(k)$ is in fact embedded. This will indeed be the case if the support of χ is sufficiently small, since the curves $\pi \circ \gamma|_{\mathcal{O}p(A)}$ are then small embedded intervals resembling a straight line.

By construction, $\alpha_1(k)$ is tangent to $\pi \circ \nu(k)$ at $t = 1/2$. This allows us to define a further homotopy $(\alpha_\theta)_{\theta \in [1,2]}$ so that $\alpha_2(k)(t) = \pi \circ \nu(k)(t)$ for every $t \in \mathcal{O}p(\{1/2\})$. This latter homotopy may be assumed to be C^1 -small and supported in an arbitrarily small neighbourhood of A .

Over $\mathcal{O}p(A)$, we have that $\pi \circ \gamma$ and α_θ are C^0 -close and of similar length. It follows from Lemma 4.4.7 that there is a family of ε -horizontal curves β_θ lifting α_θ , that is C^0 -close to γ . Applying Lemma 4.4.5 allows us to assume that β_θ agrees with γ outside a neighbourhood of A . We can then apply Lemma 4.4.8 to β_2 in order to horizontalise. This yields a homotopy to some ε -horizontal family β_3 that close to A agrees with ν .

4.4.4 Transversalisation

In this subsection we explain the transverse analogues of the results presented in Subsection 4.4.3. We fix a distribution (M, \mathcal{D}) . We write K for a compact manifold and I for $[0, 1]$. $A \subset K \times I$ is a stratified set transverse to the second factor.

4.4.4.1 Transversalisation of almost-transverse curves

The following lemma explains that almost transverse curves can be pushed slightly to become transverse.

Lemma 4.4.11 *Suppose \mathcal{D} is of corank 1. Given a family of curves $\gamma : K \rightarrow \mathfrak{Emb}_{\mathcal{AT}}([0, 1]; M, \mathcal{D})$, there exists a C^1 -deformation*

$$\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}_{\mathcal{AT}}([0, 1]; M, \mathcal{D})$$

such that

- $\tilde{\gamma}(k, 0) = \gamma(k)$.
- $\tilde{\gamma}(k, 1)$ is transverse.
- Assume that γ is transverse along $\mathcal{O}p(\partial(K \times I))$. Then this homotopy is relative to the boundary.

Proof. The argument is carried out one adapted chart at a time. If $K \times I$ is covered by sufficiently small opens, we can pass to ODE charts (Subsection 4.3.2), where the statement is obvious and relative.

Do observe that this process may not be assumed to be relative if the starting family was purely horizontal. In fact, the argument will certainly displace the endpoint of the curves upwards.




Figure 4.6: We depict transversalisation of a horizontal curve. On a graphical model, it is easy to construct the desired transverse curve $\tilde{\gamma}(-, 1)$ (in magenta) by adding a small slope. This is not relative to the final endpoint.

Remark 4.4.12 *From this lemma it follows that there is a weak homotopy equivalence between $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})$ and the subspace of $\mathfrak{Emb}_{\mathcal{AT}}(M, \mathcal{D})$ consisting of curves that are somewhere (positively) transverse. This can be refined further to include those curves of $\mathfrak{Emb}_{\mathcal{AT}}(M, \mathcal{D})$ that are regular horizontal. We leave this as an exercise for the reader.*

4.4.4.2 Transversalisation of ε -transverse curves

The following lemma achieves the transverse condition in a neighbourhood of A .

Lemma 4.4.13 *Suppose that \mathcal{D} is of corank-1 and cooriented. Given a family $\gamma : K \rightarrow \mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}(I; M, \mathcal{D})$, there exists a family*

$$\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}_{\mathcal{T}}^{\varepsilon}(I; M, \mathcal{D})$$




Figure 4.7: Transversalisation for almost horizontal curves, relative to the boundary. The argument reduces to adding slope in an ODE model.

such that:

- i. $\tilde{\gamma}(k, 0) = \gamma(k)$.
- ii. $\tilde{\gamma}(k)(t) = \gamma(k)(t)$ if $(k, t) \in (K \times I) \setminus \mathcal{O}pA$
- iii. $\tilde{\gamma}(k, s)$ is C^0 -close to $\gamma(k)$, for all s .
- iii'. the length of $\tilde{\gamma}(k, s)$ is close to the length of $\gamma(k)$, for all s .
- iv. $\tilde{\gamma}(k, 1)$ is transverse close to A .

Proof. As in the proof of Lemma 4.4.8, we proceed inductively on the strata of A , each of which is in turn processed one simplex at a time. This reduces the proof to the analogous statement in which (M, \mathcal{D}) is a graphical model, K is \mathbb{D}^k , and the curves of γ have arbitrarily small length and image. Due to the ε -transverse condition, we have that the curves $\gamma(k)$ are then either (positively) vertical with respect to the base projection, in which case we do not need to do anything, or graphical over \mathcal{D} . In the latter case we work in an ODE model and add positive slope. This is relative in the parameter and domain.

4.4.4.3 Transversalisation of formally transverse embeddings

We also need a transversalisation statement, in the spirit of Lemma 4.4.13, that applies instead to formal transverse embeddings:

Lemma 4.4.14 *Given a family $\gamma : K \rightarrow \mathfrak{Emb}_{\mathcal{T}^f}(I; M, \mathcal{D})$, there exists a family $\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}_{\mathcal{T}^\varepsilon}(I; M, \mathcal{D})$ such that:*

- i. $\tilde{\gamma}(k, 0) = \gamma(k)$.
- ii. $\tilde{\gamma}(k)(t) = \gamma(k)(t)$ if $(k, t) \in (K \times I) \setminus \mathcal{O}pA$
- iii. $\tilde{\gamma}(k, s)$ is C^0 -close to $\gamma(k)$, for all s .
- iii'. the length of $\tilde{\gamma}(k, s)$ is close to the length of $\gamma(k)$, for all s .
- iv. $\tilde{\gamma}(k, 1)$ is transverse close to A .

Proof. We work inductively over the strata of A and inductively over the simplices of a triangulation of each stratum. This reduces the proof to a local and relative statement happening in an adapted chart. Then, the conclusion follows as in the proof of Lemma 4.4.10. Namely, the tangential homotopy given by γ can be used to rotate the velocity vectors of γ along A to make them transverse to \mathcal{D} .

4.5 Tangles

In this section we introduce *tangles*. These are particular local models for curves in the base \mathbb{R}^q of a graphical model (V, \mathcal{D}) . Upon lifting, they act as building blocks for horizontal curves. The reader should think of them as analogues of the stabilization in Contact Topology, seen in the Lagrangian projection.

Remark 4.5.1 *We have chosen the name “tangle” because they are reminiscent of tangles in 3-dimensional Knot Theory [28]. In the classical sense, a tangle $(\mathbb{D}^3, \mathcal{T})$ consists of a ball \mathbb{D}^3 with a finite number of properly imbedded disjoint arcs \mathcal{T} . This allows for the factorisation of knots into elementary pieces [13].*

Our tangles are similar: they are presented as boxes containing a homotopy of curves, with fixed endpoints. This allows us to attach them to any given family of curves in \mathbb{R}^q .

The construction of a tangle amounts to concatenating suitable flows and smoothing the resulting flowlines, taking care of the embedding condition of the lift. This is a natural approach: afterall, the bracket-generating condition explains how to produce motion in arbitrary directions by considering commutators of flows tangent to the distribution. We recommend that the reader takes a quick look at Appendix 4.9, which recaps some elementary results in this direction. In Subsection 4.5.1 we introduce some further notation about bracket-expressions and concatenating flows.

Pretangles are defined in Subsection 4.5.2. These are simply curves in \mathbb{R}^q given as flowlines of commutators of coordinate vector fields. These curves are just piecewise smooth. In order to address this, we introduce smoothening. This is done, for simple bracket-expressions, using *s-pretangles* (Subsection 4.5.2.2). We then introduce attaching models (Subsection 4.5.3) which will allow us to smooth out more complicated configurations of curves (Subsections 4.5.4 and 4.5.5). These are shown to be embedded and we explain how to insert them into existing curves. In Subsection 4.5.6 we explain how to manipulate these models to adjust the lifting of their endpoints.

Tangles are finally introduced in Subsection 4.5.7.

4.5.1 Flows

This subsection introduces some of the notation about flows that will be used later in this section.

4.5.1.1 Concatenation of flows

Let ϕ_t be a flow, possibly time-dependent. We write

$$(\phi_{a \rightarrow b})_t := \phi_{t+a} \circ \phi_a^{-1}$$

for the flow in the interval $[a, b]$, shifted so that $\phi_0^{a \rightarrow b}$ is the identity.

Fix a second flow ψ_t and real numbers $a < b$ and $c < d$. Then, we define the **concatenation** of $\phi_{a \rightarrow b}$ and $\psi_{c \rightarrow d}$ to be:

$$(\phi_{a \rightarrow b} \# \psi_{c \rightarrow d})_t := \begin{cases} (\phi_{a \rightarrow b})_t & t \in [0, b - a], \\ (\psi_{c \rightarrow d})_{t - (b - a)} \circ (\phi_{a \rightarrow b})_{b - a} & t \in [b - a, (b - a) + (d - c)]. \end{cases}$$

This is a time-dependent flow that is piecewise smooth in t , due to the switch at $t = b - a$.

In general, given flows $(\phi_i)_{i=1, \dots, k}$ and real numbers $(a_i < b_i)_{i=1, \dots, k}$ we can iterate the previous construction:

$$\#_{i=1, \dots, l} (\phi_i)_{a_i \rightarrow b_i} := (\#_{i=1, \dots, l-1} (\phi_i)_{a_i \rightarrow b_i}) \# (\phi_l)_{a_l \rightarrow b_l}.$$

4.5.1.2 Generalised bracket-expressions

We now generalise the formal bracket-expressions from Subsection 1.5.0.1. The aim is to consider iterates of formal bracket-expressions.

Definition 4.5.2 We say that the pair (a, k) , written as $a^{\#k}$, depending on the variable a and the integer k , is a **generalised bracket expression** of length 1. Similarly, we say that the expression $[a_1, a_2]^{\#k}$, depending on the variables a_1, a_2 and the integer k , is a generalised bracket expression of length 2. Inductively, we define a generalised bracket expression of length n to be an expression of the form

$$[A(a_1, \dots, a_j), B(a_{j+1}, a_n)]^{\# \ell}, \quad 0 < j < n$$

with A and B generalised bracket expressions of lengths j and $n - j$, respectively.

4.5.2 Pretangles and s -pretangles

We now fix a graphical model (V, \mathcal{D}) . All the constructions in this section take place within it. We write $\pi : V \rightarrow \mathbb{R}^q$ for its projection to the base. The framing reads $\{X_1, \dots, X_n\}$, with $\{X_1, \dots, X_q\}$ a framing of \mathcal{D} lifting the coordinate framing $\{\partial_1, \dots, \partial_q\}$ of \mathbb{R}^q . We write ϕ_i for the flow of ∂_i , here $i = 1, \dots, q$. The flow of X_i is denoted by Φ_i .

4.5.2.1 Pretangles

The following construction produces time-dependent flows that are iterates of a given commutator:

Definition 4.5.3 Let $A = [a, b]^{\#m}$ be a generalised bracket expression of length 2. Let ϕ and ψ be flows. We define

$$A(s) := \left(\left(\phi_{0 \rightarrow \frac{s}{\sqrt{m}}} \right) \# \left(\psi_{0 \rightarrow \frac{s}{\sqrt{m}}} \right) \# \left(\phi_{0 \rightarrow \frac{s}{\sqrt{m}}}^{-1} \right) \# \left(\psi_{0 \rightarrow \frac{s}{\sqrt{m}}}^{-1} \right) \right)^{\#m}.$$

The superindex $\#m$ denotes concatenating the described commutator m times.

We can introduce an analogous definition for bracket expressions of greater length, inductively:

Definition 4.5.4 Let $B = [a_1, A(a_2, \dots, a_l)]^m$ a generalised bracket expression. Consider flows $(\phi_i)_{i=1, \dots, l}$. Then we denote:

$$B(s) := \left(\left(\phi_{0 \rightarrow \frac{s}{\sqrt{m}}} \right) \# A(s/\sqrt{m}) \# \left(\psi_{0 \rightarrow \frac{s}{\sqrt{m}}}^{-1} \right) \# A(s/\sqrt{m})^{-1} \right)^{\#m}.$$

The following is the main definition of this subsection:

Definition 4.5.5 Let A be a generalised bracket-expression with ϕ_1, \dots, ϕ_ℓ as inputs. An integral curve, depending on the parameter s , of the flow defined by the expression $A(s)$ is called a **pretangle**. We will denote such a curve by $\gamma^{A(s)}$.

4.5.2.2 S-pretangles

We will introduce the notion of **S-pretangles**, where the “s” stands for “smooth”.

We can describe a way of smoothing the corners where the previously defined pretangles failed to be smooth. We will define two different ways of smoothing a corner. Define first the following time-dependent vector field:

$$Y_t^{(i,j),\delta} := \frac{\delta-t}{\delta} \cdot X_i + \frac{t}{\delta} \cdot X_j, \quad t \in [0, \delta].$$

Denote by $s_{0 \rightarrow \delta}^{i,j}$ the flow associated to the vector field $Y_t^{(i,j),\delta}$. Note that the concatenation of flows $\phi_{0 \rightarrow \eta-\delta/2}^i \# s_{0 \rightarrow \delta}^{i,j} \# \phi_{\delta/2 \rightarrow \tau}^j$ or, in short, $\phi_{0 \rightarrow \eta-\delta/2}^i \#_{i,j}^\delta \phi_{\delta/2 \rightarrow \tau}^j$, can be made C^∞ -close to $\phi_{0 \rightarrow \eta}^i \# \phi_{0 \rightarrow \tau}^j$ by taking δ small enough:

$$\phi_{0 \rightarrow \eta-\delta/2}^i \#_{i,j}^\delta \phi_{\delta/2 \rightarrow \tau}^j \xrightarrow[\delta \rightarrow 0]{\|\cdot\|_{C^\infty}} \phi^i \# \phi^j.$$




Figure 4.8: Schematic description of $\phi_{0 \rightarrow t-\delta/2}^i \#_{i,j}^\delta \phi_{\delta/2 \rightarrow t}^j$. The flow $s_\delta^{i,j}$ provides a way of smoothing the previously defined concatenation of two given vector field flows.

The flows $s_\delta^{i,j}$ play the role of smoothing the concatenation of two vector field flows when concatenated in between. Indeed, note that $\phi_{0 \rightarrow \eta-\delta/2}^i \#_{i,j}^\delta \phi_{\delta/2 \rightarrow \tau}^j$ is a C^∞ -flow.

Let us introduce now a different way of smoothing a corner. Denote by $\delta' := \delta/4 - \delta/50$, $\delta'' := \delta/25$ and $\delta''' = \delta/50$. Consider the following flow (see Figure 4.9):

$$d_\delta^{j,i} := \phi_{0 \rightarrow \delta'}^j \#_{j,i}^{\delta''} \phi_{\delta''' \rightarrow \delta'}^i \#_{i,-j}^{\delta''} \phi_{\delta''' \rightarrow \delta'}^{-j} \#_{-j,-i}^{\delta''} \phi_{\delta''' \rightarrow \delta/4}^{-i}$$

We now define smooth pretangles $\gamma_{SPT}^{t,\delta}$ for two given vector field flows ϕ_t^i, ϕ_t^j , as the curve in Figure 4.11, defined as a pretangle for $[\phi^i, \phi^j]$ whose corners have been smoothed by using the flows




Figure 4.9: Schematic description of $\phi_{0 \rightarrow \sqrt{t}-\delta/2}^j \# d_{\delta}^{j,i} \# \phi_{\delta/2 \rightarrow \sqrt{t}}^i$. The flow $d_{\delta}^{j,i}$ provides a essentially different way of smoothing the previously defined concatenation of two given vector field flows.

$\#_{\pm i, \pm j}^{\delta}$ and $d_{\delta}^{j,i}$ (see the upright corner). A precise formula for the curve can be given. Indeed, $\gamma_{\mathcal{SP}\mathcal{T}}^{t,\delta}$ is an integral curve of the flow:

$$\mathcal{SP}\mathcal{T}_{[\phi^i, \phi^j]}^{t,\delta} := \phi_{0 \rightarrow t-\delta/2}^i \#_{i,j}^{\delta} \phi_{\delta/2 \rightarrow t-\delta/2}^j \# d_{\delta}^{j,-i} \# \phi_{\delta/2 \rightarrow t-\delta/2}^{-i} \#_{-i,-j}^{\delta} \phi_{\delta/2 \rightarrow t-3\delta/2}^{-j} \#_{-j,i}^{\delta} \# \phi_{\delta/2 \rightarrow t-\delta}^i \#_{i,-j}^{\delta} \#_{-j,i}^{\delta}$$

We refer to δ as the **smoothing parameter of the s-pretangle** $\mathcal{SP}\mathcal{T}_{[\phi^i, \phi^j]}^{t,\delta}$. (See Figure 4.11, where an integral curve of such a flow is depicted inside the grey box).

4.5.3 Attaching models

As we saw earlier, pretangles can be interpreted as a local model that can be attached to a family of curves in the base of a graphical model in order to quantitatively control the endpoints of the lift (see Proposition 4.5.18). Nonetheless, we face two fundamental problems when we try to do that: these curves are not smooth. Furthermore, they are not (topologically) embedded and it is not readily apparent whether their lifts are embedded.

We now introduce some alternate models by carefully modifying our prior constructions. These models will depend on certain small “smoothing parameters” and will converge to pretangles, in the C^0 -norm, as we make these parameters tend to zero.

The general definition reads:

Definition 4.5.6 Let $\gamma : I \rightarrow \mathbb{R}^q$ be a curve that is integral for a coordinate vector field X_i in the adapted frame. We call an **attaching model** with axis ∂_i the choice of:

- i) a size $\eta > 0$,
- ii) two attaching points $p_1 = \gamma(t_1), p_2 = \gamma(t_2)$ at distance $d_g(p_1, p_2) = 2\eta$,
- iii) a hypercube $\mathcal{B}(\eta) \subset \mathbb{R}^q$, called the box of the model, of side 2η so that p_1 and p_2 are in opposite faces of \mathcal{B} .
- iv) a curve β with endpoints p_1, p_2 satisfying:
 - iv.a) its image lies inside $\mathcal{B}(\eta)$.
 - iv.b) The curve $\tilde{\gamma} : I \rightarrow \mathbb{R}^q$ defined as $\tilde{\gamma}|_{I \setminus [t_1, t_2]} = \gamma|_{I \setminus [t_1, t_2]}$, $\tilde{\gamma}|_{[t_1, t_2]} = \beta$ is continuous. If it has C^r regularity, we say that the model is C^r -regular.




Figure 4.10: Schematic depiction of an **attaching model** for certain choice of $\eta > 0$, attaching points p_1, p_2 , a hypercube $\mathcal{B}(\eta)$ and a curve β .

4.5.3.1 Pretangle models

A concrete instance of Definition 4.5.6 to be used in the next section reads:

Definition 4.5.7 Let A be a generalised bracket expression of the form $A(\phi_i, \dots, \phi_n) = [\phi_i, B(\phi_\ell, \dots, \phi_n)]^{\#k}$ with inputs flows ϕ_i, \dots, ϕ_n , and where B is a bracket expression of smaller length. A **pretangle model** associated to A is an attaching model where the curve $\beta : [t_1, t_2] \rightarrow \mathbb{R}^q$ inside the box satisfies:

- it is a pretangle for $t \in (t_1 + \epsilon, t_2 - \epsilon)$,
- it coincides with the straight segment in the direction ∂_i joining p_1, p_2 for $t \in (t_1, t_1 + \epsilon) \cup (t_2 - \epsilon, t_2)$.

We call length of the pretangle model the length of the pretangle inside the box.




Figure 4.11: Pretangle model.

4.5.4 Length-2 tangle models

We now introduce the building blocks of the main objects of interest in the section. We present first a specific type of attaching model that we call basic length 2 tangle model. These are meant to be better behaved than pretangle models, whose regularity is only C^0 .

Definition 4.5.8 Let A be a generalised bracket expression of the form $A(\phi_i, \dots, \phi_n) = [\phi_i, B(\phi_\ell, \dots, \phi_n)]^{\#k}$ with inputs flows ϕ_i, \dots, ϕ_n , and where B is a bracket expression of smaller length. A **length-2 base tangle model** associated to A is the attaching model described by Figure 4.12.

The curve $\beta_\delta : [t_1, t_2] \rightarrow \mathbb{R}^q$ inside the box is immersed and satisfies:

- it is a smooth pretangle $\mathcal{SPT}_{[\phi^i, \phi^j]}^{\mu, \delta}$ for $t \in (t_1 + \epsilon, t_2 - \epsilon)$,
- it is δ -close to the straight segment in the direction ∂_i joining the points p_1, p_2 for $t \in (t_1, t_1 + \epsilon) \cup (t_2 - \epsilon, t_2)$,
- the size of the box of the attaching model is $\mu + 2\delta$.

We say that δ is the smoothing parameter of the model.

Whenever it is clear from the context we will refer as the tangle to the curve β in the tangle model.




Figure 4.12: Length 2 base tangle model. The curve inside the box is a smooth pretangle $\mathcal{SPT}_{[\phi^i, \phi^j]}^{\mu, \delta}$ for $t \in (t_1 + \epsilon, t_2 - \epsilon)$ and coincides with the straight segment in the direction ∂_i joining the points p_1, p_2 elsewhere.

Remark 4.5.9 Note that any length 2 pretangle model can be C^0 -approximated by a length 2 base tangle model by taking the smoothing parameter δ small enough.

4.5.4.1 Birth homotopy for length-2 base tangle models

We first state the following result:

Proposition 4.5.10 Assume $[X_i, X_j = B(X_1, \dots, X_\ell)] = X_z$. Denote by α_z be the covector dual to X_z .

Then, any curve $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^q$ enclosing area A in the plane $\mathbb{R}^2 = \langle \partial_i, \partial_j \rangle$ lifts to \mathcal{D} as a curve $\tilde{\gamma} : [0, 1] \rightarrow \mathbb{R}^n$ satisfying

$$\int_{\gamma} \alpha_z = A (1 + O(r)).$$

Proof. Denote by $\Gamma_{(\tilde{\gamma}(1), \tilde{\gamma}(0))}$ the oriented segment connecting the points $\tilde{\gamma}(1)$ and $\tilde{\gamma}(0)$. Denote by $\beta := \tilde{\gamma} \# \Gamma_{(\tilde{\gamma}(1), \tilde{\gamma}(0))}$ the concatenation of the curves $\tilde{\gamma}$ and $\Gamma_{(\tilde{\gamma}(1), \tilde{\gamma}(0))}$. Note that

$$\int_{\beta} \alpha_z = - \int_{\Gamma_{(\tilde{\gamma}(1), \tilde{\gamma}(0))}} \alpha_z$$

and, thus, this integral measures the difference of the ∂_z -coordinate values of the points $\tilde{\gamma}(1)$ and $\tilde{\gamma}(0)$.

Consider a topological disk $\mathcal{D}_{\tilde{\gamma}}$ bounded by β and whose boundary gets projected to γ in the projection onto the plane $\langle \partial_i, \partial_j \rangle$. By Stokes' theorem,

$$\int_{\beta} \alpha_z = \int_{\mathcal{D}_{\tilde{\gamma}}} d\alpha_z.$$

By Cartan's formula we have that

$$d\alpha_z(X_i, X_j) = \alpha_z([X_j, X_i]).$$

Thus, if we particularize this equation at the point $p \in M$, we get that

$$d\alpha_z(p)(X_i(p), X_j(p)) = 1,$$

and it vanishes when evaluated at any other combination of two elements of the framing associated to the coordinate chart. Thus, the 2-form $d\alpha$ coincides with $dx_i \wedge dx_j$ in the origin at the level of 0-jets. As an application of Taylor's Remainder Theorem we get that

$$\int_{\mathcal{D}_{\tilde{\gamma}}} d\alpha_z = \int_{\mathcal{D}_{\tilde{\gamma}}} dx_i \wedge dx_j + O(r) = A + A \cdot O(r),$$

yielding the claim. □

Our goal is to present a homotopy that introduced a length-2 base tangle. The following statement holds:

Proposition 4.5.11 (Birth homotopy for length-2 base tangle models) *Let X_i, X_j be two elements in the adapted framing such that $[X_i(p), X_j(p)] = X_\ell(p)$. Let $\gamma : [-\delta, \delta] \rightarrow \mathbb{R}^q$ be a horizontal curve in a graphical model.*

There exists a homotopy of embedded horizontal curves $(\gamma^u)^{u \in [0, 1]}$ such that:

- i) $\gamma^0 = \gamma$.
- ii) $\pi \circ \gamma^1|_{[-\delta/2, \delta/2]}$ is endowed with a length-2 base tangle model associated to the generalised bracket expression $[\phi_i, \phi_j]$.
- iii) $\pi \circ \gamma^u(t) = \pi \circ \gamma(t)$ for $t \in \mathcal{O}p(\{-\delta, \delta\})$ and all $u \in [0, 1]$.

Property iii) guarantees that this homotopy, when projected into the base, is relative to both endpoints. This, in particular, implies that the lifted homotopy γ^u through horizontal curves is also relative to the starting point.

Proof. We construct the homotopy in the base; i.e. we will define $\pi \circ \gamma^u$ and, being the lifting onto the connection unique, the claim will follow.

Since iterated models are constructed iteratively on the base length 2 model, it suffices to show the result for that the latter. We first locally homotope $\pi \circ \gamma$ to an integral curve for X_i in the base, any segment around $\gamma(t_0)$ is as described by the first frame in Figure 4.13. Consider the local isotopy of immersed cuves in the base described by the Figure 4.13.




Figure 4.13: Birth homotopy for a length 2 base tangle model $\mathcal{T}_{[\phi^i, \phi^j]}^t$.

The first three depicted frames in the movie correspond to the isotopy $(\pi \circ \gamma^u)^{u \in [0, \frac{1}{2}]}$, while the fourth one completes it to $(\pi \circ \gamma^u)^{u \in [0, 1]}$.

Points *i*), *ii*) and *iii*) readily follow from the isotopy taking place in the projected plane $\langle \partial_i, \partial_j \rangle$. Therefore all we have to check is that embeddedness holds when we lift the curve to \mathcal{D} . We will verify that any pair of intersection points taking place in the base (at most two pairs, depending on the value of $u \in [0, 1]$) lift to different points upstairs.

Note that this fact can be achieved trivially if the rank of the distribution \mathcal{D} is greater than 2, since we can use an additional coordinate ∂_z in order to perform the homotopy while remaining embedded already in the base. This fact is depicted in Figure 4.14, where it is shown how to avoid any crossing in the 3-plane $\langle \partial_i, \partial_j, \partial_z \rangle$ during the isotopy.




Figure 4.14: Increase of the additional coordinate ∂_z during the homotopy in order to achieve embeddedness in distributions \mathcal{D} of rank greater than 2.

So, let us assume that the distribution is of rank 2 and, thus, because of $\dim(M) > 3$ and the bracket-generating condition, we can assume that either $[X_i(p), X_\ell(p)] = X_z(p)$ or $[X_i(p), X_\ell(p)] =$

$X_z(p)$, where X_z is some other element in the adapted frame. Without loss of generality, we assume $[X_i(p), X_\ell(p)] = X_z(p)$.

Let us denote by $(\pi \circ \gamma^u(t_1), \pi \circ \gamma^u(t_2))$ the 1-parametric family of pairs of points corresponding to the upper-right autointersection in the homotopy in Figure 4.13. By Proposition 4.5.10 the difference in the values of the ∂_ℓ -coordinate between the liftings of the points $\pi \circ \gamma^u(t_1)$ and $\pi \circ \gamma^u(t_2)$ is $A^u(1 + O(r))$, where A^u is the area enclosed by the curve $\pi \circ \gamma^u|_{[t_1, t_2]}$ in the plane $\langle \partial_i, \partial_j \rangle$. Therefore for a sufficiently small choice of $r > 0$, $A^u(1 + O(r))$ is a positive number.

Denote by $(\pi \circ \gamma^u(t'_1), \pi \circ \gamma^u(t'_2))$ the 1-parametric family of pair of points corresponding to the other autointersection in Figure 4.13. If the lifting $\gamma|_{[t'_1, t'_2]}$ of the curve $\pi \circ \gamma|_{[t'_1, t'_2]}$ projects onto the plane $\langle \partial_i, \partial_\ell \rangle$ as an opened curve then we are done, since this means that the ∂_z -coordinates of the points $\gamma^u(t'_1)$ and $\gamma^u(t'_2)$ are different.

Otherwise, we get a closed loop that encloses area B^u in the plane $\langle \partial_i, \partial_\ell \rangle$ and that implies that, again by Proposition 4.5.10, $\pi \circ \gamma(t'_1)$ and $\pi \circ \gamma(t'_2)$ differ an ammount of $B^u(1 + O(r))$ in the ∂_ℓ -coordinate. We conclude then that the ∂_ℓ -coordinates of the liftings of both points are different by the same argument as above.

Properties *i*), *ii*) and *iii*) are satisfied by construction.




Figure 4.15: When we look at the projection of the curve into the plane $\langle \partial_i, \partial_\ell \rangle$ we get a closed loop that encloses area B^u .

Remark 4.5.12 Note that we can inductively choose two attaching points q_1, q_2 inside a length 2 base tangle model and a box whose boundary intersects the curve only at q_1, q_2 as in Figure 4.16. This way, we can insert to the given model another length 2 base tangle model (see Figure 4.16) which is 2δ -close, in the C^0 -norm, to the given one. We call ***k times nested length 2 tangle curves*** with smoothing parameter δ to the curve obtained after repeating this process k times.




Figure 4.16: Length 2 base tangle model on the left with two marked attaching points q_1, q_2 and a choice of box for inserting another length 2 base tangle model inside. On the right, a 2 times nested length 2 tangle curve.

4.5.4.2 Iterated length-2 tangle models

We introduce now a variation on the previously defined model:

Definition 4.5.13 A k -times iterated length-2 tangle model associated to the generalised bracket expression $[\phi_i, \phi_j]^{\#k}$ with inputs ϕ_i, ϕ_j is an attaching model described by Figure 4.16. Note that the curve inside the box satisfies:

- it coincides with the straight segment in the direction p_1, p_2 for $t \in (t_1, t_1 + \epsilon) \cup (t_2 - \epsilon, t_2)$,
- the curve β inside the box is a k times nested length 2 tangle curve with smoothing parameter δ .
- the size of the box of the attaching model is $\mu + 2\delta$,

Note that a length-2 base tangle model is just a 1-time iterated length 2 tangle model.

Remark 4.5.14 A key remark is the following one: note that as $\tau, \delta \rightarrow 0$, any k -times iterated length 2 tangle model converges to an also k -times iterated pretangle model in the C^0 -norm.

Their birth homotopy is explained in the following proposition:

Proposition 4.5.15 Let X_i, X_j be two elements in the adapted framing such that $[X_i(p), X_j(p)] = X_\ell(p)$. Let $\gamma : [-\delta, \delta] \rightarrow \mathbb{R}^q$ be a horizontal curve in a graphical model. There exists a homotopy of embedded horizontal curves $(\gamma^u)^{u \in [0,1]}$ such that:

- $\gamma^0 = \gamma$.
- $\pi \circ \gamma^1|_{[-\delta/2, \delta/2]}$ is endowed with a k -times iterated length-2 tangle model associated to the generalised bracket expression $[\phi_i, \phi_j]^{\#k}$.
- $\pi \circ \gamma^u(t) = \pi \circ \gamma(t)$ for $t \in \mathcal{O}p(\{-\delta, \delta\})$ and all $u \in [0, 1]$.

Proof. The claim follows by inductively applying Proposition 4.5.11 starting from the outermost curve.

4.5.5 Tangle models of higher length

Consider a bracket expression of the form $A(\phi_1, \dots, \phi_m) = [\phi_i, B(\phi_r, \dots, \phi_\ell)]$ with flows ϕ_1, \dots, ϕ_m as inputs. A **length- m base tangle model** is an attaching model associated to A . It is described inductively on its length, which is the length of A . The inductive step is described in Figure 4.17.




Figure 4.17: Length $N > 2$ base tangle model. The grey boxes represent tangle models of size $\rho > 0$ of one unit smaller length. The real numbers τ, ρ are called smoothing parameters associated to the inductive step and are all greater than all the smoothing parameters defined in previous steps. The direction ∂_j is associated to the coordinate flow ϕ_j , which is the first entry appearing in the generalised bracket expression A , different from ϕ_i .

The grey boxes in Figure 4.17 represent tangle models of size $\eta > 0$ associated to the expression $B(\phi_r, \dots, \phi_\ell)$. The direction ∂_j is associated to the coordinate flow ϕ_j , which is the first entry appearing in the generalised bracket expression A , different from ϕ_i . All the model, except for the pieces inside the grey boxes, is described in the plane ∂_i, ∂_j . The real numbers τ, ρ are called smoothing parameters associated to the inductive step and are all greater than all the smoothing parameters defined in previous steps.

For length- N tangle models, we can also choose two attaching points q_1, q_2 and a box (2τ -close in the C^0 -norm to the outermost box) in such a model and iterate the construction, in the same fashion as in length 2, thus constructing another model over the given one.

Remark 4.5.16 *Note that as all the smoothing parameters of any k -times iterated length n tangle model tend to zero, the model converges to a pretangle model in the C^0 -norm. Indeed, it is clear that the result is true for length 2 models (See Remark 4.5.14). On the other hand, assuming that the grey boxes in Figure 4.18 contained pretangle models instead of tangle models, observe that as τ and ρ tend to 0, the whole construction would converge to a pretangle model. Combining both facts the claim follows. We call **the pretangle model associated to the tangle model** to such a pretangle model.*




Figure 4.18: Length $N > 2$ base tangle model on the left with two marked attaching points q_1, q_2 and a choice of box for inserting another length N base tangle model inside. On the right, a 2 times iterated length N tangle curve.

4.5.5.1 Birth homotopy for higher length tangle models

The birth homotopy is given by the following result:

Proposition 4.5.17 *Consider a bracket expression $A(\phi_1, \dots, \phi_m)$ with inputs the flows ϕ_1, \dots, ϕ_m . Consider*

$$\pi \circ \gamma : [-\delta, \delta] \rightarrow \mathbb{R}^q$$

a family of curves given by a horizontal lift γ .

Then, there exists a homotopy of embedded horizontal curves $(\gamma^u)_{u \in [0,1]}$ such that:

- i) $\gamma^0 = \gamma$.*
- ii) $\pi \circ \gamma^1|_{[-\delta/2, \delta/2]}$ is endowed with a k -times iterated length- n tangle model associated to the generalised bracket expression A .*
- iii) $\pi \circ \gamma^u(t) = \pi \circ \gamma(t)$ for $t \in \mathcal{O}p(\{-\delta, \delta\})$ and all $u \in [0, 1]$.*

Proof. It is easy to construct the homotopy in the base by defining $\pi \circ \gamma^u$. The length n case iterated model (Figure 4.18) reduces to the non-iterated model (Figure 4.17) since the birth homotopy for the former can be constructed inductively by using the birth homotopy of the latter.

Note that the birth homotopy for the non-iterated length n model can be constructed inductively. Indeed, assume we already now how to introduce length $n - 1$ -models of sufficiently small size at any given point of a curve and proceed as follows. We first homotope the given curve in the box to the curve in Figure 4.17 (but omitting the grey boxes). Now we perform the birth homotopies for the $n - 1$ tangle models in the grey boxes and we are done.

The base case corresponds to the case of length 2-tangle models, which we already explained how to do (See Proposition 4.5.11).

4.5.6 Area isotopy

Associated to a 2-length tangle realizing the bracket $[X_i, X_j]$, we have a notion of **increasing** or **decreasing** its “area” just by geometrically increasing or decreasing the area enclosed by the tangle in the $\langle \partial_i, \partial_j \rangle$ plane. In a sense, the increasing/decreasing of such area parametrizes (**controls**) the increase of the ∂_z coordinate of the lifted curve (where $\partial_z = [X_i, X_j]$ and X_k the lif of ∂_k). We will extend this notion of area controlling for higher length tangle expressions.

Assume $\partial_z(p) = A(X_1, \dots, X_n)(p)$ in a graphical model based at $p \in M$ with $\text{len}(A) = \lambda$. The following statement allows us to estimate how a pretangle controls the endpoint:

Proposition 4.5.18 *A pretangle $\gamma^{A(\mu)}$ into the graphical model associated to the generalised bracket-expression $A(\phi_s^1, \dots, \phi_s^n)$ lifts to the distribution as a curve $\tilde{\gamma} : [0, 1] \rightarrow \mathbb{R}^n$ where the difference between the endpoints $\gamma(1) - \gamma(0)$ is $\mu^\lambda (\partial_z + O(r))$.*

Proof. By Proposition 4.9.9 (Subsection 4.9) the following equality holds

$$A(\varphi_t^{X_1}, \dots, \varphi_t^{X_\lambda}) = \varepsilon_{t^\lambda} \circ \phi_{t^\lambda}^{A(X_1, \dots, X_\lambda)}.$$

On one hand we have that $\partial_z(p) = A(X_1, \dots, X_n)(p)$ and, thus, by Taylor’s Remainder Theorem we have that for nearby points $q \in \mathcal{O}p(p)$ the following equality holds

$$A(X_1, \dots, X_n)(q) = \partial_z + O(r)$$

The notation $\partial_z + O(r)$ denotes a ∂_z plus some vector of size $O(r)$; i.e. not necessarily collinear to ∂_z . Combining both inequalities:

$$A(\varphi_t^{X_1}, \dots, \varphi_t^{X_\lambda})(q) = \varepsilon_{t^\lambda} \circ \phi_{t^\lambda}^{\partial_z + O(r)} = t^\lambda \cdot (\partial_z + O(r)),$$

where the error associated to ε_{t^λ} has been incorporated by $O(r)$. Now taking $t = \mu$ implies the claim.

The lifting of a curve into a connection does not depend on its reparametrization. Therefore,

Definition 4.5.19 *By Proposition 4.5.18, we have associated to a pretangle $A(\mu)$ a real number μ^λ which is independent of its reparametrization and that we call its **total area**.*

Let $\gamma : I \rightarrow \mathbb{R}^q$ be a curve in the base of the graphical model equipped with a pretangle model associated to A with attaching points $\gamma(t_1) = q_1$ and $\gamma(t_2) = q_2$

Corollary 4.5.20 *The curve γ equipped with the pretangle model lifts to the distribution as a curve $\tilde{\gamma} : [0, 1] \rightarrow \mathbb{R}^n$ where the difference between the endpoints $\tilde{\gamma}(1)$ and $\tilde{\gamma}(0)$ is $\mu^\lambda (1 + O(r)) \cdot \partial_z$.*

Definition 4.5.21 *We define the **total area of a pretangle model** as the total area of the pretangle in the model.*

4.5.6.1 Area isotopy for pretangle models

We now describe a way of increasing/decreasing the total area of a given pretangle model by appropriately manipulating it.

Assume that the generalised bracket expression A is of the form $A(\phi_i, \dots, \phi_n) = [\phi_i, B(\phi_\ell, \dots, \phi_n)]^{\#k}$ with inputs flows ϕ_i, \dots, ϕ_n , where B is a bracket expression

of smaller length. Let $r > 0$ be the total number of times that ϕ_i appears as an entry in the expression $A(\phi_i, \dots, \phi_n)$. Consider a pretangle model \mathcal{PM}_A associated to A with box a hypercube $\mathcal{B} \subset \mathbb{R}^q$ and total area μ^λ .

Take coordinates in \mathbb{R}^q in such a way that the hypercube \mathcal{B} has its center at the origin. Take a bump function $\psi : \mathbb{R}^q \rightarrow [0, 1]$ in \mathbb{R}^q with support $\mathcal{Op}(M^{1/r} \cdot \mathcal{B})$, where $M^{1/r} \cdot \mathcal{B}$ denotes the hypercube with side $M^{1/r}$ times the one of \mathcal{B} . M denotes the size of the maximal box onto which we can extend \mathcal{B} .

Definition 4.5.22 *We define the Area isotopy $(\Psi^u)_{u \in [1, M]}$ of the pretangle model \mathcal{PM}_A as:*

$$(\Psi^u)_{u \in [0, M]} : \mathbb{R}^q \longrightarrow \mathbb{R}^q$$

$$x = (x_1, \dots, x_i, \dots, x_q) \longmapsto (x_1, \dots, \psi_{\mathcal{B}}(x) \cdot (u)^{1/r} x_i, \dots, x_q)$$

The upcoming proposition explains how the Area isotopy $\Psi^u \circ \gamma$ applied to the curve γ equipped with the pretangle model \mathcal{PM}_A behaves with respect to endpoints.

It follows from a combination of Proposition 4.5.18 together with a routinary inductive argument based on Proposition 4.5.10:

Proposition 4.5.23 (Endpoint of lifted pretangle models under the Area isotopy) $\Psi^u \circ \gamma$ lifts to the distribution as a curve $\tilde{\gamma}^u : [0, 1] \rightarrow \mathbb{R}^n$ where the difference between the lifts of the attaching points $\tilde{\gamma}(t_1)$ and $\tilde{\gamma}(t_2)$ is $(u \cdot \mu^\lambda) (\partial_z + O(r))$.

4.5.6.2 Area isotopy for tangle models

We have explained so far how the Area isotopy controls the displacement of the lifted endpoints of pretangle models. Nonetheless, we can extend the discussion to tangle models.

Recall that as all the smoothing parameters of any tangle model tend to zero, the model converges to a pretangle model in the C^0 -norm (see Remark 4.5.16). We call the **total smoothing parameter** of a given tangle model to the real number $\delta := \max_i \{\delta_i\}$, where $\{\delta_i\}_i$ is the set of all smoothing parameters of a given tangle model. Then, we have that in the limit case where $\delta \rightarrow 0$, tangle models converge to pretangle models in the C^0 -norm.

Definition 4.5.24 *We define the **total area of a tangle model** as the total area of its associated pretangle model.*

We can analogously define the Area isotopy for a tangle model:

Definition 4.5.25 *We define the **Area isotopy for a tangle model** as the Area isotopy of its associated pretangle model.*

As a consequence of the whole discussion until this point we deduce the following key result:

Proposition 4.5.26 (Endpoint of lifted tangle models under the Area isotopy) *Let $\gamma : I \rightarrow \mathbb{R}^q$ be a curve equipped with a tangle model associated to A with attaching points $\gamma(t_1) = q_1$ and $\gamma(t_2) = q_2$. Then, $\Psi^u \circ \gamma$ lifts to the distribution as a curve $\tilde{\gamma}^u : [0, 1] \rightarrow \mathbb{R}^n$ where the difference between the lifts of the attaching points $\tilde{\gamma}^u(t_1)$ and $\tilde{\gamma}^u(t_2)$ is $(u \cdot \mu^\lambda) (\partial_z + O(r)) (1 + O(\delta))$.*

Remark 4.5.27 *Note that, as the radius r of the graphical model gets close to zero, the quantity $(1 + O(r))$ becomes close to 1. The same phenomenon holds when the total smoothing parameter*




Figure 4.19: Schematic description of a length 3 tangle model. The transition from the second frame to the first one represents the isotopy Ψ^u acting on the model. The long segments X_i get expanded while the length 2 subtangle models do not get expanded nor shrunk.

δ of the tangle model gets close to zero, $(1 + O(\delta))$ becomes close to 1. Therefore, in the limit, adjusting the endpoints of lifted tangle models is practically equivalent to adjusting the endpoints of the corresponding lifted pretangle models.

4.5.7 Tangles

Let X_j be an element in the framing of TV , with A a bracket-expression generating it. We assume that A is of the form $[\phi_i, B(-)]^{\#k}$. Consider the following data:

- a size $R > 0$,
- a curve $\gamma : [0, 1] \rightarrow \mathbb{R}^q$ parallel to ∂_i
- two attaching points $\gamma(t_1) = q_1$ and $\gamma(t_2) = q_2$

Then there exists a tangle model \mathcal{TM} , associated to the bracket-expression A , and endowed with:

- Total area μ^λ (determined by R).
- Smoothing parameter $\delta > 0$.
- A birth homotopy for \mathcal{TM} , parametrised by $\theta \in [0, 1]$ and given by Proposition 4.5.17.
- An area isotopy (Definition 4.5.22) parametrised as $d \mapsto \Psi^{(d/\mu^\lambda)}$.
- An upper bound $h := M\mu^\lambda$ for the area isotopy.

We now put together all the ingredients introduced in this section:

Definition 4.5.28 An X_j -tangle is a family of curves

$$\mathcal{T} : (0, h] \times [0, 1] \longrightarrow \mathfrak{Imm}([0, 1]; \mathbb{R}^q)$$

given by the previously introduced tangle model \mathcal{TM} . It is parametrised by

- an **estimated-displacement** $d \in (0, h]$ that governs the area isotopy,

- the *birth-parameter* $\theta \in [0, 1]$.

The number h is called the *maximal-displacement* of the tangle.

4.5.7.1 Error in the displacement

The following statement bounds the difference between the estimated-displacement and the actual displacement of the endpoint upon lifting a tangle:

Lemma 4.5.29 *Lift the tangle \mathcal{T} using Lemma 4.4.3. Then the following estimate holds:*

$$\text{lift}(\mathcal{T}(d, 1))(1) = \text{lift}(\mathcal{T}(0, 1))(1) + (0, \dots, 0, d, 0, \dots, 0) + O(r + \delta)d,$$

where r is the radius of the graphical model and δ is the smoothing parameter of the tangle.

The proof is immediate from Proposition 4.5.26.

4.6 Controllers

In the previous Section 4.5 we introduced the notion of tangle. The purpose of a tangle is to displace the endpoint of a horizontal curve in a given direction. In this section we introduce the notion of *controller*. This is a sequence of tangles, located one after the other, in order to be able to control the endpoint of a curve fully.

In Subsection 4.6.1 we talk about general (finite-dimensional) families of horizontal curves. The goal is to discuss their endpoint map and make quantitative statements about their controllability. We then particularise to controllers (Subsection 4.6.2), which are specific families of horizontal curves built out of tangles. The process of adding a controller to a horizontal or ε -horizontal curve is explained in Subsection 4.6.3.

4.6.1 Controllability

We introduced the notion of regularity in Subsection 4.1.1. This meant that the endpoint map of the horizontal curve under consideration was an epimorphism, which should be understood as a form of infinitesimal controllability (every infinitesimal displacement of the endpoint can be followed by a variation of the curve). In this Subsection we pass from infinitesimal to local.

4.6.1.1 Controlling families

For our purposes we need to work on a parametric setting. Fix (M, \mathcal{D}) , a manifold endowed with a bracket-generating distribution, and a compact fibre bundle $E \rightarrow K$. We write E_k for the fibre over $k \in K$.

Given a family of horizontal curves

$$\gamma : E \longrightarrow \mathfrak{Maps}([0, 1]; M, \mathcal{D})$$

we have evaluation maps $\text{ev}_0, \text{ev}_1 : E \longrightarrow M$ defined by the expression $\text{ev}_a(e) := \gamma(e)(a)$. We require that ev_0 is constant along the fibres of E .

Definition 4.6.1 *The family γ is **controlling** (in a manner fibered over K) if $\text{ev}_1|_{E_k}$ is a submersion for all $k \in K$.*

Given a section $f : K \rightarrow E$ we can produce a family $\gamma \circ f : K \rightarrow \mathfrak{Maps}([0, 1]; M, \mathcal{D})$. We say that $\gamma \circ f$ is **controllable** and that γ is a **controlling extension**. That is, we are interested in $\gamma \circ f$ and we think of the controlling family γ as a device that allows us to control its endpoints.

4.6.1.2 Controllability

It is immediate from the parametric nature of the implicit function theorem that this implies local controllability:

Lemma 4.6.2 *Given a controlling family γ and a section $f : K \rightarrow E$, there are constants $C, \eta_0 > 0$ such that:*

- *for any $0 < \eta < \eta_0$,*
- *and any smooth choice of endpoint $q_k \in \mathbb{D}_\eta(\gamma \circ f(k)(1))$ (disk of radius η around $\gamma \circ f(k)(1)$),*

there exists a section $g : K \rightarrow E$ such that:

- *g and f are homotopic by a homotopy of C^0 -size at most $C\eta$.*
- *$\gamma \circ g(k)(1) = q_k$.*

Furthermore, if $q_k = \gamma \circ f(k)(1)$, it can be assumed that $g(k) = f(k)$.

The following variation which follows from the inverse function theorem will be useful for us:

Lemma 4.6.3 *Let γ be a controlling family such that its evaluation map $\text{ev}_1|_{E_k}$ is an equi-dimensional embedding for all $k \in K$. Let $q : K \rightarrow M$ with $q \in \text{ev}_1(E_k)$. Then, there is a unique section $f : K \rightarrow E$ such that $\gamma(f(k))(1) = q(k)$.*

4.6.1.3 Existence of controlling families

The following statement follows from standard control theoretical arguments:

Lemma 4.6.4 *A family of regular horizontal curves $\gamma : K \rightarrow \mathfrak{Maps}([0, 1]; M, \mathcal{D})$ admits a controlling family.*

In the non-parametric case, this was proven in [68, Proposition 4 and Corollary 5]. The proof amounts to choosing infinitesimal variations (given by the regularity condition) and integrating these to a controlling family. The parametric case was then proved in [89, Section 8] and requires to patch these variations parametrically in $k \in K$. The idea is that the infinitesimal variations can be “localised” in $[0, 1]$ by an appropriate use of bump functions. This can be exploited to show that variations do not interfere with each other when patching.

4.6.2 Defining controllers

A controller will depend on the following input data: a graphical model (V, \mathcal{D}) , a maximal-displacement $h > 0$, a radius $R > 0$, a size-at-rest S (this will be introduced later, it is just a name for the parameter), and a smoothing-parameter δ . Recall that $\pi : V \rightarrow \mathbb{R}^q$ is the projection to the base in the model and we have a framing $\{X_1, \dots, X_n\}$ of TV such that $\{X_1, \dots, X_q\}$ is a framing of \mathcal{D} obtained as a lift of the coordinate framing of \mathbb{R}^q .

4.6.2.1 Setup

We consider the cube $C = [-R, R]^q \subset \mathbb{R}^q$ and we write $\gamma : [0, 1] \rightarrow C$ for the curve $\gamma(t) = (2Rt - R, 0, \dots, 0)$ parametrising its first coordinate axis. We call it the *axis* of the controller to be built. For each vertical direction $i = q + 1, \dots, n$ in the model, we define a pair of points

$$x_{i,+} = \gamma\left(\frac{i-q}{n}\right), \quad x_{i,-} = \gamma\left(\frac{i-q}{n} = \frac{1}{2n}\right),$$

that are meant to serve as insertion points for tangles. We fix a box $C_{i,\pm}$ centered at $x_{i,\pm}$ and of side $1/4n$. In this manner all the boxes are disjoint.

For each $i = q + 1, \dots, n$, we insert (Proposition 4.5.17) a $(\pm X_i)$ -tangle at $x_{i,\pm}$; we denote it by $\mathcal{T}^{i,\pm}$. The radius of these tangles should be smaller than the side $1/4n$, so that they are contained in the boxes $C_{i,\pm}$. We require that the maximal-displacement of the tangles is $2h$. We write $\mathcal{T}_{d,\theta}^{j,\pm}$ whenever we need to include their estimated-displacement d and birth parameter θ in the discussion. We use δ as their smoothing-parameter.

4.6.2.2 The definition

The estimated-displacement of $\mathcal{T}^{i,+}$ pushes positively along X_i by an amount in the range $(0, 2h]$. Similarly, the estimated-displacement of $\mathcal{T}^{i,-}$ pushes along $-X_i$. We want to combine these two displacements in order to produce motion in the X_i -direction in the range $[-h, h]$. To do so, we recall the size-at-rest S parameter that we fixed earlier, and we consider a bump function $\chi_S : \mathbb{R} \rightarrow [0, 1]$ satisfying

$$\chi_S(-\infty, S] = S, \quad \chi_S|_{[2S, \infty)}(a) = a, \quad \chi'_S|_{(S, 2S)} > 0.$$

Definition 4.6.5 A *controller* is a family of curves

$$\mathcal{C} : [-h, h]^{n-q} \times [0, 1] \longrightarrow \mathfrak{Imm}([0, 1]; C \subset \mathbb{R}^q)$$

parametrised by:

- a **estimated-displacement** $d = (d_{q+1}, \dots, d_n) \in [-h, h]^{n-q}$,
- and a **birth parameter** $\theta \in [0, 1]$.

Each curve $\mathcal{C}(d, \theta)$ is obtained from γ by inserting

$$\mathcal{T}_{\chi_S(d_i), \theta}^{i,+} \quad \text{at} \quad x_{i,+}, \quad \mathcal{T}_{\chi_S(d_i), \theta}^{i,-} \quad \text{at} \quad x_{i,-}.$$

In particular, $\mathcal{C}(d, \theta)$ agrees with γ close to the boundary of C . The controller \mathcal{C} depends on the following parameters:

- the radius $R > 0$ of the box C that contains it,
- the maximal-displacement $h > 0$ bounding the estimated-displacement,
- the size-at-rest S that defines the interpolating function χ_S ,
- the smoothing-parameter $\delta > 0$ of its tangles.

4.6.3 Insertion of controllers

Let (V, \mathcal{D}) be a graphical model and K a compact manifold that serves as parameter space.

Proposition 4.6.6 *Let the following data be given:*

- a family $\gamma : K \rightarrow \mathfrak{Emb}([0, 1]; D_r \subset V, \mathcal{D})$,
- a maximal-displacement $0 < h < 2r$,
- a k -disc $A \subset K$,
- a time $t_0 \in (0, 1)$,
- a constant $\eta > 0$ and a sufficiently small $\tau > 0$,
- a sufficiently small radius $R > 0$, size-at-rest $S > 0$, and smoothing-parameter $\delta > 0$.

Then, assuming that r is sufficiently small, there are:

- a function $\theta : K \rightarrow [0, 1]$ that is identically one on the τ -neighborhood of A $\nu_\tau(A)$ and zero in the complement of $\nu_{2\tau}(A)$,
- a family $\tilde{\gamma} : K \times [-h, h]^{n-q} \times [0, 1] \rightarrow \mathfrak{Emb}([0, 1]; V, \mathcal{D})$,

such that the following statements hold:

- $\tilde{\gamma}(k, d, 0) = \gamma(k)$.
- $\tilde{\gamma}(k, d, s)(t) = \gamma(k)(t)$ if $t \in \mathcal{O}_p(0)$.
- $\pi \circ \tilde{\gamma}(k, d, s)(t) = \pi \circ \gamma(k)(t)$ outside of $\nu_{2\tau}(A) \times \nu_{2\tau}(t_0)$.
- The length of $\tilde{\gamma}(k, d, s)$ in the region $\nu_{2\tau}(t_0) \setminus \nu_\tau(t_0)$ is bounded above by η .
- For all $k \in \nu_\tau(A)$, all $t \in \nu_\tau(t_0)$, and all $s \in [1/2, 1]$, it holds that

$$\pi \circ \tilde{\gamma}(k, d, s)(t) = \mathcal{C}(d, (2s - 1)\theta(k)) \left(\frac{t - t_0 - \tau}{2\tau} \right) + \pi \circ \gamma(k)(t_0).$$

Here the left summand of the right-hand-side is a reparametrised and translated copy of the controller \mathcal{C} with radius R , size-at-rest S , maximal-displacement h , and smoothing-parameter δ .

The family $\tilde{\gamma}$ is said to be obtained from γ by inserting the controller \mathcal{C} along A . The last item asserts that \mathcal{C} has been inserted. The other items provide quantitative control for the insertion.

Proof. We first homotope the family γ in the vicinity of $A \times \{t_0\}$ in order to align its projection with the axis of the controller \mathcal{C} . This follows from an application of Lemma 4.4.10. We denote the resulting homotopy of ε -horizontal curves by $\beta(k, s)|_{s \in [0, 1/2]}$. When we apply Lemma 4.4.10, we use a constant $\tau > 0$ such that: $\beta(k, s)$ agrees with $\gamma(k, s)$ outside of $\nu_{2\tau}(A) \times \nu_{2\tau}(t_0)$. Furthermore, over $\nu_\tau(A) \times \nu_\tau(t_0)$, the projection $\pi \circ \beta(k, 1/2)$ agrees with a translation of the axis of \mathcal{C} . The

desired bound η tells us how small τ must be chosen. In particular, it should be small enough so that $\text{len}(\beta(k, s))$ is bounded above by $\text{len}(\gamma(k)) + \eta/2$.

We then apply Lemma 4.4.3 to β , relative to $t = 0$. This yields a family of horizontal curves $\tilde{\gamma}(k, d, s)|_{s \in [0, 1/2]}$ such that $\pi \circ \tilde{\gamma}(k, d, s) = \pi \circ \beta(k, s)$. Thanks to the bound η and the fact that γ is horizontal, we can invoke Lemma 4.4.4 to assert that $\tilde{\gamma}(k, d, s)$ is indeed defined for all $t \in [0, 1]$. Due to the uniqueness of lifts, $\tilde{\gamma}(k, d, 0) = \gamma(k)$. We denote $\gamma'(k) = \tilde{\gamma}(k, d, 1/2)$.

By construction, over $\nu_\tau(A) \times \nu_\tau(t_0)$, $\pi \circ \gamma'$ agrees with a translation of the axis of \mathcal{C} . Our choice of τ determines the length of the curves $\pi \circ \gamma'$ in the region $\nu_\tau(A) \times \nu_\tau(t_0)$. This is the available length for placing the axis of the controller. As such, it provides an upper bound for our choice of R . We then define

$$\pi \circ \tilde{\gamma}|_{k \in \nu_\tau(A), \quad t \in \nu_\tau(t_0), \quad s \in [1/2, 1]}$$

using the formula appearing in the last item of the statement. For other values of k and t , we set $\pi \circ \tilde{\gamma}(k, d, s) = \pi \circ \gamma'(k)$.

The family $\tilde{\gamma}(k, d, s)|_{s \in [1/2, 1]}$ itself is given from its projection by lifting horizontally. This is done with Lemmas 4.4.3 and 4.4.4. Here is where the smallness of r enters. It must be sufficiently small to control the error in the estimated-displacement of \mathcal{C} (which amounts to the error in the estimated-displacement of its tangles; Lemma 4.5.29). How small r must be depends only on the graphical model (V, \mathcal{D}) . It follows that r being small enough implies that a lift exists for all times and, due to the properties of tangles, it yields embedded curves.

4.6.3.1 Controllability

We now address how the insertion of a controller allows us to control the endpoint of the corresponding horizontal curves.

Lemma 4.6.7 *Consider the setup and conclusions of Proposition 4.6.6. Consider the endpoint map*

$$\mathbf{ep} : K \times [-h, h]^{n-q} \longrightarrow \mathbb{R}^{n-q}$$

defined by $\mathbf{ep}(k, d) := \pi| \circ \tilde{\gamma}(k, d, 1)(1)$. Then the following estimate holds:

$$\mathbf{ep}(k, d) = \gamma(k) + (0, d)(1 + O(r) + O(\delta))$$

for all $k \in A$.

Proof. This is immediate from the analogous statement about tangles, namely Lemma 4.5.29.

As in Lemma 4.3.5 we can be more precise and say that there is a constant C , depending only on the graphical model, such that the error is bounded above by $C \cdot (r + \delta)$. It follows that imposing $\delta, r < 1/C$ implies that $\mathbf{ep}(k, -)$ is an equidimensional embedding whose image contains a ball of radius $h/2$, centered at $\gamma(k)$.

4.6.3.2 Insertion in the ε -horizontal case

For our purposes, we will need the following variation of Proposition 4.6.6:

Lemma 4.6.8 *Let r, h, A, t_0, η, R, S , and δ be as in Proposition 4.6.6. Given $\gamma : K \rightarrow \mathfrak{Emb}^\varepsilon([0, 1]; D_r \subset V, \mathcal{D})$, there is a family $\tilde{\gamma} : K \times [-h, h]^{n-q} \times [0, 1] \rightarrow \mathfrak{Emb}([0, 1]; V)$ such that:*

- $\tilde{\gamma}(k, d, 0) = \gamma(k)$.
- $\tilde{\gamma}(k, d, s) = \gamma(k)$ outside of $\nu_{2\tau}(A) \times \nu_{2\tau}(t_0)$.
- The length of $\tilde{\gamma}(k, d, s)$ in the region $\nu_{2\tau}(t_0) \setminus \nu_{\tau}(t_0)$ is bounded above by η .
- For all $k \in \nu_{\tau}(A)$, all $t \in \nu_{\tau}(t_0)$, and all $s \in [1/2, 1]$, it holds that

$$\pi \circ \tilde{\gamma}(k, d, s)(t) = \mathcal{C}(d, (2s - 1)\theta(k)) \left(\frac{t - t_0 - \tau}{2\tau} \right) + \pi \circ \gamma(k)(t_0).$$

Proof. We use Lemmas 4.4.3 and 4.4.4 to horizontalise γ , yielding some family γ' . We apply to it Proposition 4.6.6, yielding a family $\tilde{\gamma}'$ that is also horizontal. We then adjust its vertical component to yield the claimed $\tilde{\gamma}$. Due to the displacement introduced by the controller, it may be the case that $\tilde{\gamma}$ is not ε -horizontal. However, it is still graphical over \mathbb{R}^q .

4.7 h-Principles for horizontal embeddings

In this section we prove our main Theorem 1.2, the classification of regular horizontal embeddings. This uses all the tools that we have presented in previous sections. The corresponding statement for immersions, Theorem 1.1, will follow from simplified versions of the same arguments. In Subsection 4.7.5 we explain how this is done.

4.7.1 Relative version of Theorem 1.2

We explained in Section 4.2 that our h -principle arguments are relative in nature. This allows us to state an analogue of Theorem 1.2 that deals with embedded regular horizontal paths and is relative in parameter and domain.

Proposition 4.7.1 *Let K be a compact manifold. Let $I = [0, 1]$. Let (M, \mathcal{D}) be a manifold of dimension $\dim(M) > 3$, endowed with a bracket-generating distribution. Suppose we are given a map $\gamma : K \rightarrow \mathfrak{Emb}^{\varepsilon}(I; M, \mathcal{D})$ satisfying:*

- $\gamma(k) \in \mathfrak{Emb}^r(I; M, \mathcal{D})$ for $k \in \mathcal{Op}(\partial K)$.
- $\gamma(k)(t)$ is horizontal if $t \in \mathcal{Op}(\partial I)$.

Then, there exists a homotopy $\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}^{\varepsilon}(I; M, \mathcal{D})$ satisfying:

- $\tilde{\gamma}(k, 0) = \gamma(k)$.
- $\tilde{\gamma}(k, 1)$ takes values in $\mathfrak{Emb}^r(I; M, \mathcal{D})$.
- The homotopy is relative to $k \in \mathcal{Op}(\partial K)$. Also, $\tilde{\gamma}(k, s)(t) = \gamma(k)(t)$ for all k, s and for all $t \in \mathcal{Op}(\partial I)$.
- $\tilde{\gamma}(k, s)$ is C^0 -close to $\gamma(k)$ for all $s \in [0, 1]$.

Do note that the analogous statement where we consider formal horizontal embeddings instead of ε -horizontal ones follows from this one, thanks to the results in Subsection 4.2.2.

Now, as usual, the absolute statement follows from the relative one:

Proof (of Theorem 1.2 from Proposition 4.7.1). The statement follows, according to Subsection 4.2.2, from the vanishing of the relative homotopy groups of the pair

$$(\mathfrak{Emb}^\varepsilon(M, \mathcal{D}), \mathfrak{Emb}^r(M, \mathcal{D})).$$

Consider a family $\gamma : \mathbb{D}^a \rightarrow \mathfrak{Emb}^\varepsilon(M, \mathcal{D})$ that takes values in $\mathfrak{Emb}^r(M, \mathcal{D})$ along \mathbb{S}^{a-1} . This represents a class in the a th relative homotopy group. We must deform this family to lie entirely in $\mathfrak{Emb}^r(M, \mathcal{D})$.

We may assume, by suitable reparametrisation in the parameter, that in a collar of \mathbb{S}^{a-1} , the family γ is radially constant. This provides for us an open along the boundary of \mathbb{D}^a in which all curves are regular horizontal. We then consider the product space $\mathbb{D}^a \times \mathbb{S}^1$ (this \mathbb{S}^1 factor is the domain of each loop) and we make the curves γ horizontal in a neighbourhood of the slice $\mathbb{D}^a \times \{1\}$, using Lemma 4.4.8. This is a C^0 -small process.

Regarding $\mathbb{D}^a \times \mathbb{S}^1$ as $\mathbb{D}^a \times I / \sim$ yields a family of ε -horizontal paths, horizontal at the endpoints, with all curves regular close to $\partial\mathbb{D}^a$. We apply Proposition 4.7.1 to it, relatively to $\mathcal{O}p(\mathbb{S}^{a-1})$ in the parameter, and relatively to $\mathcal{O}p(\partial I)$ in the domain. The claim follows.

The remainder of this section is mostly dedicated to the proof of Proposition 4.7.1.




Figure 4.20: Choosing the time slice $\mathbb{D}^a \times \{1\}$ identifies the product $\mathbb{D}^a \times \mathbb{S}^1$ with $\mathbb{D}^a \times [0, 1]$, allowing us to apply Proposition 4.7.1.

4.7.2 Setup for the proof of Proposition 4.7.1

Given the family of curves γ , we will produce a series of homotopies in order to obtain $\tilde{\gamma}(-, 1)$. These elementary homotopies are relative to the boundary and through ε -horizontal curves. Their concatenation will be the homotopy $\tilde{\gamma}$.

We denote $\dim(M) = n$, $\text{rank}(\mathcal{D}) = q$, and $\dim(K) = k$. We write π_K and π_I for the projections of $K \times I$ to its factors. We fix a metric on K , endow I with the euclidean metric, and endow the product $K \times I$ with the product metric. We will write $d(-, -)$ to denote distance between subsets; the metric used should be clear from context.

4.7.2.1 A fine cover of $K \times I$

Our first goal is to subdivide $K \times I$ in a manner that is nicely adapted to (M, \mathcal{D}) and γ . The aim with this is to reduce our subsequent arguments to constructions happening in very small balls in

which errors are controlled. No homotopy of γ is produced in this subsection, we are just doing some preliminary work.

Introduce a size parameter $N \in \mathbb{N}$, to be fixed as we go along in the proof. We divide I into intervals I_j of length $1/N$. Furthermore, we fix a finite cover of K by charts parametrised by the unit cube. Such cubes can themselves be divided into cubes of side $3/N$, spaced along the coordinate axes as $1/N$. We write

$$\{\phi_i : [0, 3/N]^k \longrightarrow U_i \subset K\}$$

for the resulting collection of cubical charts.

4.7.2.2 Bounding the number of intersections between charts

By construction, there is a constant C_1 such that any intersection $U_{i_1} \cap \dots \cap U_{i_{C_1}}$, involving distinct charts, is empty. This follows from the properties of cubical subdivision; a detailed argument can be found in Section 5.4.3 of this Thesis.

4.7.2.3 Covering the image by graphical models

Given the origin $k \in K$ of the chart U_i and the initial time $t_j := j/N \in I_j$, we fix an adapted chart $(V_{i,j}, \Psi_{i,j})$ centered at $\gamma(k)(t_j)$. These adapted charts are given by Lemma 4.3.5, meaning that they all have the same radius $r_0 > 0$ and the difference between their framing and the coordinate axes is controlled by some constant $C_2 > 0$.

4.7.2.4 Discussion about parameters

The constants C_1 , C_2 , and r_0 are given to us and depend on (M, \mathcal{D}) and the family γ . For convenience, we introduce new parameters $0 < r_1 < r_0$ and $0 < l$, to be fixed later in the proof. We impose $1/N \ll r_1, l$ in order to ensure that:

$$\gamma(k)(t) \in \Psi_{i,j}(\mathbb{D}_{r_1}) \text{ for all } (k, t) \in U_i \times I_j \quad (4.2)$$

$$\Psi_{i,j}^{-1} \circ \gamma(k)|_{I_j} \text{ has euclidean length bounded above by } l \quad (4.3)$$

I.e. the curves in each cube $U_i \times I_j$ are very short and are located very close to the origin of the corresponding graphical model.

4.7.3 Introducing controllers

Our next goal is to add controllers to γ over each cube $U_i \times I_j$. We continue using the notation introduced in the previous subsection. We write $\nu_r(A)$ for the r -neighbourhood of a subset A .

4.7.3.1 Boundary neighbourhoods

We fix a small enough constant $\tau > 0$ so that all curves $\gamma(k)$ with $k \in \nu_\tau(\partial K)$ are horizontal and regular. We then consider a constant $\tau/2 < \tau' < \tau$ and a subset \mathcal{U} of the cover $\{U_j \times I_i\}$ so that:

- the elements of \mathcal{U} , together with $\nu_{\tau'}(\partial(K \times I))$, cover $K \times I$,
- all elements in \mathcal{U} are disjoint from $\nu_{\tau/2}(\partial(K \times I))$.




Figure 4.21: The cover \mathcal{U} , together with the boundary neighbourhoods of radii τ and $\tau/2$. A controller is shown in one of the $U_i \times I_j$.

4.7.3.2 Introducing controllers

Given $U_i \times I_j \in \mathcal{U}$, we choose a time $t_{i,j}$ in the interior of I_j . We require that these are all distinct. We then introduce a controller $\mathcal{C}_{i,j}$ (Lemma 4.6.8) along $U_i \times \{t_{i,j}\}$. We write:

- S for the size-at-rest of all the controllers.
- $\eta > 0$ for the size of the neighbourhood of $U_i \times \{t_{i,j}\}$ in which the controllers are contained. If η is sufficiently small, the controllers do not interact with one another.

Later on in the proof we will use the estimated-displacement of the $\mathcal{C}_{i,j}$, one pair (i, j) at a time. For now we write γ' for the K -family of ε -horizontal curves that has the estimated-displacement of each controller at 0. Do note that S must be small enough to guarantee ε -horizontality. Furthermore, γ' is homotopic to γ through ε -horizontal curves, relative to the complement of all $\nu_\tau(U_i \times \{t_{i,j}\})$.

There is now a subtlety that we need to take care of: some of the $\mathcal{C}_{i,j}$ enter the region $\nu_\tau(\partial(K \times I))$, destroying the horizontality condition there. This cannot be addressed using the controllers themselves, since the collection $\{U_i\}$ does not cover $\nu_\tau(\partial K)$ completely. We address it instead using regularity. Note that this is not a technical point: due to the phenomenon of rigidity (Subsection 4.1.1), the usage of regularity at this stage cannot be avoided.

4.7.3.3 Horizontalisation

The family γ' is horizontal over $\mathcal{O}p(\partial K)$, but not necessarily over the whole band $\nu_\tau(\partial K)$, due to our insertion of controllers. To address this, we reintroduce horizontality, at the cost of losing control of the endpoints. Namely, given $\mu > 0$, any sufficiently small size-at-rest S will guarantee that there is a family of horizontal curves $(\alpha(k))_{k \in \nu_\tau(\partial K)}$ that satisfies:

- i. $|\alpha - \gamma|_{\nu_\tau(\partial K)}|_{C^0} < \mu$.
- i'. $\text{len}(\alpha(k)) < \text{len}(\gamma(k)) + \mu$.
- ii. α is homotopic to $\gamma|_{\nu_\tau(\partial K)}$ through horizontal curves.
- iii. This homotopy is relative in the parameter to $\nu_{\tau/2}(\partial K)$.
- iii'. The homotopy is relative to $\{t = 0\}$ in the domain.

The family α is constructed inductively, one chart $U_i \times I_j$ at a time, increasingly in j , and arbitrarily in i . The inductive step consists of using each adapted chart $(V_{i,j}, \Psi_{i,j})$ to see $\gamma'|_{U_i \times I_j}$ as a family of curves in the graphical model $V_{i,j}$. We can then apply the lifting Lemma 4.4.3 to the projection $\pi \circ \gamma'|_{U_i \times I_j}$. The lifting process can be completed over $[0, 1]$ because a small S means that $\pi \circ \gamma'|_{U_i \times I_j}$ is close to $\pi \circ \gamma|_{U_i \times I_j}$. This also justifies Conditions (i) and (i'). The families are homotopic to one another via the birth of the controller, proving Condition (ii). Lastly, projecting and lifting leaves horizontal curves invariant, proving Conditions (iii) and (iii').

We now choose μ small enough so that the bounds provided by Conditions (i) and (i') allows us to invoke the interpolation Lemma 4.4.5. This allows us to interpolate through ε -horizontal curves between α and γ' in the region $\{\tau' < d(k, \partial K) < \tau\}$. The resulting family of curves α' :

- agrees with γ' in the complement of $\nu_\tau(\partial K)$,
- agrees with γ in $\nu_{\tau/2}(\partial K)$,
- is horizontal in $\nu_{\tau'}(\partial K)$,
- contains controllers along $U_i \times \{t_{i,j}\}$.

The issue is that $\alpha(k)(1)$ may be different from $\gamma'(k)(1) = \gamma(k)(1)$ in the region $\{\tau/2 < d(k, \partial K) < \tau'\}$. This is a feature of the lifting process. Nonetheless, according to Lemma 4.4.3, the endpoints are μ -close.

4.7.3.4 Controllability

By hypothesis, the curves $\gamma|_{\nu_\tau(\partial K)}$ are horizontal and regular. It follows that $\gamma|_{\nu_\tau(\partial K)}$ is a controllable family³, according to Lemma 4.6.4. We deduce that there are constants $c, \delta > 0$ such that any δ -displacement of their endpoints can be followed by a homotopy of the curves themselves, through horizontal curves, that is $c\delta$ -small.

We claim that the family $\alpha'|_{\nu_{\tau'}(\partial K)}$ is also controllable, with constants $2c$ and $\delta/2$, as long as S is sufficiently small. Indeed, consider the variations F of γ that yield controllability. Then, the homotopy lifting property, applied to F and the homotopy of horizontal curves connecting $\gamma|_{\nu_{\tau'}(\partial K)}$ with $\alpha'|_{\nu_{\tau'}(\partial K)}$, yields corresponding variations for small values of the homotopy parameter. They exist for the whole homotopy if μ is assumed to be sufficiently small.

³ Do note that $\nu_\tau(\partial K)$ is not compact but, since τ is arbitrary, we can take a slightly smaller compact neighbourhood of ∂K and carry the argument there.

Then, assuming that S is sufficiently small, we have that $\mu < \delta/2$ and we can use the controllability of $\alpha'|_{\nu_{\tau'}(\partial K)}$ to yield a family $\gamma'' : K \rightarrow \mathfrak{Emb}^\varepsilon(I; M, \mathcal{D})$ that:

- agrees with γ in $\mathcal{O}p(\partial(K \times I))$.
- is horizontal and regular if $k \in \nu_{\tau'}(\partial K)$.
- has a family of controllers, still denoted by $\{\mathcal{C}_{i,j}\}$, along $U_i \times \{t_{i,j}\}$.

The situation is depicted in Figure 4.22.




Figure 4.22: For each element in the cover \mathcal{U} , a controller has been introduced. These are shown as green thin rectangles. Close to the boundary, some controllers (in red) enter its τ -neighbourhood. Horizontality is reestablished using the variations given by local controllability.

4.7.3.5 Discussion about parameters

In this subsection we added controllers $\{\mathcal{C}_{i,j}\}$ to the family γ . Each curve $\gamma(k)$ crosses at most $N \cdot C_1$ controllers; here C_1 is the upper bound for the intersections between elements in $\{U_i\}$. Inserting the controllers produces a deformation α of $\gamma|_{\nu_{\tau'}(\partial K)}$ through horizontal curves. This deformation displaces the endpoints an amount μ , which we can estimate. For each controller inserted, the endpoints move a magnitude $O(S)$, the size-at-rest. This implies that μ is bounded above by $O(S) \cdot N \cdot C_1$.

Furthermore, in Subsection 4.7.2 we showed that γ satisfies the size estimates given in Equations 4.2 and 4.3, involving r_1 and l . We want α and thus γ' to satisfy these as well. To this end, the C^0 -distance between γ , α , and γ' must be much smaller than r_1 and l .

These considerations force the choice $S \ll 1/N \ll l, r_1$.

4.7.4 Using the controllers

In this subsection we complete the proof of Proposition 4.7.1. The idea is to use projection and lifting to replace γ'' by a family of horizontal curves β , whose endpoints are incorrect. The endpoints will then be adjusted thanks to the presence of controllers.

4.7.4.1 Horizontalisation

Much like in the proof of Theorem 1.2, we first apply Lemma 4.4.8 to γ'' at each time $t_j = j/N \in I$. This can be done in a C^0 -small way, through ε -horizontal curves, by making the newly created horizontal region sufficiently small. This is relative to $\nu_{\tau'}(\partial K)$ in the parameter. The resulting

family is denoted by γ_0 . The proof now focuses on a concrete interval I_j ; the argument is identical for all of them.

4.7.4.2 Horizontalisation again

We have a family γ_0 with values in $\mathfrak{Emb}^\varepsilon(I_j; M, \mathcal{D})$ that along the boundary of $K \times I_j$ is horizontal. Suppose $1/N$ is sufficiently small. We can argue as in Subsection 4.7.3.3 to construct a K -family β with values in $\mathfrak{Emb}(I_j; M, \mathcal{D})$ such that:

- $\beta(k) = \gamma_0(k)$ if $d(k, \partial K) \leq \tau'$.
- $\beta(k)(t) = \gamma_0(k)(t)$ if $t \in \mathcal{O}p(\{t_j\})$.
- Write π for the projection to the base given by each graphical model $V_{i,j}$. Then, the C^∞ -closeness of $\pi \circ \beta$ and $\pi \circ \gamma_0$ is controlled by S .
- In particular, the two families are homotopic through ε -horizontal curves.

In particular, the controllers of the family γ_0 define controllers for β . We still denote them by $\{\mathcal{C}_{i,j}\}_i$. The goal now is to use these to produce a homotopy of horizontal curves between β and the claimed $\tilde{\gamma}|_{K \times \{1\}}$.

4.7.4.3 Adjusting the endpoint over one chart

The difference $e = |\gamma_0(k)(t_{j+1}) - \beta(k)(t_{j+1})|$ is certainly bounded above by r_1 . However, as explained in Subsection 4.3.2.1, it can also be bounded above by $C_2 \cdot r_1 \cdot l$. It follows that we should impose $l \ll 1/(C_2 \cdot r_1)$ to make e much smaller than r_1 . By making this choice, Equations 4.2 and 4.3 hold for β .

We now perform induction on i to correct this difference. We start with the base case $i = 1$, so we work over the chart U_1 . Adjusting the estimated-displacement of the controller $\mathcal{C}_{1,j}$ yields a homotopy of horizontal curves

$$\tilde{\beta} : U_1 \times A \subset \mathbb{R}^{n-q} \longrightarrow \mathfrak{Emb}(I_j; M, \mathcal{D})$$

with $\tilde{\beta}(k, 0) = \beta(k)$. The variable $a \in A$ measures the endpoint displacement introduced by the controller vertically. We package this as an endpoint map

$$\tilde{\beta}(-, -)(t_{j+1}) : U_1 \times A \longrightarrow \mathbb{R}^{n-q}$$

which satisfies the following error estimate:

$$\tilde{\beta}(k, a)(t_{j+1}) = \beta(k)(t_{j+1}) + a \cdot (1 + C_3 \cdot (r_1 + |a| + \delta)).$$

According to Lemma 4.3.5, the constant C_3 appearing as the coefficient of the error term $(r_1 + |a|)$ is independent of the adapted chart and thus independent of the controller. This forces us to choose $\delta, r_1 \ll 1/C_3$. The quantity $|a|$ will be of magnitude e and thus smaller than r_1 .

This choice tells us that $\tilde{\beta}(k, -)(t_{j+1})|_{\mathbb{D}_{2r_1}}$ is an embedding whose image contains \mathbb{D}_{r_1} , for all $k \in V_1$. In particular, it contains the desired endpoint $\gamma_0(k)(t_{j+1})$. The inverse function theorem (Lemma 4.6.3) defines for us a unique function $a : V_1 \rightarrow \mathbb{D}_{2r_1}$ so that $\tilde{\beta}(k, a(k))(t_{j+1}) = \gamma_0(k)(t_{j+1})$.

We now cut-off the function a , in order to make the construction relative to the boundary of $U_1 \times I_j$. Fix a constant $\rho > 0$ and write $W_i \subset U_i$ for a domain covering U_i up to a ρ -neighbourhood

of its boundary. We require that the family $\{W_i\}$ is an open cover of $K \setminus \nu_{\tau'}(\partial K)$. This imposes $\rho \ll 1/N, \tau$. This allows us to introduce a cut-off function $\chi_1 : K \rightarrow [0, 1]$ that is one in W_1 and zero along ∂U_1 . We set $\beta_1(k) := \tilde{\beta}(k, \chi(k).a(k))$. This is a family of horizontal curves such that:

- $\beta_1(k)$ is horizontal.
- β_1 is homotopic to β as maps into $\mathfrak{Emb}(I_j; M, \mathcal{D})$.
- The base projections of $\beta_1(k)$ and $\gamma_0(k)$ agree for all $k \in \mathcal{Op}(\partial U_1)$.
- $\beta_1(k)$ and $\gamma_0(k)$ agree over $t \in \mathcal{Op}(\partial I_j)$, for every $k \in W_1$.

The third property allows us to homotope β_1 to a family $\gamma_1 : U_1 \rightarrow \mathfrak{Emb}^\varepsilon(I_j; M, \mathcal{D})$ that is horizontal over W_1 and agrees with γ_0 in $\mathcal{Op}(\partial U_1)$. We can then use γ_0 to extend γ_1 to a family $K \rightarrow \mathfrak{Emb}^\varepsilon(I_j; M, \mathcal{D})$. The two are homotopic, relative to endpoints and to the complement of U_1 , thanks to the second property and the fact that β was the horizontal lift of γ_0 .

4.7.4.4 The inductive argument

The i_0 -th inductive step follows the exact same argument. It produces a family $\gamma_{i_0} : K \rightarrow \mathfrak{Emb}^\varepsilon(I_j; M, \mathcal{D})$ that is horizontal over the union $\cup_{i \leq i_0} W_i$. The observation to be made is the following. Suppose U_{i_0} intersects non-trivially some previous U_i . Then, in the overlap $U_{i_0} \cap W_i$, we have that the family γ_{i_0-1} is already horizontal over W_i . It follows that the associated horizontal family β agrees with γ_{i_0-1} over W_i , due to the uniqueness of horizontal lifts. In particular, when we use the controller $\mathcal{C}_{i_0,j}$ to produce a fully controllable family over W_{i_0} , we see that no adjustments must be made over W_i , since the endpoint is already correct there. This is immediate from the uniqueness provided by the inverse function theorem.

Since the W_i cover $K \setminus \nu_\tau(\partial K)$, the inductive argument produces the required homotopy $\tilde{\gamma}$. The proof of Proposition 4.7.1 is complete.

4.7.4.5 Final discussion about constants

Two quantities are left to be controlled which we will do in this section. The first is the error accumulated by the controllers. This was proportional (Lemma 4.6.7) to the radius of the graphical models (and could therefore be controlled by r_1 and by the smoothing parameter δ). The other quantity is the C^0 -distance e between β and γ_{i-1} (particularly at their endpoints). This was controlled by setting $l \ll r_1$.

We note that the embedding condition enters the discussion only in the choice of r_1 . Namely: once the error of the controllers has been bounded, it follows that the horizontal curves produced by the controller are indeed embedded. See Lemma 4.6.7.

The summary is that we require the chain of inequalities $\delta \ll S \ll 1/N \ll l \ll r_1$. \square

Remark 4.7.2 *The reader experienced in h-principles may wonder why we do not use a triangulation of $K \times I$, in general position with respect to π_K , to argue. Indeed, this would have the added advantage of localising our arguments to balls that do not interact with one another. This was not the case in the proof we presented.*

The issue with the triangulation approach is that we would have to make a first homotopy that makes our curves horizontal along the codimension-1 skeleton. This is certainly possible, but we




Figure 4.23: The inductive process. In the i th step of the induction we introduce horizontality over the region W_i , in dark blue. Appropriate cut-off functions have been introduced in $U_i \setminus W_i$ (light red) to make this homotopy relative to the boundary of the model. At a later stage, we consider some U_{i_0} overlapping with W_i . In the overlap, the step i_0 homotopy is constant, thanks to horizontality.

have no guarantee that the produced curves are themselves regular. This is absolutely necessary, since we need to be able to introduce controllers at the bottom of each top-cell.

In fact, this can be made to work. The local integrability of micro-regular curves (i.e. curves that are in particular regular over any interval) was proven in [89, 16], which would allow us to produce regular curves along the skeleton. However, it seemed preferable to us to keep the proof self-contained and not invoke additional results. \square

4.7.5 Other h -principles for horizontal curves

We now discuss how the proof of Theorem 1.2 adapts to prove Theorems 1.19 and 1.3.

Proof (of Theorem 1.19). The absolute statement can be reduced to proving the relative h -principle over the interval (i.e. the analogue of Proposition 4.7.1). The proof of the relative statement is identical to the one we presented for embeddings. The reader can check that the proof goes through line by line. Instead, it is more interesting to point out how the proof simplifies for immersions.

First note that the construction of tangles (Section 4.5) and controllers (Section 4.6) is less involved if we do not need to take care of self-intersections. In particular, we do not need to develop all the explicit models shown in Figures 4.14 and 4.15. Similarly, in Subsection 4.7.4.5, we do not need to control errors in order to ensure the controller produces embedded curves.

This crucial difference explains why the statement for immersions goes through in $\dim(M) = 3$. Achieving embeddedness parametrically was not possible in dimension 3, but one can certainly produce tangles and controllers that are immersed.

Similarly:

Proof (of Theorem 1.3). The statement reduces once again to the h -principle for horizontal paths, relative both in parameter and domain. We now indicate the differences with respect to the proof of Proposition 4.7.1.

First: since there is no first order formal data, we do not need to set up a convex integration argument to achieve ε -horizontality. Second: we use the covering arguments as they appear in

Subsection 4.7.2 but we run into issues in Subsection 4.7.3, when we try to introduce controllers. In Proposition 4.6.6 we explained how to introduce them when our curves are embedded/immersed, which may not be the case here. To address this, we use the “stopping trick”; see Figure 4.24. We explain it next.

Given a smooth horizontal arc $\gamma : [a, b] \rightarrow (M, \mathcal{D})$, we can precompose it with a non-decreasing map $\phi : [a, b] \rightarrow [a, b]$ that is the identity at the endpoints and is constant in $\mathcal{O}p(\{(a+b)/2\})$. Since ϕ is homotopic to the identity rel boundary, it defines a homotopy between γ and a horizontal curve $\gamma \circ \psi$ whose parametrisation is stationary at the middle point. The curve $\gamma \circ \psi$ is thus regular. Even more: suppose $\nu : [0, 1] \rightarrow (M, \mathcal{D})$ is some other horizontal curve with $\nu(0) = \gamma \circ \psi((a+b)/2)$. Then $\gamma \circ \psi$ is homotopic to a smooth horizontal curve γ' whose image is the concatenation

$$(\gamma \circ \psi|_{[(a+b)/2, b]}) \bullet \nu \bullet (\gamma \circ \psi|_{[a, (a+b)/2]}).$$

Since ν is arbitrary, it may be chosen to be a curve contained in a graphical model around $\gamma \circ \psi((a+b)/2)$ and projecting to the base as whatever direction we require.

The conclusion is that any family $\gamma : K \rightarrow \mathcal{L}(M)$, horizontal at the boundary, can also be assumed to be regular along ∂K , up to a homotopy through horizontal curves. Furthermore, it can be assumed to be immersed whenever controllers need to be introduced. This concludes the proof.




Figure 4.24: The stopping trick. The space of horizontal loops is extremely flexible. Any family can be homotoped to a regular family by introducing stationary points in the parametrisation. This allows us to introduce controllers.

4.8 h-Principle for transverse embeddings

In this section we prove Theorem 1.5, the classification of transverse embeddings.

4.8.1 The relative *h*-principle

The *h*-principle for transverse paths, relative in parameter and domain, reads:

Proposition 4.8.1 *Let K be a compact manifold. Let $I = [0, 1]$. Let (M, \mathcal{D}) be a manifold of dimension $\dim(M) > 3$, endowed with a bracket-generating distribution of corank 1. Suppose that we are given a map $\gamma : K \rightarrow \mathfrak{Emb}_{\mathcal{T}}^f(I; M, \mathcal{D})$ satisfying:*

- $\gamma(k) \in \mathfrak{Emb}_{\mathcal{T}}(I; M, \mathcal{D})$ for $k \in \mathcal{O}p(\partial K)$.
- $\gamma(k)(t)$ is positively transverse if $t \in \mathcal{O}p(\partial I)$.

Then, there exists a homotopy $\tilde{\gamma} : K \times [0, 1] \rightarrow \mathfrak{Emb}_{\mathcal{T}}^f(I; M, \mathcal{D})$ satisfying:

- $\tilde{\gamma}(k, 0) = \gamma(k)$.

- $\tilde{\gamma}(k, 1)$ takes values in $\mathfrak{Emb}_{\mathcal{T}}(I; M, \mathcal{D})$.
- this homotopy is relative to $k \in \mathcal{Op}(\partial K)$ and to $t \in \mathcal{Op}(\partial I)$.
- $\tilde{\gamma}(k, s)$ is C^0 -close to $\gamma(k)$ for all $s \in [0, 1]$.

Observe that we do not assume that \mathcal{D} is cooriented. We will be able to pass to the cooriented case (and thus use ε -transverse embeddings) during the proof.

Proof (Proof of Theorem 1.5 from Proposition 4.8.1). We must prove the vanishing of the relative homotopy groups of the pair

$$(\mathfrak{Emb}_{\mathcal{T}}^f(M, \mathcal{D}), \mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})).$$

Given a family γ representing a class in the a -th relative homotopy group, we must deform it to lie entirely in $\mathfrak{Emb}_{\mathcal{T}}(M, \mathcal{D})$. Close to $\partial \mathbb{D}^a$ the family is transverse because transversality is an open condition. We then fix a slice $\mathbb{D}^a \times \{1\} \subset \mathbb{D}^a \times \mathbb{S}^1$ and apply the transversalisation Lemma 4.4.13 to γ there. This reduces the argument to a family of paths and thus to Proposition 4.8.1.

We henceforth focus on the proof of Proposition 4.8.1

4.8.2 The first triangulation step

We are given a family of formally transverse curves γ . Our first goal is to pass to the ε -transverse setting. To do so, we will produce a very fine subdivision of $K \times I$ so that we can work over balls in (M, \mathcal{D}) . We can then use the local coorientability of \mathcal{D} to introduce ε -transversality.

4.8.2.1 Triangulating

We subdivide $K \times I$ using a triangulation \mathcal{T} . This should be compared to the proof of Proposition 4.7.1, which used a different scheme to localise the arguments to little balls. The reason was explained in Remark 4.7.2: Triangulating $K \times I$ would have led to issues due to the phenomenon of rigidity for horizontal curves. However, there is no rigidity for transverse curves because they are defined by an open condition.

In order to produce a triangulation, we proceed as follows. We pick a small constant $\tau > 0$ so that γ is transverse over $\nu_{\tau}(\partial(K \times I))$. We choose a closed domain $A \subset K \times I$ such that $\{A, \nu_{\tau}(\partial(K \times I))\}$ covers $K \times I$. We ask that ∂A is smooth. We then apply the jiggling Corollary 4.4.2 to $(A, \text{Ker}(d\pi_K))$. We do not need to introduce further subdivisions, we simply choose \mathcal{T} fine enough so that each top-cell $\Delta \in \mathcal{T}$ is mapped by γ to an adapted chart of (M, \mathcal{D}) .

According to Lemma 4.4.1, the following holds: a simplex Δ is either transverse to the vertical or is contained in ∂A . In the latter case, γ is already transverse in $\mathcal{Op}(\Delta)$. It follows that we can apply Lemma 4.4.14 to γ , along the codimension-1 skeleton, in a manner relative to ∂A . This yields a homotopy of formal transverse embeddings, relative to ∂A , between γ and a family γ_1 . The family γ_1 is transverse on a neighbourhood of the codimension-1 skeleton.

4.8.2.2 ε -transversality

Let $\Delta \in \mathcal{T}$ be a top-cell. Since γ_1 maps it to an adapted chart (V, Φ) , we have that $\gamma_1|_\Delta$ takes values in a manifold endowed with a coorientable distribution. Furthermore, since $\gamma_1|_\Delta$ is formally transverse, it defines a preferred coorientation for $\mathcal{D}|_{\Phi(V)}$.

According to Lemma 4.4.1, the top-cells of \mathcal{T} are homotopic to flowboxes, meaning that there is a fibre-preserving embedding

$$\Psi : \mathbb{D}^k \times [0, 1] \longrightarrow \Delta \subset A \subset K \times I$$

whose image covers most of Δ . In particular, the boundary of this embedding may be assumed to be contained in the region where γ is transverse.

The conclusion is that $\Phi^{-1} \circ \gamma_1 \circ \Psi$ is a family of formally transverse curves that:

- takes values in a cooriented graphical model (V, \mathcal{D}_V) ,
- is transverse over $\mathcal{O}p(\partial(\mathbb{D}^k \times [0, 1]))$.

This implies that we can apply the h -principle for ε -transverse curves (Subsection 4.2.4) to $\Phi^{-1} \circ \gamma_1 \circ \Psi$, yielding a family $\gamma_2 : \mathbb{D}^k \longrightarrow \mathbf{Emb}_{\mathcal{T}^\varepsilon}([0, 1]; V, \mathcal{D}_V)$ that is (positively) transverse over $\mathcal{O}p(\partial(\mathbb{D}^k \times [0, 1]))$.

We have one such family per top-cell Δ . Furthermore, their domains are disjoint. This implies that it is sufficient for us to work with each γ_2 individually, homotoping them through ε -transverse curves (and relative to $\partial(\mathbb{D}^k \times [0, 1])$) to a family of transverse curves.

4.8.3 The second triangulation step

Let us explain how the remainder of the proof goes, morally. Our goal is to apply the case of horizontal embeddings. The reasoning is that, whenever the curves $\gamma_2(k)$ are graphical over \mathcal{D}_V , we can flatten them to make them almost horizontal, which will then allow us to manipulate them through the use of controllers. In order to project towards \mathcal{D}_V , we use the projection $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^q$ to the base of the graphical model.

There are two issues with this idea. The first is that the family γ_2 may have uncontrolled length, since it was produced by the h -principle for ε -transverse curves. To address this, we will triangulate again to pass to small balls. The other issue will be explained afterwards.

4.8.3.1 Triangulating again.

We proceed as above, applying the jiggling Lemma 4.4.1 to γ_2 and $(\mathbb{D}^k \times [0, 1], \text{Ker}(d\pi_{[0, 1]}))$. We obtain a sequence of triangulations \mathcal{T}_b inducing a triangulation of the boundary (do note that it has corners, but this is not an issue). We fix b as in the horizontal case (Subsection 4.7.2): each simplex Δ should have diameter bounded above by $1/N \ll l, r_1$. In this manner, $\gamma_2(\Delta)$ is contained in an adapted chart (U, Ψ) of radius r_1 and each curve in $\gamma_2|_\Delta$ has length at most l from the perspective of the corresponding graphical model U . We still refer to these conditions as Equations 4.2 and 4.3.

We now apply Lemma 4.4.13 to γ_2 , along the codimension-1 skeleton of \mathcal{T}_b . This shows that γ_2 is homotopic, relative to the boundary, to a family γ_3 that is transverse along the codimension-1 skeleton. This is a C^0 -small process that does not increase the length of the curves much. It follows that the conditions given by Equations 4.2 and 4.3 apply to γ_3 as well.

As we did earlier, we can embed a copy of $\mathbb{D}^k \times [0, 1]$ into each top-simplex, in a fibered manner, in such a way that its boundary lies in the region where γ_3 is transverse. These domains do not interact with one another, so we argue on each of them separately.

4.8.3.2 Triangulating one more time

We now encounter the second issue. We cannot flatten γ_3 to make it almost horizontal along the locus

$$\Sigma := \{(k, t) \in D^k \times [0, 1] \mid d\pi(\gamma_3(k)') = 0\}.$$

That is, the locus where the velocity vector becomes vertical. Nonetheless, using Thom transversality [39, p. 17], we can assume that Σ is a closed submanifold of $\mathbb{D}^k \times [0, 1]$.




Figure 4.25: Schematic representation of $\mathbb{D}^k \times [0, 1]$. The blue band represents a neighbourhood of the boundary. The thin green curve is the vertical locus Σ . Its neighbourhood W is shown in red. The region B is a slight thickening of the complement.

The family γ_3 is vertical along Σ and ε -transverse in general. It follows that γ_3 is positively transverse on a neighbourhood $W \supset \Sigma$. We can therefore find a closed subdomain $B \subset \mathbb{D}^k \times [0, 1]$, disjoint from W , whose smooth boundary lies in the region where γ_3 is transverse. See Figure 4.25.

We can then proceed as above, applying Lemma 4.4.1 to γ_3 and $(B, \text{Ker}(d\pi_{[0,1]}))$. We do not need the resulting triangulation \mathcal{T}' to be thin, since we already achieved quantitative control in the previous subdivision. Applying Lemma 4.4.13 shows that γ_3 is homotopic, relative to ∂B , to a family γ_4 that is transverse along the codimension-1 skeleton. This is a C^0 -small process.

We henceforth argue on each top simplex separately, relative to the boundary. The punchline is that we have a family

$$\gamma_4 : \mathbb{D}^k \longrightarrow \mathfrak{Emb}_{\mathcal{T}^\varepsilon}([0, 1]; \mathcal{D}_U),$$

that is transverse along the boundary. Here (U, \mathcal{D}_U) is the graphical model of radius r_1 that we fixed earlier (and that we used to discuss verticality). Each curve of γ_4 has length bounded above by l . Since we avoided Σ , we can assume that each curve $\gamma_4(k)$ is graphical over \mathcal{D}_U .

4.8.4 End of the argument

Choose $\tau > 0$ small enough so that γ_4 is transverse in a τ -neighbourhood of the boundary of $\mathbb{D}^k \times [0, 1]$. We can apply Corollary 4.4.9 to γ_4 along the slice $D = \mathbb{D}_{1-\tau/2}^k \times \{\tau/2\}$, in order to yield a family that is almost transverse in $\mathbb{D}^k \times [0, \tau]$ and horizontal in $\mathcal{O}p(D)$. This allows us to introduce a controller \mathcal{C} along D ; see Lemma 4.6.8. The resulting family is called γ_5 . It consists of embedded curves as long as r_1 was sufficiently small (Lemma 4.6.7).

We now adjust the estimated-displacement of \mathcal{C} in order to obtain an almost transverse family γ_6 such that

$$e = |\gamma_6(k)(1 - \tau) - \gamma_5(k)(1 - \tau)|$$

is small. We require that

$$e \ll h = |\gamma_5(k)(1) - \gamma_5(k)(1 - \tau)|.$$

If that is the case, we can invoke the fact that $\gamma_5(k)$ is transverse in the interval $[1 - \tau, 1]$, to extend $\gamma_6(k)$ to a curve that is transverse in $[1 - \tau, 1]$ and agrees at the end with $\gamma_5(k)$. This preserves embeddedness.

The proof concludes invoking Lemma 4.4.11, which adds a C^1 -small perturbation to γ_6 to yield a family that is transverse. \square

Proof (of Theorem 1.19). The proof follows from the same arguments. As we already observed for horizontal embeddings, handling the controller becomes easier in the immersion case, since self-intersections do not need to be avoided. Because of this reason, the statement also holds in dimension 3.

4.9 Appendix: Technical lemmas on commutators

Given a vector field Z on M we write ϕ_t^Z for its flow at time t . Given a pair of vector fields X and Y on M , we want to compare, in a quantitative manner, the flow of their Lie bracket $\phi_t^{[X,Y]}$ with the commutator of their flows ϕ_t^X and ϕ_t^Y . The contents of this appendix will be an important technical ingredient in the proof of our main theorems.

Given 1-parameter families (not necessarily subgroups) $(\varphi_t)_{t \in \mathbb{R}}$ and $(\psi_t)_{t \in \mathbb{R}}$ in the diffeomorphism group $\text{Diff}(M)$, we can define in a given local chart the map $[\psi(t), \varphi(s)] = \varphi_s \circ \psi_t \circ \varphi_s^{-1} \circ \psi_t^{-1}(x)$. Note that if we take $s = t$ then this map is the commutator of the families taken for each time t , and we denote it by $[\psi_t, \varphi_t] := [\psi(t), \varphi(t)]$.

Lemma 4.9.1 Write $X = \frac{\partial}{\partial t} \Big|_{t=0} \phi_t$ and $Y = \frac{\partial}{\partial t} \Big|_{t=0} \psi_t$ and assume $\phi_0 = \psi_0 = \text{Id}$. Then the following statements hold:

i) There exists a 2-parametric family of diffeomorphisms $\varepsilon_{ts} = o(ts)$ (where $\varepsilon_0 = \text{Id}$ ⁴) such that

$$\psi_t \circ \phi_s(x) = \varepsilon_{ts} \circ \varphi_{ts}^{X+Y}(x)$$

ii) $\frac{\partial^k}{\partial t^k} \Big|_{t,s=0} [\psi(t), \phi(s)](x) = 0$ for any $k \in \mathbb{N}$,

iii) $\frac{\partial^k}{\partial s^k} \Big|_{t,s=0} [\psi(t), \phi(s)](x) = 0$ for any $k \in \mathbb{N}$,

⁴ This will be the case whenever we write ε_{ts} or ε_t

$$iv) \frac{1}{2} \frac{\partial^2}{\partial t \partial s} \Big|_{t=0} [\psi(t), \phi(s)](x) = [X, Y].$$

v) There exists a 2-parametric family of diffeomorphisms $\varepsilon(ts) = o(ts)$ such that

$$[\phi_t, \psi_s] = \varepsilon_{ts} \circ \varphi_{ts}^{[X,Y]}.$$

Proof. Part i) follows from Taylor's Remainder Theorem applied to the composition map $\psi_t \circ \phi_s(x)$, ii) and iii) are obvious, iv) follows from the definition of Lie Bracket and v) follows from an application of Taylor's Remainder Theorem together with ii), iii) and iv).

Remark 4.9.2 A trivial but rather useful observation that we will eventually make use of is the following one. If ϕ_t, ψ_t , are two flows such that for some 1-parametric family of diffeomorphisms ε_t

$$\phi_t = \varepsilon_t \circ \psi_t$$

then ε_t is also a flow since it is the composition of two flows $\varepsilon_t = \psi_t^{-1} \circ \phi_t$. Moreover, if $\varepsilon = o(t)$, then there exists another flow $\tilde{\varepsilon} = o(t)$ such that

$$\psi_t = \tilde{\varepsilon}_t \circ \phi_t.$$

For this last statement just note that

$$\phi_t = \varepsilon_t \circ \psi_t \implies \phi_t \circ \psi_t^{-1} = o(t) \implies \psi_t \circ \phi_t^{-1} = o(t)$$

and thus there exists some flow $\tilde{\varepsilon}_t = o(t)$ such that $\psi_t = \tilde{\varepsilon}_t \circ \phi_t$.

The following Lemma formalizes how errors inside a commutator of 1-parametric families of diffeomorphisms can be taken out the bracket expressions when comparing to the flow of the respective brackets.

Lemma 4.9.3 Consider a flow $\varepsilon(t) = o(t)$ and write $X = \frac{\partial}{\partial t} \Big|_{t=0} \phi_t$, $Y = \frac{\partial}{\partial t} \Big|_{t=0} \psi_t$ and assume $\phi_0 = \psi_0 = Id$. Then there exists a 2-parametric family of diffeomorphisms $\tilde{\varepsilon}(ts) = o(ts)$ such that

$$[\varphi_s^X, \varepsilon_t \circ \varphi_t^Y] = \tilde{\varepsilon}_{ts} \circ \varphi_{ts}^{[A,B]}.$$

Proof. The result follows from a direct application of points i) and v) from Lemma 4.9.1.

Remark 4.9.4 In particular, if we take $s = t$ in this previous lemma, we get that $\tilde{\varepsilon}_{t^2}$ is an actual flow, since it is the composition of two flows as in Remark 4.9.2,

$$\tilde{\varepsilon}_{t^2} = \left(\varphi^{[A,B]} \right)_{t^2}^{-1} \circ [\varphi_t^X, \varepsilon_t \circ \varphi_t^Y]$$

The same remark is true for the family ε_{ts} in v) from Lemma 4.9.1.

Lemma 4.9.5 Consider $X = \frac{\partial}{\partial t} \Big|_{t=0} \phi_t$, $Y = \frac{\partial}{\partial t} \Big|_{t=0} \psi_t$ and assume $\phi_0 = \psi_0 = Id$. Then, if $\phi_t = \varepsilon_t \circ \psi_t$ for certain $\varepsilon_t = o(t)$, then there exists some other $\tilde{\varepsilon}_t = o(t)$ such that

$$\phi_t = \psi_t \circ \tilde{\varepsilon}_t$$

Proof. By Lemma 4.9.3 there exists some $h_t = o(t)$ such that $h_{t^2} \circ \varepsilon_t^{-1} \circ \psi_t^{-1} \circ \varepsilon_t \circ \psi_t = Id$. So, we have that $h_{t^2} \circ \varepsilon_t^{-1} \circ \psi_t^{-1} = \phi_t^{-1}$. The result follows from taking inverse flows at both side of the equation.

Lemma 4.9.6 Write $X = \frac{\partial}{\partial t} \Big|_{t=0} \phi_t$ and $Y = \frac{\partial}{\partial t} \Big|_{t=0} \psi_t$ and assume $\phi_0 = \psi_0 = Id$. Then we have the following estimation:

$$\phi_t^{-1} \circ \psi_s \circ \phi_t = \varepsilon_{ts} \circ \varphi_s^{X+t[X,Y]}.$$

Proof. First, note that

$$\phi_t^{-1} \circ \psi_s \circ \phi_t = \psi_s \circ [\phi_t, \psi_s] = \psi_s \circ \tilde{\varepsilon}_{ts} \circ \varphi_{ts}^{[X,Y]} = \psi_s \circ \tilde{\varepsilon}_{ts} \circ \varphi_s^{t[X,Y]},$$

where the second equality follows by Lemma 4.9.1. Nevertheless, this implies the existence of a 2-parametric family of diffeomorphisms $\epsilon_{ts} = o(ts)$ such that $\psi_s \circ \varphi_s^{t[X,Y]} \circ \epsilon_{ts}$. But, from point i) in Lemma 4.9.1 applied to the composition $\psi_s \circ \varphi_s^{t[X,Y]}$, there exist some other $\tilde{\epsilon}_{ts} = o(ts)$, $\varepsilon_{ts} = o(ts)$ such that

$$\psi_s \circ \varphi_s^{t[X,Y]} \circ \epsilon_{ts} = \varphi_s^{Y+t[X,Y]} \circ \tilde{\epsilon}_{ts} = \varepsilon_{ts} \circ \varphi_s^{Y+t[X,Y]}$$

thus yielding the claim.

Definition 4.9.7 Given two 1-parametric families of diffeomorphisms φ_s, ψ_t , we define their k -th iterated commutator $[\psi_t, \varphi_s]^{\#k}$ as follows

$$[\psi_t, \varphi_s]^{\#k} := \left(\left(\varphi_{\frac{s}{\sqrt{k}}}^{-1} \right) \circ \left(\psi_{\frac{t}{\sqrt{k}}}^{-1} \right) \circ \left(\varphi_{\frac{s}{\sqrt{k}}} \right) \circ \left(\psi_{\frac{t}{\sqrt{k}}} \right) \right)^k.$$

The reason for parametrizing the flows by $\frac{s}{\sqrt{k}}$ in the definition of iterated commutator is justified by the following proposition.

Proposition 4.9.8 Write $X = \frac{\partial}{\partial t} \Big|_{t=0} \phi_t$ and $Y = \frac{\partial}{\partial t} \Big|_{t=0} \psi_t$ and assume $\phi_0 = \psi_0 = Id$. Then there exists $\varepsilon_t = o(t)$ such that

$$[\psi_t, \varphi_s]^{\#k} = \varepsilon_{ts} \circ \varphi_{ts}^{[X,Y]}$$

Proof. By the definition of the k -th iterated commutator of flows and by Lemma 4.9.1, there exist flows $\varepsilon_t^1 = o(t), \dots, \varepsilon_t^k = o(t)$ such that

$$[\psi_t, \varphi_t]^{\#k} = \left(\varepsilon_{ts/k}^1 \circ \varphi_{ts/k}^{[X,Y]} \right) \circ \dots \circ \left(\varepsilon_{ts/k}^k \circ \varphi_{ts/k}^{[X,Y]} \right)$$

But by Lemma 4.9.5 there exist flows $\tilde{\varepsilon}_t^1 = o(t), \dots, \tilde{\varepsilon}_t^k = o(t)$ such that

$$\left(\varepsilon_{ts/k}^1 \circ \varphi_{ts/k}^{[X,Y]} \right) \circ \dots \circ \left(\varepsilon_{ts/k}^k \circ \varphi_{ts/k}^{[X,Y]} \right) = \tilde{\varepsilon}_{ts/k}^1 \circ \dots \circ \tilde{\varepsilon}_{ts/k}^k \circ \left(\varphi_{ts/k}^{[X,Y]} \right)^k$$

and so the claim follows.

With this battery of technical results at our disposal, we can compare how taking a given bracket expression behaves with respect to taking flows [79, Theorem 1]. Our case includes the iterated ($\#^m$ case) :

Proposition 4.9.9 Let $X_1, X_2, \dots, X_\lambda$ be (possibly repeated) vector fields on a manifold M . Then, for any bracket expression $A(-, \dots, -)$ of length λ there exists a flow $\varepsilon_t = o(t)$ such that

$$A\left(\varphi_t^{X_1}, \dots, \varphi_t^{X_\lambda}\right) = \varepsilon_{t^\lambda} \circ \phi_{t^\lambda}^{A(X_1, \dots, X_\lambda)}.$$

Proof. We proceed by induction on the length of the formal bracket-expression:

- For $k = 2$ the result holds by Proposition 4.9.8.
- The Induction Hypothesis (IH) says that the statement holds for all expressions of length $k' < k$. By definition, if $A(-, \dots, -)$ is an expression of length k , there exists $i < k$ and an integer m such that $A(X_1, \dots, X_k) = [B(X_1, \dots, X_i), C(X_{i+1}, \dots, X_k)]^{\#m}$, with $B()$ of length i and $C()$ of length $k - i$. Computing we see that there are flows $f_t = o(t)$ and $g_t = o(t)$ such that:

$$\begin{aligned} A(\phi_t^{X_1}, \dots, \phi_t^{X_k}) &= \left[B(\phi_t^{X_1}, \dots, \phi_t^{X_i}), C(\phi_t^{X_{i+1}}, \dots, \phi_t^{X_k}) \right]^{\#m} \\ &\stackrel{\text{IH}}{=} \left[f_{t^i} \circ \phi_{t^i}^{B(X_1, \dots, X_i)}, g_{t^{k-i}} \circ \phi_{t^{k-i}}^{C(X_{i+1}, \dots, X_k)} \right]^{\#m} \end{aligned} \quad (4.4)$$

By an application of Proposition 4.9.8 first, and by Lemma 4.9.3, there exists a flow $\varepsilon_t = o(t)$ such that

$$\left[f_t \circ \phi_{t^i}^{B(X_1, \dots, X_i)}, g_t \circ \phi_{t^{k-i}}^{C(X_{i+1}, \dots, X_k)} \right]^{\#m} = \varepsilon_{t^k} \circ \left[\phi_{t^i}^{B(X_1, \dots, X_i)}, \phi_{t^{k-i}}^{C(X_{i+1}, \dots, X_k)} \right] \quad (4.5)$$

But, since $\left[\phi_{t^i}^{B(X_1, \dots, X_i)}, \phi_{t^{k-i}}^{C(X_{i+1}, \dots, X_k)} \right] = \phi_{t^k}^{A(X_1, \dots, X_k)}$ the result follows from combining (4.4) and (4.5).

Part III

Spaces of bracket-generating distributions

Chapter 5

Convex integration with avoidance

5.1 Jets and relations

We now introduce jet spaces (Subsection 5.1.1). We put particular emphasis in the geometry behind principal directions and subspaces (Subsection 5.1.2). This will allow us to discuss differential relations and over-relations and study them “one direction at a time” (Subsection 5.1.3). In Subsection 5.1.4 we introduce foliated analogues of these concepts; these will be used to phrase parametric statements in later Sections.

5.1.1 Jet spaces

Given a smooth fibre bundle $X \rightarrow M$, we denote by $J^r(X)$ its associated space of r -jets. We have projections from the space of r -jets to the space of r' -jets, $r' < r$:

$$\pi_{r'}^r : J^r(X) \rightarrow J^{r'}(X).$$

Furthermore, we write $\pi_M^r : J^r(X) \rightarrow M$ for the base projection. The projection π_{r-1}^r is an affine fibration. Given a section $f : M \rightarrow X$, we write $j^r f : M \rightarrow J^r(X)$ for its r -jet, which is a holonomic section of jet space.

5.1.1.1 Jet spaces in local coordinates

If we work in a local chart of X , we can identify M with a vector space V and the fibres of X with a vector space W . Doing so allows us to identify, locally:

$$J^r(X) \supset J^r(V \times W) \cong V \times W \times \text{Hom}(V, W) \times \text{Sym}^2(V, W) \times \cdots \times \text{Sym}^r(V, W),$$

by sending a Taylor polynomial at a given point in V (an element of the right-hand side) to the jet it represents in $J^r(X)$. Here $\text{Sym}^r(V, W)$ denotes the space of symmetric tensors of order r (i.e. homogeneous polynomials) with entries in V and values in W . In particular, we are identifying the (affine) fibres of π_{r-1}^r with their underlying vector space $\text{Sym}^r(V, W)$.

5.1.2 Principal subspaces

The following notion formalises the idea of two r -jets that agree except along a pure derivative of order r :

Definition 5.1.1 Given a hyperplane $\tau \subset T_p M$, we say that two sections $f, g : M \rightarrow X$ have the same $\perp(\tau, r)$ -jet at $p \in M$ if

$$D_p|_\tau j^{r-1} f = D_p|_\tau j^{r-1} g,$$

where $D_p|_\tau$ means taking the differential at p and restricting it to τ .

When τ is a hyperplane field, the $\perp(\tau, r)$ -jets form a bundle, which we denote by

$$J^{\perp(\tau, r)}(X).$$

There are affine fibrations:

$$\begin{aligned} \pi_{\perp(\tau, r)}^r : J^r(X) &\rightarrow J^{\perp(\tau, r)}(X), \\ \pi_{r-1}^{\perp(\tau, r)} : J^{\perp(\tau, r)}(X) &\rightarrow J^{r-1}(X). \end{aligned}$$

In practice, the hyperplane field τ may be defined only over an open subset U of M , but we will still write $J^{\perp(\tau, r)}(X)$ instead of $J^{\perp(\tau, r)}(X|_U)$. Given a section $f : M \rightarrow X$, we write

$$j^{\perp(\tau, r)} f : M \rightarrow J^{\perp(\tau, r)}(X)$$

for the corresponding section of $\perp(\tau, r)$ -jets. A section of this form is said to be *holonomic*.

Definition 5.1.2 The fibers of $\pi_{\perp(\tau, r)}^r$ are said to be the **principal subspaces** associated to τ (and r). They are all affine subspaces parallel to one another. Given $z \in J^r(X)$, we write

$$\text{Pr}_{\tau, z} := (\pi_{\perp(\tau, r)}^r)^{-1} \left(\pi_{\perp(\tau, r)}^r(z) \right)$$

for the principal subspace that contains it.

Instead of talking about hyperplanes, it is often convenient to talk about covectors $\lambda \in T^*M$. We then write $\perp(\lambda, r) := \perp(\text{Ker}(\lambda), r)$. When λ is a global 1-form, the bundle $J^{\perp(\lambda, r)}(X)$ is only defined in the open set $\{\lambda \neq 0\}$. However, we can define $\text{Pr}_{\lambda, z}$ everywhere by setting $\text{Pr}_{\lambda, z} = \{z\}$ if $\lambda = 0$.

In the context of convex integration, we will attempt to add to $f : M \rightarrow X$ oscillations of order r along codirections λ ; this will amount to pushing $j^r f$ along $\text{Pr}_{\lambda, j^r f}$.

5.1.2.1 Principal subspaces in coordinates

As in Subsubsection 5.1.1.1, we use vector spaces V and W as the local models for M and the fibre of $X \rightarrow M$, respectively.

Lemma 5.1.3 Consider a codirection $\lambda \in T_0^*V$ and an element $z \in J^r(V \times W)$ lying over $0 \in V$. The principal subspace $\text{Pr}_{\lambda, z}$ is the image of the affine map:

$$\begin{aligned} W &\longrightarrow J^r(V \times W) \\ w &\longmapsto z + \lambda^{\otimes r} \otimes w. \end{aligned}$$

In particular, the vector subspace underlying $\text{Pr}_{\lambda, z}$ is

$$\{\lambda^{\otimes r} \otimes w \mid w \in W\} \subset \text{Sym}^r(V, W),$$

which we call the **principal direction** Pr_λ . Do note that the map from Lemma 5.1.3 is defined for all λ , and in fact varies smoothly with λ , but is an affine isomorphism between W and the principal subspace if and only if $\lambda \neq 0$. An element in Pr_λ is said to be **principal** or **pure**. We recall:

Lemma 5.1.4 $\text{Sym}^r(V, W)$ admits a basis consisting of principal elements.

The Lemma says that any two elements $F, G \in J^r(X)$ lying over the same $H \in J^{r-1}(X)$ differ by a finite sequence of modifications along principal subspaces of order r . See Figure 5.1.




Figure 5.1: The principal cone in $J^2(\mathbb{R}^2, \mathbb{R})$. We use coordinates (x, y) in the base \mathbb{R}^2 . We identify the fibre of $J^2 \rightarrow J^1$ with the symmetric bilinear maps $\text{Sym}^2(\mathbb{R}^2, \mathbb{R})$; this is what we depict. The pure directions are then of the form $\alpha^{\otimes 2}$, with $\alpha \in T^*\mathbb{R}^2$. They are shown forming a blue cone, on the right hand side. We single out three principal vectors: $dx \otimes dx$, $dy \otimes dy$, and $\frac{dx+dy}{\sqrt{2}} \otimes \frac{dx+dy}{\sqrt{2}}$. The cone linearly spans the whole fibre, since these three vectors form a principal basis. The vectors contained in the left-hand-side cone are not principal, as they are of the form $-\alpha^{\otimes 2}$.

5.1.2.2 Principal paths

We have formalised the idea that two jets differ from one another along a single pure derivative by saying that they have same underlying $\perp (\tau, r)$ -jet. We can similarly define the notion of two jets differing by a finite sequence of changes along pure derivatives:

Definition 5.1.5 Fix $z \in J^{r-1}(X)$. A **principal path** over z is a sequence

$$\{F_i \in (\pi_{r-1}^r)^{-1}(z)\}_{i=0, \dots, I}$$

such that $F_{i+1} - F_i$ is principal. We say that I is the **(principal) length** of the path.

Do note that, unless $F_i = F_{i+1}$, the pair (F_i, F_{i+1}) uniquely determines the principal subspace $\text{Pr}_{\lambda, F_i} = \text{Pr}_{\lambda, F_{i+1}}$ containing both.

Fix $z \in J^{r-1}(X)$ and set $p = \pi_M^{r-1}(z) \in M$. According to Lemma 5.1.4, we can fix an ordered collection of hyperplanes $\{\tau_i \subset T_p M\}_{i=0, \dots, I}$ such that the corresponding principal directions span the fibre $(\pi_{r-1}^r)^{-1}(z) \cong \text{Sym}^r(V, W)$. If this collection is minimal, we say that it is a **(principal) basis**; in this case $I = \dim(\text{Sym}^r(V, \mathbb{R}))$. It follows that any two elements $F, G \in (\pi_{r-1}^r)^{-1}(z)$ can be connected by a principal path of length exactly I . Choosing a basis uniquely determines a principal path between F and G . See Figure 5.2.




Figure 5.2: Any direction in the fibre $J^r \rightarrow J^{r-1}$ can be written in a unique manner once a principal basis has been fixed. In this example, as in Figure 5.1, we work in $J^2(\mathbb{R}^2, \mathbb{R})$. Its fibre is identified with $\text{Sym}^2(\mathbb{R}^2, \mathbb{R})$ and we fix $\{(dx + dy) \otimes (dx + dy), dx \otimes dx, dy \otimes dy\}$ as principal basis. This provides a preferred principal path between any two given vectors. In the figure we show a vector in green and how it connects to the origin using a path (which in this concrete case is of length 2).

5.1.3 Over-relations

We are interested in finding and classifying solutions of differential relations. More generally, we define:

Definition 5.1.6 Let $X \rightarrow M$ be a fibre bundle. An **over-relation** of order r is a smooth manifold \mathcal{R} endowed with a smooth map

$$\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X),$$

that we sometimes call the **anchor**. If $\iota_{\mathcal{R}}$ is an inclusion, we will say that \mathcal{R} is a **differential relation**. The over-relation \mathcal{R} is said to be **open** if the map $\iota_{\mathcal{R}}$ is a submersion.

A **formal solution** of \mathcal{R} is a section $F : M \rightarrow \mathcal{R}$ of $\pi_M^r \circ \iota_{\mathcal{R}}$. A formal solution is a genuine **solution** if $\iota_{\mathcal{R}}(F)$ is holonomic.

Observe that the map $\pi_M^r \circ \iota_{\mathcal{R}} : \mathcal{R} \rightarrow M$ need not be a fibration if \mathcal{R} is open. It is, however, a *microfibration*, meaning that the homotopy lifting property holds for small times.

Remark 5.1.7 More generally, Gromov defines open over-relations as those where $\iota_{\mathcal{R}}$ is a microfibration [59, p. 175] but not necessarily submersive. Such generality is unnecessary for our purposes.

A trivial but key observation is the following:

Lemma 5.1.8 Let $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ be an open over-relation and let $\mathcal{R}' \subset \mathcal{R}$ be an open subset. Then $\iota_{\mathcal{R}}|_{\mathcal{R}'}$ is an open over-relation as well.

The main motivating example for our usage of over-relations is the following:

Example 5.1.9 Given an over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ of order r , the projection

$$\pi_{r'}^r \circ \iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^{r'}(X)$$

is an over-relation of order $r' < r$. That is, over-relations are crucial if we want to construct solutions inductively on r .

Observe that openness of $\iota_{\mathcal{R}}$ implies openness of \mathcal{R} , since the maps π_r^r are submersive.

We need to understand how over-relations $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ relate to our idea of introducing oscillations along a given principal subspace. Given $z \in J^r(E)$, we write

$$\mathcal{R}_{\tau,z} := \iota_{\mathcal{R}}^{-1}(\text{Pr}_{\tau,z})$$

for the restriction of \mathcal{R} to the principal subspace containing z . Given $F \in \mathcal{R}$, we write

$$\text{Pr}_{\tau,F} := \text{Pr}_{\tau,\iota_{\mathcal{R}}(F)} \quad \mathcal{R}_{\tau,F} := \iota_{\mathcal{R}}^{-1}(\text{Pr}_{\tau,F}).$$

We use analogous notation when dealing with covectors λ instead of hyperplanes τ .

5.1.4 The foliated setting

One can generalise all the discussion up to this point to differential relations that vary with respect to a parameter. The language of foliations is convenient for this purpose.

We fix a foliated manifold (N, \mathcal{F}) and a bundle $Y \rightarrow N$. We write $J^r(Y; \mathcal{F})$ for the bundle of leafwise r -jets (i.e. equivalence classes of sections of Y up to r -order tangency along the leaves of \mathcal{F}). Note that $J^r(Y; \mathcal{F})$ is the disjoint union of all the individual $J^r(Y|_L)$, with L ranging over the leaves of \mathcal{F} , endowed with the natural smooth structure.

Apart from the usual projections among these bundles for varying r , we have a forgetful map

$$\pi^{\mathcal{F}} : J^r(Y) \rightarrow J^r(Y; \mathcal{F}),$$

that just remembers the leafwise jets. In particular, if L is a leaf of \mathcal{F} we obtain a projection map $J^r(Y)|_L \rightarrow J^r(Y|_L)$.

A section of $J^r(Y; \mathcal{F})$ is holonomic if its restriction to each leaf is holonomic. A section F of $J^r(Y)$ is leafwise holonomic if the corresponding $\pi^{\mathcal{F}} \circ F|_L$ are holonomic; this is weaker than F itself being holonomic.

Definition 5.1.10 A *foliated over-relation* is a map

$$\iota_{\mathcal{S}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$$

where \mathcal{S} is smooth manifold. We say it is **open** if it is submersive.

We can restrict \mathcal{S} to a leaf L of \mathcal{F} and yield a (standard) over-relation $\mathcal{S}|_L \rightarrow J^r(Y|_L)$.

5.1.4.1 Parametric lifts of over-relations

The most important example of foliated over-relation is the following:

Definition 5.1.11 Let $X \rightarrow M$ be a bundle, $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(Y)$ an over-relation, and K a compact manifold serving as parameter space. Set $N := M \times K$ and write π_M and π_K for the projections mapping to M and K , respectively. Endow N with the foliation \mathcal{F} consisting of the fibres of π_K . Set $Y := X \times K = \pi_M^* X$ and $\mathcal{S} := \mathcal{R} \times K = \pi_M^* \mathcal{R}$.

The **parametric lift** of \mathcal{R} to $M \times K$ is the foliated over-relation

$$\pi_M^* \iota_{\mathcal{R}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F}).$$

Do note that the leaves of \mathcal{F} are copies $M \times \{k\}$ of M and, identifying both using π_M , we have that \mathcal{S} restricts to \mathcal{R} along each $M \times \{k\}$. Families $(F_k)_{k \in K}$ of formal solutions of $M \rightarrow \mathcal{R}$ are then equivalent to formal solutions $F : N \rightarrow \mathcal{S}$. The family $(F_k)_{k \in K}$ consists of holonomic sections if and only if F is (leafwise) holonomic.

We remark:

Lemma 5.1.12 *The parametric lift of an open over-relation is open.*

5.1.4.2 Non-foliated preimage

Any $\iota_{\mathcal{S}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$ defines an over-relation in $J^r(Y)$ by pullback. This will be relevant later on, because it will allow us to rephrase statements about \mathcal{S} to statements about the pullback (therefore reducing the foliated theory to the non-foliated one).

Definition 5.1.13 *Let $\iota_{\mathcal{S}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$ be a foliated over-relation. Its **non-foliated preimage** is the over-relation*

$$\mathcal{S}^* := \{(F, z) \in \mathcal{S} \times J^r(Y) \mid \iota_{\mathcal{S}}(F) = \pi^{\mathcal{F}}(z)\},$$

with anchor map $\iota_{\mathcal{S}}^ : \mathcal{S}^* \rightarrow J^r(Y)$ defined by the expression $\iota_{\mathcal{S}}^*(F, z) := z$.*

I.e. \mathcal{S}^* is the pullback of $J^r(Y) \rightarrow J^r(Y; \mathcal{F})$ to \mathcal{S} . It follows that:

Lemma 5.1.14 *The non-foliated preimage of an open, foliated over-relation is open.*

5.2 Ampleness and convex integration

In this Section we recall some key ideas behind convex integration. Our goal is not to be comprehensive, but rather to fix notation and discuss its different incarnations, as introduced by Gromov [59]. We also borrow from Spring [96] and Eliashberg-Mishachev [39].

We first recall the three standard flavours of ampleness; each of them is the basis of a concrete implementation of convex integration. Classic ampleness is explained in Subsection 5.2.2. Ampleness along principal frames (often called ampleness in coordinate directions) is explained in Subsection 5.2.3. Ampleness in the sense of convex hull extensions appears in Subsection 5.2.4.

We then compare them in Subsection 5.2.5. This will clarify how ampleness up to avoidance (to appear in Section 5.3) fits within this greater context.

5.2.1 Ampleness in affine spaces

We define ampleness for subsets of affine spaces first. We adapt it to relations in jet spaces in upcoming Subsections.

Definition 5.2.1 *Let A be an affine space. Then:*

- *Let $B \subset A$ be a subset. Given $b \in B$, we write B_b for the path-component containing it. We say that B is **ample** if the convex hull $\text{Conv}(B, b) := \text{Conv}(B_b)$ of each $B_b \subset B$ is the whole of A .*

- Let C be a topological space and $\iota : C \rightarrow A$ be a continuous map. The map ι is **ample** if $\text{Conv}(C, c) := \text{Conv}(\iota(C_c)) = A$ for each $c \in C$.

Furthermore, we say that *ampleness holds trivially* if for each $c \in C$ either $\iota(C_c) = \emptyset$ or $\iota(C_c) = A$.

A particularly relevant case in the examples to come is the following:

Example 5.2.2 A stratified subset $\Sigma \subset \mathbb{R}^n$ of codimension at least 2 is said to be *thin*. Its complement is ample. See Figure 5.3.




Figure 5.3: Example of a thin set $L \subset \mathbb{R}^3$. Any point in L is a convex combination of points in the complement. One such example is shown in the image, where three black points in $\mathbb{R}^3 \setminus L$ convexly generate a point in L (surrounded by red lines).

Not all ample subsets have thin complements. The following example shows an ample subset whose complement has codimension one:

Example 5.2.3 The subset of \mathbb{R}^3 defined by

$$\mathcal{H}^- := \{(x, y, z) \in \mathbb{R}^3 \mid xy - z^2 < 0\}$$

is the outer-component of a cone. It is ample and thus any point $p \in \mathbb{R}^3$ can be expressed as a convex combination of points in \mathcal{H}^- . The remaining part of the complement of the cone

$$\mathcal{H}^+ := \{(x, y, z) \in \mathbb{R}^3 \mid xy - z^2 > 0\}$$

consists of two components, neither of which is ample.

The set \mathcal{H}^- will reappear later on in our study of hyperbolic $(4, 6)$ distributions. \mathcal{H}^+ appears in the study of elliptic $(4, 6)$ distributions. See Section 7.2.

5.2.2 Ampleness in all principal directions

We now define the most commonly used notion of ampleness for differential relations. It is also the most restrictive one.

Definition 5.2.4 Fix a bundle $X \rightarrow M$ and an over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$. Let $\lambda \in T_p^*M$ be a covector. We say that $\iota_{\mathcal{R}}$ is

- **ample along the principal direction determined by λ** if, for every $F \in \mathcal{R}$ projecting to p , the map $\iota_{\mathcal{R}} : \mathcal{R}_{\lambda,F} \rightarrow \text{Pr}_{\lambda,F}$ is ample.
- **ample in all principal directions** if the over-relations $(\pi_{r'}^T \circ \iota_{\mathcal{R}})_{r'=1,\dots,r}$ are ample along all non-zero covectors λ .

Being the most commonly used flavour, we sometimes just say that $\iota_{\mathcal{R}}$ is **ample**. Gromov's convex integration is usually stated as:

Theorem 5.1. *The complete C^0 -close h -principle holds for any open over-relation that is ample in all principal directions.*

This result was first proven, only for first order, in [58, Corollary 1.3.2]. The statement for all orders appeared later in [59, Section 2.4, p. 180]. The first order case is treated as well in [39, Part 4].

5.2.2.1 The foliated setting

Fix a bundle $Y \rightarrow (N, \mathcal{F})$ and a foliated over-relation $\iota_{\mathcal{S}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$. We say that \mathcal{S} is **ample along all foliated principal directions** if, for each leaf L , the restriction $\mathcal{S}|_L$ satisfies Definition 5.2.4.

By construction, the ampleness of the non-foliated preimage $\mathcal{S}^* \rightarrow J^r(Y)$ of \mathcal{S} can be read purely along \mathcal{F} :

Lemma 5.2.5 Fix a leaf L , a point $p \in L$, a formal datum $z \in J^r(Y)|_p$, and a codirection $\lambda \in T_p^*N$. Write $z' \in J^r(Y|_L)$ for the leafwise jet of z and λ' for the restriction $\lambda|_L \in T_p^*L$.

The following conditions are equivalent:

- \mathcal{S}^* is ample along the principal subspace $\text{Pr}_{\lambda,z} \subset J^r(Y)$.
- $\mathcal{S}|_L$ is ample along the principal subspace $\text{Pr}_{\lambda',z'} \subset J^r(Y|_L)$.

It immediately follows that:

Corollary 5.2.6 *The complete C^0 -close h -principle holds for any open, foliated over-relation that is ample in all foliated principal directions.*

A particularly beautiful consequence of these statements is the following. Suppose K is a compact manifold, $X \rightarrow M$ is a bundle, and $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ is an open over-relation that is ample along all principal directions. Then, the parametric lift $\mathcal{R} \times K$ is, by definition, ample along all foliated principal directions. Applying Lemma 5.2.5 we deduce that $(\mathcal{R} \times K)^* \rightarrow J^r(X \times K)$ is ample along all principal directions. It follows that, in order to prove Theorem 5.1 for arbitrary parameters (and relatively in parameter and domain), it is sufficient to prove the non-parametric version (relatively in domain). Indeed: the K -parametric statement for \mathcal{R} is just the 0-parametric statement for $\mathcal{R} \times K$.

5.2.3 Ampleness along principal frames

As we pointed out in the Introduction, we do not need ampleness in all directions, since it is sufficient to be able to proceed over a base in the space of derivatives. This motivates the following definitions.

Definition 5.2.7 A *locally-defined hyperplane field* is a pair (U, τ) consisting of an open set $U \subset M$ and a germ of hyperplane field τ along the closure \bar{U} .

The hyperplane field (U, τ) is **integrable** if τ integrates to a codimension-1 foliation.

Our hyperplane fields will live on charts and therefore they will always be locally-defined. The condition that τ is a germ along the closure \bar{U} is included to make some of our later statements cleaner. The reader can think of \bar{U} as being some closed ball in M . Often, we just write τ and we leave U implicit; we say that U is the *support* of τ .

Definition 5.2.8 A *principal frame* of order r is a collection C of locally-defined hyperplane fields satisfying:

- All of the fields in C are integrable and have the same support U .
- C is a principal basis in each of the fibres of π_{r-1}^r lying over U .

A *principal cover* of order r is a collection \mathcal{C} of principal frames of order r whose supports cover M .

A principal direction/subspace defined by a hyperplane in \mathcal{C}/C will be called a \mathcal{C}/C -principal direction/subspace.

The second flavour of ampleness reads:

Definition 5.2.9 An over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ is **ample along principal frames in order r** if each point $p \in M$ admits an r -order principal frame C with support $U \ni p$ such that \mathcal{R} is ample along all C -principal directions.

The over-relation $\iota_{\mathcal{R}}$ is **ample along the principal cover \mathcal{C}** if \mathcal{R} is ample along all \mathcal{C} -principal directions.

The over-relation $\iota_{\mathcal{R}}$ is **ample along principal frames** if the relations $(\pi_{r'}^r \circ \iota_{\mathcal{R}})_{r'=1, \dots, r}$ are ample along principal frames in their respective order.

If $r = 1$ (or if $r > 1$ but $\dim(M) = 1$), a principal cover can be obtained by picking a covering of M and in each chart setting $\{\tau_i = \text{Ker}(dx_i)\}$, where $\{dx_i\}$ is the coordinate coframe. For this reason, when one deals with first order jets, ampleness along a principal cover is also called *ampleness along coordinate directions*; see [58, Definition 1.2.6] and [39, p. 167].

Theorem 5.2. The complete C^0 -close h -principle holds for any open over-relation that is ample along principal frames.

For first order, this is the main result in [58, Theorem 1.3.1]; it appears in [39, p. 172] as well. For arbitrary order, it follows from [59, p. 179, Principal Stability Theorem C]. An alternate implementation for arbitrary order appeared in the master thesis [?]; it avoids convex hull extensions and adapts instead the idea from [58].

5.2.3.1 The foliated setting

Consider a bundle $Y \rightarrow (N, \mathcal{F})$ and a foliated over-relation $\mathcal{S} \rightarrow J^r(Y; \mathcal{F})$. It is possible to adapt Definition 5.2.9 to the foliated setting by relying on principal covers that consist of leafwise hyperplane fields. This is ultimately unnecessary for us, so we leave the details to the reader. The same comment applies to the next section.

5.2.4 Ampleness in the sense of convex hull extensions

If an (over)-relation is ample along principal frames, all formal solutions can be made holonomic, one derivative at a time, by introducing oscillations along the codirections given by the frames. However, one can imagine a situation where different formal solutions need oscillations along different principal frames or even oscillations along collections of codirections that do not form a frame at all.

In order to formalise this idea, we introduce the concept of convex hull extensions:

Definition 5.2.10 *Let $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ be an over-relation. Its **convex hull extension** is the set*

$$\text{Conv}(\mathcal{R}) := \{(F, \lambda, z) \in \mathcal{R} \times_M T^*M \times_M J^r(X) \mid z \in \text{Conv}(\mathcal{R}_{\lambda, F}, F)\}$$

with anchor map $(F, \lambda, z) \mapsto z$.

Remark 5.2.11 *Definition 5.2.10 differs slightly from the definition of convex hull extension given in [59, p. 177]. The reason is that we want our open over-relations to be manifolds that submerge onto jet space (instead of more general microfibrations with domain a (quasi)topological space). Assuming \mathcal{R} itself was a manifold, Gromov's convex hull extension would still yield instead a topological space with conical singularities. The upcoming convex integration statements are unaffected by this change.*

We observe:

Lemma 5.2.12 *Suppose \mathcal{R} is open. Then, $\text{Conv}(\mathcal{R})$ is an open over-relation. In particular, its underlying space is a smooth manifold.*

Proof. Let $W \rightarrow J^r(X)$ be the pullback of the vertical tangent space of X ; i.e. the subspace of TX consisting of vectors tangent to the fibres of $X \rightarrow M$. Using Lemma 5.1.3 we observe that the space

$$A = \{(F, \lambda, z) \in \mathcal{R} \times_M T^*M \times_M J^r(X) \mid z \in \text{Pr}_{\lambda, F}\}$$

is a smooth fibre bundle over $\mathcal{R} \times_M T_p^*M$ with affine fibre isomorphic to W . Using the Lemma once more, we see that the anchor map $A \rightarrow J^r(X)$ given by $(F, \lambda, z) \mapsto z$ can equivalently be written as

$$(F, \lambda, z = \iota_{\mathcal{R}}(F) + \lambda^{\otimes r} \otimes w) \mapsto \iota_{\mathcal{R}}(F) + \lambda^{\otimes r} \otimes w,$$

which is a submersion because \mathcal{R} itself was. The proof is complete noting that $\text{Conv}(\mathcal{R})$ is an open subset of A , due to the openness of \mathcal{R} .

We write $\text{Conv}^l(\mathcal{R})$ for the l -fold convex hull extension of \mathcal{R} . An element in $\text{Conv}^l(\mathcal{R})$ is then an element $F \in \mathcal{R}$, together with a principal path of length l starting at $\iota_{\mathcal{R}}(F)$. A section of $\text{Conv}^l(\mathcal{R})$ is thus a smoothly-varying choice of principal path at each point. Do note that the hyperplanes associated to such paths vary smoothly, but need not be integrable; for this reason, it is convenient to restrict our attention to the following nice subclass of sections:

Definition 5.2.13 A section $(F, \lambda_1, z_1, \dots, \lambda_l, z_l) : M \rightarrow \text{Conv}^l(\mathcal{R})$ is said to be **integrable** if the λ_i are integrable.

Do note that the λ_i are allowed to vanish and thus we speak of integrability in the locus $\{\lambda_i \neq 0\}$.

The following definition corresponds to the idea of being able to connect, using convex hull extensions, a formal datum F to the holonomic section corresponding to its zero order part $\pi_0^r \circ \iota_{\mathcal{R}} \circ F$.

Definition 5.2.14 A formal solution $F : M \rightarrow \mathcal{R}$ is (integrably) **short** if, for some l , there is a (integrable) holonomic solution $G : M \rightarrow \text{Conv}^l(\mathcal{R})$ of the form $(F, \dots, j^r(\pi_0^r \circ \iota_{\mathcal{R}} \circ F))$.

Assume that F is holonomic in a neighbourhood of a closed subset $M' \subset M$. Then, F is **short relative to** M' if the codirections $\{\lambda_i\}_{i=1}^l$ associated to $G = (F, \lambda_1, z_1, \dots, \lambda_l, z_l)$ can be chosen to be zero over $\mathcal{O}_p(M')$.

Do note that if a solution is short it is already a solution of $\pi_{r-1}^r \circ \iota_{\mathcal{R}}$. That we need such an assumption is not surprising, since the convex hull extension machinery works purely in order r . As such, in order to provide a full flexibility statement, we need to consider convex hull extensions for each $\pi_{r'}^r \circ \iota_{\mathcal{R}}$, $r' = 1, \dots, r$.

The last flavour of ampleness comes in two slightly different incarnations:

Definition 5.2.15 An over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ is **ample in the sense of (integrable) convex-hull extensions** if the following property holds: Fix

- An order $r' = 1, \dots, r$,
- A compact manifold K ,
- A K -family of formal solutions $F : M \times K \rightarrow \mathcal{R} \times K$ that is holonomic of order $r' - 1$.

Then, the family F is (integrably) short for $\pi_{r'}^r \circ \iota_{\mathcal{R}}$, relative to the regions in which it is already r' -holonomic.

Convex integration, in full generality, reads:

Theorem 5.3. The complete C^0 -close h -principle holds for any open over-relation that is ample in the sense of (integrable) convex hull extensions.

This result, assuming integrability, is presented in detail in [96, p. 123, Theorem 8.4]. The statement, without integrability, was already implicit in [59, p. 179, Principal Stability Theorem C]. The integrability hypothesis is restrictive; we discuss it further in the next Subsection.

5.2.5 A comparison of the different incarnations of ampleness

As stated earlier, classic ampleness (ampleness in all principal directions) is the most restrictive of the notions we have introduced. Indeed, it is immediate that Theorem 5.2 implies Theorem 5.1. Furthermore, Theorem 5.3 implies both: Ampleness along a frame says that we can connect any formal solution F to $j^r(\pi_0^r \circ \iota_{\mathcal{R}} \circ F)$ using the given principal frames, proving integrable shortness of F . This works for families and relatively as well.

It is obvious that ampleness in the sense of integrable convex hull extensions is more restrictive than the version without integrability. In particular, ampleness in the sense of convex hull extensions is the most general of the four definitions given.

As we wrote above, Theorem 5.3, without integrability, is contained in Gromov's text implicitly; it can be deduced from [59, p. 179, Principal Stability Theorem C]. In [96], Spring works always under integrability assumptions; this allows him to directly invoke 1-dimensional convex integration in the foliation charts associated to the integrable hyperplane fields appearing in the definition of integrable shortness.

The key claim that Gromov uses to drop integrability, see [59, p. 177], is that any *continuous* hyperplane field can be piecewise approximated by foliation charts. This can then be used to approximate any section of $\text{Conv}^l(\mathcal{R})$ by an integrable one (at the expense of increasing l). We interpret this as an *h-principle without homotopical assumptions* (see also Remark 5.4.5) saying that there is a weak equivalence between the space of sections and the subspace of integrable ones. We have not checked this claim in detail. In fact, it is not important for our results:

Remark 5.2.16 *None of the results from this thesis rely on Definition 5.2.15 or Theorem 5.3. Our arguments reduce the h-principle for relations that are ample up to avoidance to the h-principle for relations that are ample along a principal frame (Theorem 5.2). Nonetheless, in Corollary 5.4.4 we prove that a relation that is ample up to avoidance is ample in the sense of integrable convex hull extensions.*

5.2.5.1 Computability of ampleness in all principal directions

Ampleness in all principal directions is the most restrictive but easiest to check of the four incarnations. The reason is that it is *pointwise* in nature: we just go through each fibre of $J^r(X)$ checking ampleness, one principal direction at a time. In practice, one often deals with Diff-invariant relations described as the complement of some fibrewise (semi-)algebraic condition (which we call the *singularity*). It is then sufficient to check a single fibre of $J^r(X)$ and a single codirection; the problem boils down then to checking the intersection of the singularity with each principal subspace. In practice, this can be already quite involved unless the relation is relatively simple.

Remark 5.2.17 *In [78], P. Massot and M. Theillière prove that convex integration can be used to prove holonomic approximation in spaces of 1-jets. This is a beautiful application of classic convex integration in which checking ampleness is highly non-trivial.*

Classic ampleness turns out to be limited in its applications. In practice, we only encounter it if all formal solutions $F \in \mathcal{R}$ present some form of symmetry guaranteeing that they sit equally nicely with respect to all codirections $\lambda \in T^*M$; we will see this in examples in Section 6.1. The relation defining hyperbolic (4,6) distributions, despite being Diff-invariant, does not satisfy this. We expect most differential relations of codimension-1 not to satisfy it.

5.2.5.2 Computability of ampleness along principal frames

Ampleness along principal frames turns out to be not so different from ampleness in all directions. The two are equivalent if we assume Diff-invariance. In terms of computability, once a concrete principal cover \mathcal{C} is given, checking \mathcal{C} -ampleness is, by definition, easier than checking it in all directions.

5.2.5.3 Computability of ampleness in the sense of convex hull extensions

Ampleness in the sense of convex hull extensions is incredibly general, but notoriously difficult to check. The reason is that it is not a pointwise condition: A formal solution F is short if we can connect it, using a smooth family of principal paths G , to its underlying holonomic section; both F and G are global objects.

Suppose we want to construct G and thus prove that F is short. Convex integration is local in nature, so we try to find a suitable cover of M to proceed. Given a point $p \in M$, we may be able to define $G(p)$ and then extend it locally, by openness, to some open neighbourhood U_p . Finding $G(p)$ is a pointwise process and does not need ampleness in all principal *directions*; it is sufficient to find a suitable sequence of ample principal *subspaces* starting from F . We do this for all p and we extract a cover $\{U_i\}$ with a section G_i defined over U_i . In order to patch these up, we start with G_1 and we glue it with F using a cut-off close to ∂U_1 . The problem now is that the principal subspaces that behaved nicely with respect to F need not behave nicely with respect to G_1 . In particular, G_2 may not help us at all in the overlap $U_1 \cap U_2$.

Furthermore, unlike the previous two flavours, ampleness in the sense of convex hull extensions is not readily parametric. To deal with families one has to prove that the family in question, as a whole, is short.

Ampleness up to avoidance is designed to deal with these considerations and make the aforementioned sketch of argument work. It is also computable pointwise, as we explain in Subsection 5.3.2.1.

5.3 Avoidance

Let $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ be an open over-relation. Our goal is to construct a so-called *avoidance template* \mathcal{A} associated to \mathcal{R} ; if we succeed in constructing \mathcal{A} , we will say that \mathcal{R} is *ample up to avoidance*. Our main Theorem 1.8, whose proof we postpone to the next Section, says that this is a sufficient condition for the h -principle to hold.

Templates (and more general objects called pre-templates) are introduced in Subsections 5.3.2 and 5.3.3. These definitions require us to introduce some auxiliary notation about configurations of hyperplanes; this is done in Subsection 5.3.1. In Subsection 5.3.5 we present some simple constructions of pre-templates. These constructions can yield empty pre-templates when \mathcal{R} is very far from being ample; this is explained in Subsection 5.3.6.

5.3.1 Configurations of hyperplanes

Given a positive integer a and a vector space V , we write

$$\text{H-Conf}_a(V) := \{(H_1, \dots, H_a) \in (\mathbb{P}V^*)^a \mid H_i \neq H_j, \text{ for all } i \neq j\} / \Sigma_a.$$

I.e. the smooth, non-compact manifold consisting of all unordered configurations $[H_1, \dots, H_a]$ of a distinct hyperplanes in V . Its non-compactness is due to collisions (i.e. any sequence in which H_i approaches H_j has no convergent subsequence). In order to consider collections of arbitrary finite cardinality, we consider the union:

$$\text{H-Conf}(V) := \coprod_{a=0}^{\infty} \text{H-Conf}_a(V),$$

where $\text{H-Conf}_0(V) := \{\emptyset\}$ is the space containing only the empty configuration.

Given two configurations $\Xi, \Xi' \in \text{H-Conf}(V)$ we will write $\Xi' \subset \Xi$ if every hyperplane in the former is contained in the latter.

5.3.1.1 Repetitions

In practice, we will deal with ordered collections of hyperplanes that may have repetitions. Concretely, these correspond to points in the closed manifold

$$\overline{\text{H-Conf}}_a(V) := (\mathbb{P}V^*)^a.$$

Consider the open dense subset $\overline{\text{H-Conf}}_a^*(V) \subset \overline{\text{H-Conf}}_a(V)$ consisting of those collections with no repetitions. Its complement is an algebraic subvariety. By construction, we have a quotient map

$$\pi : \overline{\text{H-Conf}}_a^*(V) \longrightarrow \text{H-Conf}_a(V)$$

whose fibres are isomorphic to the symmetric group Σ_a . As before, we write

$$\overline{\text{H-Conf}}(V) := \prod_{a=0}^{\infty} \overline{\text{H-Conf}}_a(V), \quad \overline{\text{H-Conf}}^*(V) := \prod_{a=0}^{\infty} \overline{\text{H-Conf}}_a^*(V),$$

where $\overline{\text{H-Conf}}_0(V)$ and $\overline{\text{H-Conf}}_0^*(V)$ are the singleton set $\{\emptyset\}$.

5.3.1.2 Bundles of configurations

Fix a manifold M . We write

$$\text{H-Conf}(TM) \rightarrow M$$

for the smooth fibre bundle with fibre $\text{H-Conf}(T_p M)$ at a given $p \in M$. Similarly, we write $\overline{\text{H-Conf}}(TM)$ and $\overline{\text{H-Conf}}^*(TM)$. By construction, we have a quotient map

$$\overline{\text{H-Conf}}^*(TM) \longrightarrow \text{H-Conf}(TM)$$

given by the fibrewise action of the symmetric groups.

5.3.2 Avoidance templates and ampleness

Fix a bundle $X \rightarrow M$, an over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$, and a subset \mathcal{A} of the fibered product $\mathcal{R} \times_M \text{H-Conf}(TM)$.

Given a family of hyperplanes $\Xi \in \text{H-Conf}(T_p M)$, we write

$$\mathcal{A}(\Xi) := \mathcal{A} \cap (\mathcal{R} \times_M \{\Xi\}).$$

Using the canonical identification $\mathcal{R} \times_M \{\Xi\} \cong \mathcal{R}_p$, we regard $\mathcal{A}(\Xi)$ as a subset of the fibre \mathcal{R}_p lying over $p \in M$. If we are given a collection of hyperplane fields $\Xi : M \rightarrow \text{H-Conf}(TM)$ instead, we will similarly write $\mathcal{A}(\Xi)$ for the union of all the subsets $\mathcal{A}(\Xi(p))$ as p ranges over the entirety of M . In this case, $\mathcal{A}(\Xi)$ is a subset of \mathcal{R} . If it is a smooth submanifold, the map $\iota_{\mathcal{R}} : \mathcal{A}(\Xi) \rightarrow J^r(E)$ is an over-relation.

Given some $F \in \mathcal{R}$ lying over a point p , we similarly denote

$$\mathcal{A}(F) := \mathcal{A} \cap (\{F\} \times \text{H-Conf}(T_p M)).$$

As before, we regard $\mathcal{A}(F)$ as the subset of $\text{H-Conf}(T_p M)$ consisting of those Ξ such that $F \in \mathcal{A}(\Xi)$. If F is instead a section $M \rightarrow \mathcal{R}$, $\mathcal{A}(F)$ will be the subset of $\text{H-Conf}(TM)$ given by the union of all $\mathcal{A}(F(p))$, as p ranges over all points in M .

Definition 5.3.1 *An open subset $\mathcal{A} \subset \mathcal{R} \times_M \text{H-Conf}(TM)$ is an **(avoidance) pre-template** if the following property holds:*

I. If $\Xi' \subset \Xi \in \text{H-Conf}(TM)$ is a subconfiguration, then $\mathcal{A}(\Xi) \subset \mathcal{A}(\Xi')$.

*The pre-template \mathcal{A} is an **(avoidance) template** if, additionally:*

II. Given $\Xi \in \text{H-Conf}(TM)$, $\mathcal{A}(\Xi)$ is ample along the principal directions determined by Ξ .

III. Given $F \in \mathcal{R}$ lying over $p \in M$, $\mathcal{A}(F)$ is dense in each $\text{H-Conf}_m(T_p M)$.

Property (I) guarantees coherence: removing hyperplanes from Ξ makes the relation $\mathcal{A}(\Xi)$ bigger. In particular, if $\mathcal{A}(\Xi)$ is ample along Ξ , then $\mathcal{A}(\Xi')$ is ample along $\Xi' \subset \Xi$.

Our main definition reads:

Definition 5.3.2 *An open over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ is said to be **ample up to avoidance** if each of the over-relations*

$$(\pi_{r'}^r \circ \iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^{r'}(X))_{r'=1, \dots, r}$$

admits an avoidance template.

Observe that $\mathcal{R} \times_M \text{H-Conf}(TM)$ is an avoidance template if and only if \mathcal{R} is ample in all principal directions.

5.3.2.1 Computability of avoidance

We stated in Subsection 5.2.5 that ampleness up to avoidance is as computable as classic ampleness. There are two parts to this claim.

First we note that verifying whether a given open subset $\mathcal{A} \subset \mathcal{R} \times_M \text{H-Conf}(TM)$ is a template boils down to pointwise checks. Property (I) is often given by construction. Property (III) is often checked together with openness and follows as soon as the complement of $\mathcal{A}(F)$ is given, fibrewise, by some algebraic equality. Property (II) is the most involved, but it is no different from checking ampleness along a principal frame.

The second part of the claim is that the construction of templates is algorithmic. Indeed, we present two possible constructions in Subsection 5.3.5. However, the reader should just think of these as rough guidelines. In practice (for instance, in the proof of Theorem 1.12), one needs to make adjustments in order to produce a template. Still, the adjustments that need to be made are somewhat standard; see Remark 7.2.8.

Lastly, we observe that ampleness up to avoidance is parametric in nature, much like classic convex integration. Namely, given a template \mathcal{A} for \mathcal{R} , we can define an associated foliated template for any parametric lift $\mathcal{R} \times K$; see Subsection 5.3.4. The parametric version of Theorem 1.8 will follow then from the non-parametric one.

5.3.3 Lifted avoidance templates

Definition 5.3.1 is intuitive conceptually but, in practice (see the proofs of Propositions 5.4.2 and 5.4.3), it is often more convenient to deal with the following notion:

Definition 5.3.3 Let π be the quotient map $\overline{\text{H-Conf}}^*(TM) \rightarrow \text{H-Conf}(TM)$. We write

$$\bar{\mathcal{A}} \subset \mathcal{R} \times_M \overline{\text{H-Conf}}^*(TM) \subset \mathcal{R} \times_M \overline{\text{H-Conf}}(TM)$$

for the preimage of a given subset

$$\mathcal{A} \subset \mathcal{R} \times_M \text{H-Conf}(TM).$$

Given $\Xi \in \overline{\text{H-Conf}}(TM)$, we write $\mathcal{A}(\Xi) := \mathcal{A}(\pi(\Xi))$. Similarly, given $F \in \mathcal{R}$, we write $\bar{\mathcal{A}}(F)$ for the preimage by π of $\mathcal{A}(F)$.

We remark:

Lemma 5.3.4 Fix a subset $\mathcal{A} \subset \mathcal{R} \times_M \text{H-Conf}(TM)$. Then, \mathcal{A} is a pre-template if and only if

- $\bar{\mathcal{A}}$ is open.
- $\bar{\mathcal{A}}$ is invariant under the action of the permutation groups Σ_* .
- I. Consider $\Xi', \Xi \in \overline{\text{H-Conf}}(TM)$. Suppose $\pi(\Xi')$ is a subconfiguration of $\pi(\Xi)$. Then $\bar{\mathcal{A}}(\Xi) \subset \bar{\mathcal{A}}(\Xi')$.

Furthermore, \mathcal{A} is a template if and only if, additionally:

- II. Given $\Xi \in \text{H-Conf}(TM)$, $\bar{\mathcal{A}}(\Xi)$ is ample along the principal directions determined by Ξ .
- III. Given $F \in \mathcal{R}$ lying over $p \in M$, $\bar{\mathcal{A}}(F)$ is dense in each $\overline{\text{H-Conf}}_m(T_p M)$.

Proof. First note that $\overline{\text{H-Conf}}^*(TM) \subset \overline{\text{H-Conf}}(TM)$ is open. Its complement, which is an algebraic variety and thus of positive codimension, consists of all configurations that involve repetitions. The claim follows from this fact and the observation that π is a quotient map.

Conversely, any open, Σ_* -invariant subset of $\mathcal{R} \times_M \overline{\text{H-Conf}}^*(TM)$ is the $\bar{\mathcal{A}}$ of some template \mathcal{A} as long as Properties (I), (II) and (III) hold.

5.3.4 Foliated templates

We will prove in Section 5.4 that the parametric analogue of Theorem 1.8 follows from Theorem 1.8 itself. Compare this to Theorem 5.1 and Corollary 5.2.6. This is best implemented using the foliated setting, which we now introduce.

Fix a foliated manifold (N, \mathcal{F}) , a bundle $Y \rightarrow N$, and an over-relation $\iota_{\mathcal{S}} : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$. We look at subsets $\mathcal{A} \subset \mathcal{S} \times_N \text{H-Conf}(\mathcal{F})$. We define $\mathcal{A}(\Xi)$ and $\mathcal{A}(F)$ in the obvious manner. Then:

Definition 5.3.5 An open subset $\mathcal{A} \subset \mathcal{S} \times_N \text{H-Conf}(\mathcal{F})$ is a **foliated pre-template** if the following property holds:

- I. If $\Xi' \subset \Xi \in \text{H-Conf}(\mathcal{F})$ is a subconfiguration, then $\mathcal{A}(\Xi) \subset \mathcal{A}(\Xi')$.

The pre-template \mathcal{A} is a **foliated template** if, additionally:

- II. Given $\Xi \in \text{H-Conf}(\mathcal{F})$, $\mathcal{A}(\Xi)$ is ample along the principal directions determined by Ξ .
- III. Given $F \in \mathcal{S}$ lying over $p \in N$, $\mathcal{A}(F)$ is dense in each $\text{H-Conf}_m(\mathcal{F}_p)$.

The following observation follows immediately from the leafwise nature of Definition 5.3.5:

Lemma 5.3.6 *Let $X \rightarrow M$ be a bundle and $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ an over-relation. Fix a compact manifold K . Suppose \mathcal{R} admits a template \mathcal{A} . Then the parametric lift $\mathcal{R} \times K$ admits a foliated template $\mathcal{A} \times K$.*

Furthermore:

Lemma 5.3.7 *Let $Y \rightarrow (N, \mathcal{F})$ be a bundle over a foliated manifold. If an over-relation $\mathcal{S} \rightarrow J^r(Y; \mathcal{F})$ admits a foliated template \mathcal{A} , its non-foliated preimage $\mathcal{S}^* \rightarrow J^r(Y)$ admits a template \mathcal{A}^* .*

Proof. We define \mathcal{A}^* as a subset of $\mathcal{S}^* \times_N \text{H-Conf}(TN)$. Consider the subspace $\text{H-Conf}'(TN)$ of $\text{H-Conf}(TN)$ consisting of those configurations $[H_1, \dots, H_a] \in \text{H-Conf}(TN)$ that satisfy:

- All $H_i \in [H_1, \dots, H_a]$ intersect \mathcal{F} transversely.
- For all $i \neq j$, the intersections $H_i \cap \mathcal{F}$ and $H_j \cap \mathcal{F}$ are distinct.

Then, the intersection with \mathcal{F} defines a surjection $\text{H-Conf}'(TN) \rightarrow \text{H-Conf}(\mathcal{F})$ which can easily be shown to be submersive. In fact, it is a proper map with compact fibres isomorphic to a product of projective spaces, showing that

$$\pi : \mathcal{S}^* \times_N \text{H-Conf}'(TN) \longrightarrow \mathcal{S} \times_N \text{H-Conf}(\mathcal{F}),$$

is a fibration. This allows us to define

$$\mathcal{A}^* := \pi^{-1}(\mathcal{A}) \subset \mathcal{S}^* \times_N \text{H-Conf}'(TN) \subset \mathcal{S}^* \times_N \text{H-Conf}(TN).$$

The openness of \mathcal{A}^* , as well as Properties (I), (II), and (III), follow from the analogous properties for \mathcal{A} . Concretely: Property (I) follows from $\pi : \mathcal{A}^* \rightarrow \mathcal{A}$ being a fibration. Openness and Property (III) are a consequence of the fact that $\text{H-Conf}'(TN) \subset \text{H-Conf}(TN)$ is open and its fibrewise complement is an algebraic subvariety (and thus of positive codimension). Property (II) follows from Lemma 5.2.5.

5.3.5 Removing processes

The most straightforward way of producing templates consists of iteratively removing those principal subspaces along which the relation is not ample.

Definition 5.3.8 *Let $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ be an over-relation. We set:*

$$\text{Avoid}^0(\mathcal{R}) := \mathcal{R} \times_M \text{H-Conf}(TM).$$

Inductively, we define $\text{Avoid}^{l+1}(\mathcal{R})$ to be the complement in $\text{Avoid}^l(\mathcal{R})$ of the closure of

$$\{(F, \Xi) \mid \text{for some } \tau \in \Xi, \text{ the component of } F \text{ in } \text{Avoid}^l(\mathcal{R})(\Xi)_{\tau, F} \text{ is not ample}\}.$$

Do note that, crucially, $\text{Avoid}^1(\mathcal{R})$ need not be a template. Indeed, upon removing elements from $\text{Avoid}^0(\mathcal{R})$, we may have lost ampleness along subspaces that were not problematic previously. This justifies the necessity of iterating the construction.

Definition 5.3.9 *Suppose that the process just described terminates, meaning that there is a step l_0 such that*

$$\text{Avoid}^l(\mathcal{R}) = \text{Avoid}^{l_0}(\mathcal{R}) \quad \text{for every } l \geq l_0.$$

*Then, $\text{Avoid}^\infty(\mathcal{R}) := \text{Avoid}^{l_0}(\mathcal{R})$ is the **standard pre-template** associated to $\iota_{\mathcal{R}}$.*

By construction:

Lemma 5.3.10 *Each $\text{Avoid}^l(\mathcal{R})$ is a pre-template. Additionally, $\text{Avoid}^\infty(\mathcal{R})$ satisfies Property (II) in the definition of a template.*

Proof. Openness follows from the fact that we are inductively removing closed sets. For Property (I) we reason inductively as well: By induction hypothesis, $\text{Avoid}^l(\mathcal{R})(\Xi)$ is contained in $\text{Avoid}^l(\mathcal{R})(\Xi')$ whenever $\Xi' \subset \Xi$. Suppose F is an element of both. Then, the analogous statement for the components of F in $\text{Avoid}^l(\mathcal{R})(\Xi)_{\tau, F}$ and $\text{Avoid}^l(\mathcal{R})(\Xi')_{\tau, F}$ is also true. In particular, if the latter is not ample, neither is the former. I.e. if (F, Ξ') is removed, so is (F, Ξ) , proving the claim.

The second statement follows by definition of the removal process terminating.

As we will observe in examples, $\text{Avoid}^\infty(\mathcal{R})$ need not satisfy Property (III); whether it does needs to be checked in each concrete application.

5.3.5.1 Thinning

In applications, the following more restrictive notion can also be useful.

Definition 5.3.11 *Let $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ be an over-relation. We write $\text{Thin}(\mathcal{R})$ for the complement in $\mathcal{R} \times_M \text{H-Conf}(TM)$ of the closure of*

$$\{(F, \Xi) \mid \text{for some hyperplane } \tau \in \Xi, \text{ the complement of } \mathcal{R}_{\tau, F} \text{ is not thin}\}.$$

We denote the l -fold iterate of this construction by $\text{Thin}^l(\mathcal{R})$.

Definition 5.3.12 *Assuming that there is a step l_0 in which this process stabilises, we say that $\text{Thin}^\infty(\mathcal{R}) := \text{Thin}^{l_0}(\mathcal{R})$ is the **thinning pre-template** of \mathcal{R} .*

Much like earlier:

Lemma 5.3.13 *$\text{Thin}^l(\mathcal{R})$ is a pre-template. Additionally, $\text{Thin}^\infty(\mathcal{R})$ satisfies Property (II) in the definition of a template.*

One can also conceive removing pieces from \mathcal{R} using schemes different from those presented in Definitions 5.3.9 and 5.3.12. In fact, this will be necessary for our main application Theorem 1.12; see Section 7.2.

5.3.6 Trivial pre-templates

It is unclear to the authors whether the standard avoidance/thinning processes always terminate regardless of what \mathcal{R} is. One could imagine a situation where we keep removing pieces from \mathcal{R} but never stabilise. Furthermore, even if they terminate, they may produce uninteresting results. This is not surprising, as many relations are simply not ample up to avoidance:

Lemma 5.3.14 *Fix a fibre bundle $X \rightarrow M$ and a differential relation $\mathcal{R} \subset J^r(X)$. Assume that:*

- *Each $\mathcal{R}_{\tau,z}$ is trivially ample or all its components are non-ample.*
- *Each fibre of π_{r-1}^r contains an element not in \mathcal{R} .*

Then, $\text{Avoid}^\infty(\mathcal{R})(\Xi) = \emptyset$, where Ξ is any tuple that includes a principal basis.

Do note that $\text{Avoid}^\infty(\mathcal{R})$ is only interesting for those Ξ that include a basis. Otherwise there are not enough directions to span the complete fibre of π_{r-1}^r .

Proof. Consider $\Xi \in \text{H-Conf}(T_p M)$ including a principal basis. We work over a fixed fiber of π_{r-1}^r lying over p . The Lemma follows as a consequence of the following inductive claim:

- Let z_k differ from some $z_0 \notin \mathcal{R}$ by a Ξ -principal path of length k . Then, it follows that $z_k \notin \text{Avoid}^k(\mathcal{R})(\Xi)$.

The base case $k = 0$ is definitionally true.

Consider the inductive step k . Given z_k , there is some z_{k-1} such that $\tau := z_k - z_{k-1}$ is principal and z_{k-1} differs from z_0 by a principal path ν of length $k - 1$.

Due to our assumptions on \mathcal{R} , either \mathcal{R}_{τ,z_0} is empty or its components are not ample. We can then take its complement \mathcal{R}_{τ,z_0}^c and note that the shift

$$\mathcal{R}_{\tau,z_0}^c + \nu \subset \text{Pr}_{\tau,z_{k-1}}$$

is, by inductive hypothesis, disjoint from $\text{Avoid}^{k-1}(\mathcal{R})(\Xi)_{\tau,z_{k-1}}$. It follows that $\text{Avoid}^{k-1}(\mathcal{R})(\Xi)_{\tau,z_{k-1}}$ is empty or its components are non-ample. Therefore, $\text{Avoid}^k(\mathcal{R})(\Xi)_{\tau,z_{k-1}}$ is empty. In particular, z_k is not in $\text{Avoid}^k(\mathcal{R})(\Xi)$.

A couple of concrete instances where Lemma 5.3.14 applies are the relation defining functions without critical points (Subsection 5.6.1) and the relation defining contact structures (Lemma 6.1.8).

Exactly the same reasoning shows:

Lemma 5.3.15 *Let $\mathcal{R} \subset J^r(X)$ be a differential relation such that:*

- *Every $\mathcal{R}_{\tau,z}$ is trivially ample or has a complement that is not thin.*
- *Each fibre of π_{r-1}^r contains an element not in \mathcal{R} .*

Then $\text{Thin}^\infty(\mathcal{R})(\Xi) = \emptyset$, where Ξ is any tuple that includes a principal basis.

5.4 Proof of the main Theorem

In this Section we tackle the proof of Theorem 1.8. We restate it now in a slightly more general form that applies to over-relations:

Theorem 5.4. *Fix a smooth bundle $X \rightarrow M$ and an open over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$. Suppose that $\iota_{\mathcal{R}}$ is ample up to avoidance. Then, the full C^0 -close h -principle applies to \mathcal{R} .*

The proof consists of two local-to-global steps. The starting point is our assumption that a template \mathcal{A} exists.

1. Given any principal cover \mathcal{C} , we use the pointwise data given by \mathcal{A} to produce an over-relation $\mathcal{A}(\mathcal{C}) \rightarrow J^r(X)$ globally on M . By construction, $\mathcal{A}(\mathcal{C})$ will be ample with respect to \mathcal{C} . This is the content of Proposition 5.4.2.
2. Given a formal solution $F : M \rightarrow \mathcal{R}$, we choose a cover \mathcal{C} of M such that F is still a formal solution of $\mathcal{A}(\mathcal{C})$. This follows from a jiggling-type argument that is explained in Proposition 5.4.3.

Both steps are rather discontinuous in nature. This is not surprising, since covers are discontinuous objects themselves. One of the consequences of this is that the over-relation $\mathcal{A}(\mathcal{C})$ may not be a fibration (even if \mathcal{R} and \mathcal{A} were).

This sketch of argument proves that all formal solutions F are short for \mathcal{R} . From this, and the parametric nature of avoidance templates, we deduce the full h -principle for \mathcal{R} . We put all these pieces together in Subsection 5.4.2.

5.4.1 Avoidance relations associated to principal covers

Fix a smooth bundle $X \rightarrow M$, an open over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$, an avoidance template \mathcal{A} , and a principal cover \mathcal{C} . Since the elements of \mathcal{C} are defined only locally, the cardinality of \mathcal{C} may change from point to point. This implies that we cannot regard \mathcal{C} as a smooth section $M \rightarrow \overline{\text{H-Conf}}(TM)$.

Nonetheless, for our purposes, the following discontinuous construction is enough. To each subset of codirections $C \subset \mathcal{C}$ (not necessarily a principal frame) we associate the closed set:

$$T_C := \bigcap_{\tau \in C} \overline{U_{\tau}} \subset M,$$

where U_{τ} is the support of the hyperplane field τ . Recall that each $\tau \in \mathcal{C}$ is defined as a germ along the closure $\overline{U_{\tau}}$ of its support. In particular, once we pick some order for the elements of C , we can think of C as a germ of smooth section

$$C|_{\mathcal{O}p(T_C)} : \mathcal{O}p(T_C) \rightarrow \overline{\text{H-Conf}}(TM).$$

In particular, the expression $\overline{\mathcal{A}}(C)|_{\mathcal{O}p(T_C)}$ denotes a well-defined subset of $\mathcal{R}|_{\mathcal{O}p(T_C)}$. Here $\overline{\mathcal{A}}$ is the lift of \mathcal{A} to $\mathcal{R} \times_M \overline{\text{H-Conf}}(TM)$.

Definition 5.4.1 *The **avoidance over-relation** associated to \mathcal{A} and \mathcal{C} is the set*

$$\mathcal{A}(\mathcal{C}) := \mathcal{R} \setminus \left(\bigcup_{C \subset \mathcal{C}} \overline{\mathcal{A}}(C)^c|_{T_C} \right),$$

where the superscript c denotes taking complement. As a subset of \mathcal{R} , the anchor of $\mathcal{A}(\mathcal{C})$ into $J^r(X)$ is $\iota_{\mathcal{R}}$.

5.4.2 Proof of Theorem 1.8

Before we get to the proof we introduce two key auxiliary results. The first one states that avoidance relations are open and ample:

Proposition 5.4.2 *Let \mathcal{A} be a template and \mathcal{C} be a principal cover. Then, the avoidance over-relation $\iota_{\mathcal{R}} : \mathcal{A}(\mathcal{C}) \rightarrow J^r(X)$ is an open over-relation ample along \mathcal{C} .*

The second one says that we can choose avoidance relations $\mathcal{A}(\mathcal{C})$ adapted to a given formal datum:

Proposition 5.4.3 *Fix an smooth bundle $X \rightarrow M$, an open over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$, an avoidance template \mathcal{A} , and a formal solution $F : M \rightarrow \mathcal{R}$. Then, there is a principal cover \mathcal{C} such that F takes values in $\mathcal{A}(\mathcal{C})$.*

This result is proven in Subsection 5.4.3.

Proof (Proof of Proposition 5.4.2). Recall the three pointwise Properties in the definition of a template (Definition 5.3.1).

Using the openness of \mathcal{A} and the closedness of each T_C , we see that $\mathcal{A}(\mathcal{C})$ is the complement in \mathcal{R} of a finite union of closed subsets $\overline{\mathcal{A}}(C)^c|_{T_C}$. As such, it is an open subset, and thus an open over-relation with respect to $\iota_{\mathcal{R}}$.

Ampleness can now be checked at each point $p \in M$ individually. We note that there is a maximal subset $C \subset \mathcal{C}$ such that T_C contains p . According to coherence Property (I) in Definition 5.3.1, $\overline{\mathcal{A}}(C(p))$ is the smallest among all sets $\overline{\mathcal{A}}(C'(p))$ as $C' \subset \mathcal{C}$ ranges over all the subcollections satisfying $p \in T_{C'}$. It follows that

$$\bigcup_{C' \subset \mathcal{C}} \overline{\mathcal{A}}(C'(p))^c = \overline{\mathcal{A}}(C(p))^c$$

and therefore we deduce $\mathcal{A}(\mathcal{C})(p) = \overline{\mathcal{A}}(C(p))$. The ampleness of the former follows then from the ampleness of the latter, which is given by Property (II).

Note that we have not made use of Property (III). It only plays a role in the proof of Proposition 5.4.3.

Proof (of Theorems 1.8 and 5.4 assuming Proposition 5.4.3).

We want to be able to homotope any given compact family of formal solutions $(F_k)_{k \in K} : M \rightarrow \mathcal{R}$ to a family of genuine solutions. We regard the family as a formal solution $F : M \times K \rightarrow \mathcal{R} \times K$ of the parametric lift, as in Subsection 5.1.4.1.

Since \mathcal{R} is ample up to avoidance we can apply Lemmas 5.3.6 and 5.3.7 to deduce that its parametric lift $\mathcal{R} \times K$ is also ample up to avoidance. It follows that each of the foliated relations $(\pi_{r'}^r \circ \iota_{\mathcal{R} \times K})_{r'=1, \dots, r}$ admits an avoidance template $\mathcal{A}_{r'}$.

We apply Proposition 5.4.3 to F , $\pi_1^r \circ \iota_{\mathcal{R} \times K}$, and \mathcal{A}_1 to deduce that there is a principal cover \mathcal{C}_1 of $M \times K$ such that $\mathcal{A}_1(\mathcal{C}_1)$ is ample along principal directions and F is a formal solution. In particular, F is short for $\pi_1^r \circ \iota_{\mathcal{R} \times K}$.

We then apply convex integration along a principal cover (Theorem 5.2). It follows that F is homotopic to a formal solution

$$G_1 : M \times K \rightarrow \mathcal{A}_1(\mathcal{C}_1) \subset \mathcal{R} \times K$$

that is holonomic up to first order. Applying this reasoning inductively on r' we produce a holonomic solution $G : M \times K \rightarrow \mathcal{R} \times K$ homotopic to F . The section G is equivalent to a family of holonomic solutions $(G_k)_{k \in K} : M \rightarrow \mathcal{R}$ homotopic to $(F_k)_{k \in K}$. This concludes the non-relative proof.

For the relative case we observe that, according to Theorem 5.2, the homotopy connecting F and G can be assumed to be constant along any closed set in which F was already holonomic. Since we are working in the foliated setting, this proves the parametric nature of the h -principle both in parameter (K) and domain (M) .

Corollary 5.4.4 *If \mathcal{R} is ample up to avoidance, it is ample in the sense of integrable convex hull extensions.*

Proof. Fix some arbitrary formal solution $F : M \times K \rightarrow \mathcal{R} \times K$, holonomic of order r' . During the proof of Theorem 1.8 we have shown that F is a formal solution of the avoidance relation $\mathcal{A}_{r'+1}(\mathcal{C}_{r'+1})$, which is ample along principal frames. It follows that all formal solutions are integrably short, so Definition 5.2.15 applies to \mathcal{R} .

5.4.3 Jiggling for principal covers

In this Subsection we prove Proposition 5.4.3, completing the proof of Theorems 1.8 and 5.4. Our goal is to find a principal cover \mathcal{C} compatible with a given avoidance template \mathcal{A} and a formal solution F . We construct \mathcal{C} using a *jiggling* argument. Namely, we start with an (arbitrary) principal cover \mathcal{C}' which we then subdivide repeatedly. When the subdivision is fine enough, we tilt/jiggle the corresponding principal frames in order to obtain the claimed \mathcal{C} .

This argument is (strongly) reminiscent of the classic version of jiggling due to W. Thurston [101]. For completeness, we recall it in Subsection 5.4.4; its contents are not really needed for our arguments and can be skipped. Our goal with this is to highlight the similarities between the two schemes. Despite of the many parallels, it is unclear to the authors whether there is some natural generalisation subsuming both results.

Remark 5.4.5 *We think of both jiggling arguments (both Thurston's and ours) as **h -principles without homotopical assumptions**.*

Namely, being transverse to a given distribution ξ is a differential relation for submanifolds of (N, ξ) . It may not be possible, in general, to find solutions of this relation. However, by dropping the smoothness assumption on the submanifold (allowing it to have instead triangulation-like singularities), Thurston produces solutions. Similarly, given a formal solution $F : M \rightarrow \mathcal{R}$, we can define a first order differential relation for tuples of functions $\{f_i : M \rightarrow \mathbb{R}\}$ by requiring $(F, \{df_i\}) \in \mathcal{A}$. By allowing the functions to be defined only locally (as coordinate codirections of charts), we are effectively introducing discontinuities; this is the flexibility we need to find a suitable \mathcal{C} .

In both cases, the main point is that, due to the presence of discontinuities, there is no formal data associated to the objects we consider. h -Principles without homotopical assumptions play now a central role in Symplectic and Contact Topology through the arborealisation programme [97, 8, 6, 7].

Recall the setup of Proposition 5.4.3: We are given a manifold M , a bundle $X \rightarrow M$, an over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$, an avoidance template \mathcal{A} , and a formal solution $F : M \rightarrow \mathcal{R}$. We want to find a principal cover \mathcal{C} such that $\mathcal{A}(\mathcal{C})$ is ample along \mathcal{C} and F is still a formal solution of $\mathcal{A}(\mathcal{C})$.

We will assume that M is compact. If not, the upcoming argument can be adapted to use an exhaustion by compacts.

5.4.3.1 Picking an atlas

We pick an arbitrary atlas \mathcal{U} of M . We require \mathcal{U} to use closed, cubical charts, i.e. each $(U, \phi) \in \mathcal{U}$ has image $[-1, 1]^n \subset \mathbb{R}^n$. Due to compactness, we may assume that \mathcal{U} is finite. We will still write ϕ to mean an arbitrary but fixed extension of ϕ to an open neighbourhood of U . We pick $\phi^{-1}(0) \in U$ as a marked point for each $(U, \phi) \in \mathcal{U}$.

To each ordered pair $((U, \phi_U), (V, \phi_V))$ in $\mathcal{U} \times \mathcal{U}$ we associate the transition function $\phi_{UV} := \phi_V \circ \phi_U^{-1}$. Its domain and codomain are the images of $U \cap V$.

5.4.3.2 Choosing principal frames

Given $(U, \phi) \in \mathcal{U}$, we pick a principal frame Ξ_U with support in U . We write e for the cardinality of this frame, which is the dimension of the fibres of π_{r-1}^r . We require $\phi_* \Xi_U$ to be invariant with respect to the translations in \mathbb{R}^n . Such an invariant principal frame is in correspondence with a principal basis at the marked point. During our arguments we think of the two interexchangeably.

The collection of all principal frames Ξ_U , as we range over the different $(U, \psi) \in \mathcal{U}$, defines a principal cover \mathcal{C}' of M .

5.4.3.3 Subdivision

Fix some real number $C > 1$. Let c be a positive integer to be fixed later on during the proof.

We subdivide $[-1, 1]^n$ into $(2c)^n$ cubes of side $1/c$, homothetic to the original. Given $(U, \phi) \in \mathcal{U}$, we apply this subdivision to U using ϕ . This yields a new collection of cubical charts, which we denote by $\mathcal{U}(c)$. A cube $V \in \mathcal{U}(c)$ is said to be the *child* of a *parent* cube $U \in \mathcal{U}$ if it is obtained from U by subdivision. Two children of the same parent are *siblings*.

Each child V inherits the parent chart ϕ , mapping now to a small cube of side $1/c$ contained in $[-1, 1]^n$. The marked point of V is the preimage by ϕ of the center of its image. The transition function between two given cubes in $\mathcal{U}(c)$ is inherited from the parents. In particular, if two cubes are siblings, the transition function between them is the identity (restricted to their overlap).

$\mathcal{U}(c)$ need not be a cover, since siblings overlap along sets with empty interior. To obtain a cover $\mathcal{V}(c)$, we dilate each cube in $\mathcal{U}(c)$, with respect to its center, by C . If c is sufficiently large, dilating by C makes sense even for children close to the boundary of the parent. This is why we extended the charts in \mathcal{U} to slightly bigger opens. After C -dilation, each child chart $(V, \phi) \in \mathcal{V}(c)$ has for image a cube of side C/c . The domain of the transition functions is changed accordingly. See Figure 5.4.

Lastly, we attach to each cube in $\mathcal{V}(c)$ a principal frame. It is simply the restriction of the principal frame of the parent. The collection of all these principal frames, as we range over $\mathcal{V}(c)$, is a principal cover that we call $\mathcal{C}'(c)$.

We now prove a number of quantitative properties for $\mathcal{V}(c)$, as we take c to infinity. We fix a fibrewise metric on $\mathbb{P}T^*M$. This defines a fibrewise metric in $\overline{\text{H-Conf}}(TM)$, since its fibres are simply products of projective spaces.

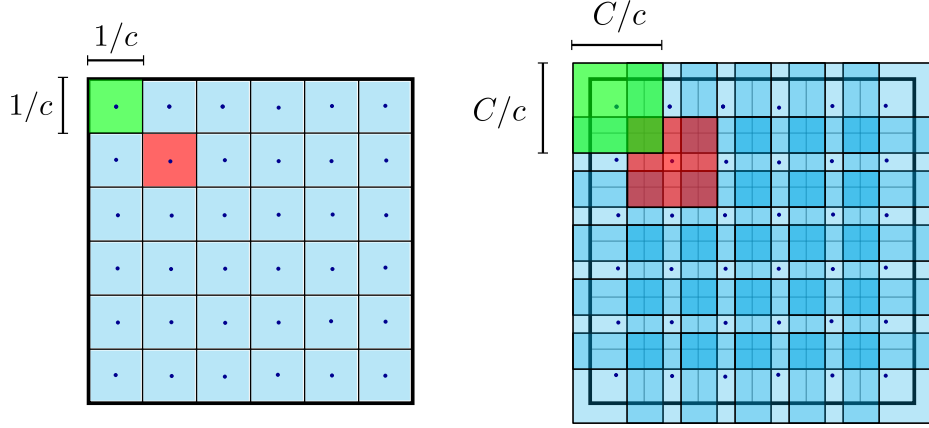


Figure 5.4: The big cube on the left represents the image $[-1, 1]^n$ of a chart $(U, \phi) \in \mathcal{U}$, before the C -dilation is introduced. We show how it is subdivided into smaller cubes, as well as the (image by ϕ of the) marked point of each smaller cube. Two of the children are marked in green and red. On the right, we depict the C -scaling of each child cube. The dilated green and red cubes, which originally met at a single point, now have an intersection with non-empty interior.

5.4.3.4 A bound on the number of overlapping cubes

Given $U \in \mathcal{V}(c)$ we write $\mathcal{V}_U(c)$ for the collection of cubes in $\mathcal{V}(c)$ that intersect U non-trivially.

Lemma 5.4.6 *There is an upper bound d_1 , independent of c and U , for the cardinality of $\mathcal{V}_U(c)$.*

Proof. This is trivially true when we restrict ourselves to sibling cubes. For unrelated cubes we reason as follows: Write $P \in \mathcal{U}$ for the parent of U and consider the image $\phi_{PP'}(U) \subset P'$ in some other cube $P' \in \mathcal{U}$, where $\phi_{PP'}$ is the transition function between P and P' . Since $\phi_{PP'}$ is C^1 -bounded by compactness, the diameter of $\phi_{PP'}(U)$ behaves as $O(1/c)$. The children of P' form a lattice spaced $1/c$. It follows that $\phi_{PP'}(U)$ can only intersect an amount $O(1)$ of them. The claim is complete since \mathcal{U} has finite cardinality.

5.4.3.5 Colouring

Lemma 5.4.7 *There is an integer d_2 , independent of c , such that we can partition $\mathcal{V}(c)$ into d_2 colours $\{\mathcal{V}(c)^{(i)}\}_{i=1, \dots, d_2}$ with the following property: If $U, V \in \mathcal{V}(c)$ belong to different colours, then they have no common neighbours (i.e. elements $W \in \mathcal{V}(c)$ overlapping non-trivially with both).*

Proof. This property is clear if we restrict to children of a fixed parent in \mathcal{U} , since children are spaced uniformly as $O(1/c)$ and have size $O(1/c)$. Then, by finiteness of \mathcal{U} , the claim follows if we use different sets of colours for each parent in \mathcal{U} .

5.4.3.6 Trivialising the configuration bundles

Given $(U, \phi) \in \mathcal{V}(c)$, we look at the bundle of hyperplanes $\mathbb{P}T^*U$. Consider the marked point $u \in U$ and the corresponding fibre $\mathbb{P}T_u^*U$. Using the parallel transport provided by the translations in the image of ϕ , we trivialise:

$$\mathbb{P}T^*U \cong U \times \mathbb{P}T_u^*U.$$

We denote the resulting projection by

$$\pi_U : \mathbb{P}T^*U \rightarrow \mathbb{P}T_u^*U.$$

Similarly, we trivialise the bundle $\overline{\text{H-Conf}}(TU)$ as $U \times \overline{\text{H-Conf}}(T_uU)$. This produces again a projection

$$\overline{\text{H-Conf}}(TU) \rightarrow \overline{\text{H-Conf}}(T_uU).$$

We abuse notation and also call π_U ; it should be apparent from context which of the two we mean.

Due to the compactness of M , the charts, projective space, and $\overline{\text{H-Conf}}$, we have that:

Lemma 5.4.8 *Fix a positive integer j_0 . Then, there is a constant $A > 1$, independent of c and $(U, \phi) \in \mathcal{V}(c)$, such that the following holds:*

*Fix a point $x \in U$. Identify $\mathbb{P}T_u^*U$ with $\mathbb{P}T_x^*U$ using π_U . Then, the fibrewise metrics at x and u bound each other from above up to a factor of A .*

The same statement holds, for every $j \leq j_0$, for the metrics in $\overline{\text{H-Conf}}_j(T_uU)$ and $\overline{\text{H-Conf}}_j(T_xU)$.

5.4.3.7 The diameter of a hyperplane field

Given $(U, \phi) \in \mathcal{V}(c)$, we look at the principal directions coming from neighbouring cubes. We want to show that these form a set whose diameter goes to zero as $c \rightarrow \infty$. We formalise this as follows, using the notation from the previous item.

Lemma 5.4.9 *Fix a second cube $(V, \psi) \in \mathcal{V}(c)$. Fix an integrable hyperplane field $\tau : V \rightarrow \mathbb{P}T^*V$, invariant under the translations in (V, ψ) . Consider the composition $\pi_U \circ \tau : U \cap V \rightarrow \mathbb{P}T_u^*M$.*

The diameter of $\text{image}(\pi_U \circ \tau)$ behaves like $O(1/c)$ and the constants involved do not depend on (U, ϕ) , (V, ψ) , or τ .

Proof. First note that the claim is automatic if U and V are siblings. Indeed, τ is then translation-invariant for the parent and thus for U , so $\pi_U \circ \tau$ is constant. Otherwise, write P for the parent of U and R for the parent of V . Let ϕ_{RP} be the transition function between the two; it restricts to the transition function between V and U .

The Taylor remainder theorem states that

$$d_{u+h}\phi_{RP} = d_u\phi_{RP} + O(h)$$

and the remainder is controlled by the second derivatives of ϕ_{RP} , which are bounded independently of c , U , V and τ . Since the diameter of U is $O(1/c)$, we have that

$$d_x\phi_{RP}(\pi_U \circ \tau) = d_u\phi_{RP}(\pi_U \circ \tau) + O(1/c) \quad \text{for all } x \in U \cap V,$$

proving the claim.

5.4.3.8 The diameter of a principal cover

We now look at covers instead of individual hyperplane fields. Fix $(U, \phi) \in \mathcal{V}(c)$ and consider all the neighbouring cubes. For each cube $(V, \psi) \in \mathcal{V}_U(c)$, suppose a principal frame ξ_V is given (not necessarily the one in $\mathcal{C}'(c)$ we fixed earlier). We may assume that ξ_V is defined over the whole of U simply by temporarily dilating V (a factor of 2 is sufficient).

The cardinality of $\mathcal{V}_U(c)$ is at most d_2 and the cardinality of each ξ_V is exactly e . By concatenating all the principal frames $(\xi_V)_{V \in \mathcal{V}_U(c)}$, we can regard them as a section

$$s_U : U \rightarrow \overline{\text{H-Conf}}_j(TU), \quad \text{for some } j \leq d_2.e.$$

Using π_U , we see s_U as a map $U \rightarrow \overline{\text{H-Conf}}_j(T_u U)$.

Lemma 5.4.10 *The diameter of $\text{image}(s_U)$ behaves like $O(1/c)$. The constants involved do not depend on (U, ϕ) nor on $(\xi_V)_{V \in \mathcal{V}_U(c)}$.*

Proof. The metric on $\overline{\text{H-Conf}}(TM)$ is just the product metric inherited from the metric in $\mathbb{P}T^*M$. Then the claim follows from Lemma 5.4.9 due to the finiteness of j .

5.4.3.9 Density bounds on the avoidance template

The discussion up to this point referred only to coverings and principal frames. The avoidance template \mathcal{A} and the formal solution $F : M \rightarrow \mathcal{R}$ enter the proof now. We will make use of openness and Property (III) in the definition of a template. Our goal is to provide a quantitative estimate regarding the size of the balls contained in $\overline{\mathcal{A}}(F)$ that one can find on a given ε -ball in $\overline{\text{H-Conf}}(TM)$.

Fix $(U, \phi) \in \mathcal{V}(c)$ with marked point u . Using the projection π_U we can associate to $\overline{\mathcal{A}}(F)$ the singularity:

$$\Sigma_U := \pi_U(\overline{\mathcal{A}}(F)^c \cap \overline{\text{H-Conf}}(TU)) \subset \overline{\text{H-Conf}}(T_u U),$$

where the superscript c denotes taking complement.

Lemma 5.4.11 *Let $\varepsilon > 0$ be given. Then, there exists $\delta > 0$ such that, for any sufficiently large c and any $j \leq d_2.e$, the following property holds:*

Fix $(U, \phi) \in \mathcal{V}(c)$ with marked point u . Each ε -ball in $\overline{\text{H-Conf}}_j(T_u U)$ contains a δ -ball disjoint from Σ_U .

Proof. Consider $x \in M$ arbitrary but fixed. We claim that there is $\delta_x > 0$ such that every ε -ball in $\overline{\text{H-Conf}}_j(T_x M)$ contains a δ_x -ball fully contained in $\overline{\mathcal{A}}(F)$; see Figure 5.5. Indeed: suppose B is a $\varepsilon/2$ -ball in $\overline{\text{H-Conf}}_j(T_x M)$. Since $\overline{\mathcal{A}}(F)$ is fibrewise dense, there exists $\Xi \in \overline{\mathcal{A}}(F) \cap \overline{\text{H-Conf}}_j(T_x M)$. By openness of $\overline{\mathcal{A}}(F)$ there is a δ_Ξ -ball $D \subset \overline{\mathcal{A}}(F) \cap \overline{\text{H-Conf}}_j(T_x M)$ centered at Ξ and contained in B . We now use the compactness of $\overline{\text{H-Conf}}_j(T_x M)$ to extract a finite cover by $\varepsilon/2$ -balls $\{B_i^x\}$. There are corresponding δ_x -balls $\{D_i^x\}$ contained in $\overline{\mathcal{A}}(F) \cap B_i^x$. Any ε -ball in $\overline{\text{H-Conf}}_j(T_x M)$ contains one of the B_i^x and thus the corresponding D_i^x , as claimed.

Before we address the statement, let us introduce some notation. Fix $(P, \phi_P) \in \mathcal{U}$ and $p \in P$, not necessarily the marked point. We use ϕ_P to trivialise $\overline{\text{H-Conf}}_j(TP) = P \times \overline{\text{H-Conf}}_j(T_p P)$, allowing us to speak of the $\alpha \times \beta$ -polydiscs given by such a trivialisation. These are products of an α -disc along P (measured by the euclidean metric of the chart) and a β -disc along the fibre $\overline{\text{H-Conf}}_j(T_p P)$ (measured by the fibrewise metric at p). By definition, a $\alpha \times \beta$ -polydisc is obtained from a fibrewise β -disc by parallel transport to the nearby fibres. We can now use the openness of $\overline{\mathcal{A}}(F)$ to thicken the collection $\{D_i^x\}$ to a family of $\rho_p \times \delta_p$ -polydiscs contained in $\overline{\mathcal{A}}(F)$, for some $\rho_p > 0$. We abuse notation and still denote these thickenings by $\{D_i^p\}$.

Using the finiteness of \mathcal{U} and the compactness of each $(P, \phi_P) \in \mathcal{U}$ we can then find constants $\rho, \delta > 0$, and lattices of points $\{p_l^P \in P\}_{l \in L, P \in \mathcal{U}}$ spaced as $\rho/2$, such that $\rho < \rho_{p_l^P}/2$ and $\delta < \delta_{p_l^P}/A$, for all $l \in L$ and $P \in \mathcal{U}$. Here A is the dilation factor given in Lemma 5.4.8.

If c is sufficiently large, any child $(U, \phi) \in \mathcal{V}(c)$ of a given $(P, \phi_P) \in \mathcal{U}$ is contained in the ρ -disc centered at some $p_l^P \in P$. In particular, any $\rho \times \varepsilon$ -polydisc centered at $\overline{\text{H-Conf}}_j(T_u M)$, with u the




Figure 5.5: The vertical line depicts the fibre of $\overline{\text{H-Conf}}_j(TM)$ over a given $x \in M$. The branching set running more or less horizontally is the complement of $\overline{\mathcal{A}}(F)$. Given a ball B of radius $\varepsilon/2$ in the fibre, we find $\Xi \in \overline{\text{H-Conf}}_j(T_x M)$ and a ball $D \subset B$ around it such that D is fully contained in $\overline{\mathcal{A}}(F)$. This argument uses only that $\overline{\mathcal{A}}(F)$ is dense.

marked point of U , is contained in a $2\rho \times \varepsilon$ -polydisc $B_i^{p_i^P}$ centered at $\overline{\text{H-Conf}}_j(T_{p_i^P} M)$. Then, upon projecting with π_U , the corresponding $2\rho \times \delta$ -polydisc $D_i^{p_i^P}$ provides the claimed δ -ball disjoint from Σ_U . See Figure 5.6.

5.4.3.10 Jiggling

The proof concludes by applying jiggling. We fix an arbitrary constant $\varepsilon_0 > 0$. By making it smaller we will be proving that the jiggling can be assumed to be as small as we want. We then define a sequence of constants (as many as colours):

$$\varepsilon_0 > \varepsilon_1 > \cdots > \varepsilon_{d_2} > 0$$

by iteratively applying Lemma 5.4.11. Namely, $4\varepsilon_{i+1}$ should be the “ δ ” corresponding to ε_i . Furthermore, we impose for each ε_i to be much bigger than the subsequent ones. Concretely, the following inequality should hold:

$$\varepsilon_i > 2.A \sum_{j>i} \varepsilon_j. \quad (5.1)$$

The successive applications of Lemma 5.4.11 provide us then with a lower bound for c .

We start with the first colour $\mathcal{V}(c)^{(1)}$, working simultaneously with all its elements. Let $(U, \phi) \in \mathcal{V}(c)^{(1)}$ with marked point u . Consider all the neighbouring $(V, \psi) \in \mathcal{V}_U(c)$, each with a corresponding principal frame Ξ_V . Together, these define a map $s_U : U \rightarrow \overline{\text{H-Conf}}(T_u M)$, as in Subsection 5.4.3.8. According to Lemma 5.4.10, the image of s_U has diameter $O(1/c)$. In particular, if c is sufficiently large (of magnitude $O(1/\varepsilon_1)$), we can assume that this diameter is smaller than ε_1 . Using Lemma 5.4.11 we can perturb each Ξ_V to a nearby frame Ξ'_V such that the corresponding map $s'_U : U \rightarrow \overline{\text{H-Conf}}(T_p M)$ satisfies:

- The C^0 -distance between s'_U and s_U is bounded above by $\varepsilon_0/2$.
- The ε_1 -neighbourhood of image(s'_U) is contained in $\overline{\mathcal{A}}(F)$.




Figure 5.6: The chart $U \in \mathcal{V}(c)$ is fully contained in the ρ -ball centered at some p_l^P . The black vertical line depicts the fibre $\overline{\text{H-Conf}}_j(T_u M)$ over its marked point $u \in U$. The branching set in light blue is the complement of $\overline{\mathcal{A}}(F)$. Given a $\rho \times \varepsilon$ -polydisc (with orange border) centered somewhere in $\overline{\text{H-Conf}}_j(T_u M)$, there is a $2\rho \times \varepsilon$ -polydisc (in blue) that contains it. In turn, the latter contains a $2\rho \times \delta$ -polydisc (in purple) which is disjoint from $\overline{\mathcal{A}}(F)^c$. This exhibits a δ -ball disjoint from Σ_U .

I.e. we have jiggled all the frames in the vicinity of U , producing new frames whose distance to the complement of $\overline{\mathcal{A}}(F)$ is controlled.

We do the same inductively on the number of colours. At step i we look at all the $(U, \phi) \in \mathcal{V}^{(i)}(c)$ at once. Using Lemma 5.4.11 we perturb the neighbouring frames $s_U : U \rightarrow \overline{\text{H-Conf}}(T_u U)$ to a nearby section s'_U satisfying:

- The C^0 -distance between s'_U and s_U is bounded above by ε_i .
- The ε_{i+1} -neighbourhood of $\text{image}(s'_U)$ is contained in $\overline{\mathcal{A}}(F)$.

This is possible as long as c is large enough; concretely, of magnitude $O(1/\varepsilon_{d_2})$. The second item gives us a lower bound for the distance to $\overline{\mathcal{A}}(F)$ which, in light of Equation 5.1, is not destroyed in later steps thanks to the first item. After d_2 steps, the proofs of Proposition 5.4.3, Theorem 1.8, and Theorem 5.4 are complete. \square

5.4.4 Detour: Thurston's jiggling

We invite the reader to compare the upcoming discussion to the proof of Proposition 5.4.3. Let us stress once more that this Subsection is only included for the sake of pointing out the parallels between the two.

The classic jiggling procedure was introduced by W. Thurston in [101]. It allows the user to produce a triangulation whose simplices are transverse to a given distribution ξ . We work with a manifold N of dimension n . We consider triangulations \mathcal{T} all whose i -simplices $\Delta \in \mathcal{T}^{(i)}$ are endowed with a parametrisation identifying them with the standard simplex in \mathbb{R}^i . Furthermore, they come with a parametrised neighbourhood germ, which is then identified with a neighbourhood of the standard simplex in $\mathbb{R}^i \subset \mathbb{R}^n$.

We say that Δ is in **general position** with respect to ξ if the quotient map

$$\text{mod}(\xi_p) : \Delta \subset \mathbb{R}^i \subset \mathbb{R}^n \rightarrow \mathbb{R}^{n-k}$$

that quotients by ξ_p has for image a subset diffeomorphic to a simplex of dimension $\min(i, n - k)$. Here we are restricting ξ to the aforementioned coordinates around Δ . Do note that the general position condition is then an strengthening of the condition that Δ is transverse to ξ . A triangulation is in general position if all its simplices are.

Proposition 5.4.12 *Let (N, ξ) be an n -manifold endowed with a distribution. Then, there exists a triangulation \mathcal{T} in **general position** with respect to ξ .*

Furthermore, if a constant $\varepsilon > 0$ and a triangulation \mathcal{T}' of N are given, we may assume that \mathcal{T} is obtained from \mathcal{T}' by applying finitely many cubical subdivisions and then perturbing the vertices of the resulting triangulation by an amount no larger than ε .

Proof. We break the argument into steps, much like we did for Proposition 5.4.3. First, we fix a locally finite atlas \mathcal{U} of N .

Upon subdividing c times, we may assume that \mathcal{T}' is subordinated to \mathcal{U} . Cubical subdivision guarantees that the radius of the simplices goes to zero as $c \rightarrow \infty$, while the cardinality of the star of each simplex is bounded independently of c .

A key property, which follows from the previous paragraph, is that there is a number d_2 , independent of c , such that we can colour the set of vertices into d_2 colours so that no two vertices of the same colour are contained in the same simplex.

We straighten out the simplices so that every $\Delta \in \mathcal{T}'$ is linear with respect to some chart in \mathcal{U} .

We tilt/jiggle the positions of the vertices of \mathcal{T}' in order to change how all simplices are embedded and thus achieve transversality. We do this inductively colour by colour. It follows that, at each inductive step, the vertices we tilt do not interfere with one another.

For a given vertex p , we consider those simplices that contain p and whose other vertices are from previous colours. When we tilt p , the good choices are those that make said simplices transverse. As $c \rightarrow \infty$, the measure of the subset of bad choices goes to zero (since the radii of all simplices go to zero).

Given $\varepsilon_0 > 0$ and any $\varepsilon_1 > 0$ sufficiently small, we can always take c large enough so that, in each ε_0 -ball of choices, there is a ε_1 -ball of good choices. This reasoning can be repeated with ε_1 and some $\varepsilon_2 > 0$. Repeating it d_2 times yields a sequence of positive constants $(\varepsilon_i)_{0, \dots, d_2}$. Our tilt at step $i+1$ can then be taken to be ε_i/d_2 small and contained in a ε_{i+1} -ball of good choices as long as c is sufficiently large. It follows that the transversality achieved in a given step is not destroyed in subsequent ones. The proof concludes after d_2 steps.

5.5 An example: exact, linearly-independent differential forms

In this section we introduce a differential relation \mathcal{R} for which the thinning process (Subsection 5.3.5.1) terminates producing an avoidance template. We let M be a 3-dimensional manifold and

we set $X := T^*M \times T^*M$. The relation $\mathcal{R} \subset J^1(X)$ is of first order and Diff-invariant. Concretely, \mathcal{R} consists of pairs $(F_1, F_2) \in J^1(X)$, where each F_i is a first-order Taylor polynomial of 1-form, such that the 2-forms dF_1 and dF_2 are linearly independent.

In Subsection 5.5.1 we recall some notation, involving the exterior differential, that will also be helpful in later sections. In Subsection 5.5.2 we formalise our claim about the thinning of \mathcal{R} , which we then prove in Subsection 5.5.3.

In Subsection 5.5.4 we will observe that the present example is a bit artificial: a linear algebra lemma due to Gromov shows that \mathcal{R} was already ample in all directions (but not thin). Nonetheless, we choose to include it as an easy incarnation of the avoidance/thinning approach.

5.5.1 Intermezzo: The symbol of the exterior differential

Let N be a manifold. We focus on the bundle $\bigwedge^k T^*N$ of k -forms. The corresponding space of 1-jets $J^1(\bigwedge^k T^*N)$ is an affine fibration

$$\pi_0 : J^1(\bigwedge^k T^*N) \longrightarrow \bigwedge^k T^*N$$

whose model vector bundle is $\text{Hom}(TN, \bigwedge^k T^*N) \cong T^*N \otimes \bigwedge^k T^*N$.

It follows that, given a formal datum $F \in J^1(\bigwedge^k T^*N)$ and a codirection $\lambda \in T_p^*N$, both based at the same point $p \in N$, the principal space associated to F and λ can be expressed explicitly as:

$$\text{Pr}_{\lambda, F} := \{F + (0, \lambda \otimes \beta) \mid \beta \in \bigwedge^k T_p^*N\}.$$

We are interested in discussing differential relations defined in terms of the exterior differential d . We will abuse notation and still denote its **symbol** by:

$$d : J^1(\bigwedge^k T^*N) \longrightarrow \bigwedge^{k+1} T^*N.$$

Given $F \in J^1(\bigwedge^k T^*N)$, the symbol maps $F + (0, \sum_i \lambda_i \otimes \beta_i)$ to $dF + \sum_i \lambda_i \wedge \beta_i$. We also introduce the **extended symbol**:

$$\begin{aligned} \text{id} \oplus d : J^1(\bigwedge^k T^*N) &\longrightarrow \bigwedge^k T^*N \oplus \bigwedge^{k+1} T^*N, \\ F &\mapsto (\pi_0(F), dF). \end{aligned}$$

5.5.2 The statement

We now restate the setup of our example. We fix a 3-manifold M and we focus on the first-order differential relation $\mathcal{R} \subset J^1(T^*M \oplus T^*M)$ defined as:

$$\mathcal{R} := \{(F_1, F_2) \in J^1(T^*M) \oplus J^1(T^*M) \mid dF_1 \text{ and } dF_2 \text{ are linearly independent}\}.$$

The main result of this Section reads:

Proposition 5.5.1 *The thinning process for \mathcal{R} terminates in one step and the resulting thinning pre-template $\text{Thin}(\mathcal{R})$ is an avoidance template.*

In particular, \mathcal{R} is ample up to avoidance and thus, according to Theorem 1.8, the h -principle holds for \mathcal{R} .

The proof of Proposition 5.5.1 requires an analysis of the structure of \mathcal{R} along principal subspaces (Subsection 5.5.3.1). This will then allow us to describe $\text{Thin}(\mathcal{R})$ explicitly and deduce that it is an avoidance template (Subsection 5.5.3.3).

5.5.3 The proof

We introduce some auxiliary notation: Fix a point $p \in M$. Given a principal direction $\lambda \in T_p^*M$, we define the singularity

$$\Sigma(\lambda) := \{F = (F_1, F_2) \in J_p^1(T^*M \oplus T^*M) \mid \text{Both } dF_i \text{ are multiples of } \lambda\}.$$

The complement of $\Sigma(\lambda)$ in $J_p^1(T^*M \oplus T^*M)$ is a first-order differential constraint, defined only at the point p . Nonetheless, we can still talk about it being ample; in fact, we will prove that $\Sigma(\lambda)$ is a thin singularity (see Lemma 5.5.4 below).

5.5.3.1 The thinning step

The following criterion will allow us to compute $\text{Thin}(\mathcal{R})$:

Lemma 5.5.2 *Fix p , λ and F as above. The following two conditions are equivalent:*

- $\mathcal{R}_{\lambda,F} \subset \text{Pr}_{\lambda,F}$ has thin complement.
- $F \notin \Sigma(\lambda)$.

Proof. Recall that the principal space defined by λ and F is given by:

$$\text{Pr}_{\lambda,F} = \{(F_1, F_2) + ((0, \lambda \otimes \beta_1), (0, \lambda \otimes \beta_2)) \mid \beta_i \in T_p^*M\}$$

In particular, it is parametrised by the pairs $\beta_1 \times \beta_2 \in T_p^*M \times T_p^*M$. The restriction of the relation \mathcal{R} to $\text{Pr}_{\lambda,F}$ reads:

$$\mathcal{R}_{\lambda,F} = \{(F_1, F_2) + ((0, \lambda \otimes \beta_1), (0, \lambda \otimes \beta_2)) \mid \text{The forms } dF_i + \lambda \wedge \beta_i \text{ are linearly independent}\}.$$

The symbol of the exterior differential yields then a map

$$d : \text{Pr}_{\lambda,F} \longrightarrow \wedge^2 T_p^*M \times \wedge^2 T_p^*M$$

whose image $d\text{Pr}_{\lambda,F}$ is 4-dimensional. It is the product $L_1 \times L_2$ of the plane of 2-forms

$$L_1 := \{dF_1 + \lambda \wedge \beta_1 \mid \beta_1 \in T_p^*M\},$$

passing through dF_1 , and the plane

$$L_2 := \{dF_2 + \lambda \wedge \beta_2 \mid \beta_2 \in T_p^*M\}$$

passing through dF_2 . These two planes are parallel to the distinguished plane L consisting of those 2-forms proportional to λ .

There are two possible situations. The first possibility (Figure 5.7) is that both dF_i are contained in L (equivalently, $L = L_1 = L_2$; equivalently, $F \in \Sigma(\lambda)$). In this case, there are covectors ν_i such that $dF_i = \lambda \wedge \nu_i$. Then we can identify $d\text{Pr}_{\lambda,F}$ with the 2-by-2 matrices:

$$\begin{aligned} \tau : \quad d\text{Pr}_{\lambda,F} &\longrightarrow \text{Ker}(\lambda)^* \times \text{Ker}(\lambda)^* \cong \mathcal{M}_{2 \times 2} \\ (dF_i + \lambda \wedge \beta_i)_{i=1,2} &\mapsto ((\nu_i + \beta_i)|_{\text{Ker}(\lambda)})_{i=1,2}. \end{aligned}$$

The subset $\mathcal{R}_{\lambda,F}$ maps precisely to the linearly independent pairs $((\nu_i + \beta_i)|_{\text{Ker}(\lambda)})_{i=1,2}$. That is, $\tau \circ d$ is an affine submersion of $\mathcal{R}_{\lambda,F}$ onto $\text{GL}(2)$. We then conclude that the complement of $\mathcal{R}_{\lambda,F}$ is




Figure 5.7: The first possibility in the proof of Lemma 5.5.2: Both dF_i are contained in L .

not thin, because the complement of $\mathrm{GL}(2)$ (the zero set of the determinant, which has codimension 1) is not. We will see in Subsection 5.5.4 that $\mathrm{GL}(2)$ is nonetheless ample, and so is $\mathcal{R}_{\lambda,F}$.

The other possibility is that one of the dF_i is not proportional to λ (equivalently, one of the L_i is different from L ; equivalently, $F \notin \Sigma(\lambda)$). Suppose, without loss of generality, that it is dF_1 . Then, an element $dF_1 + \lambda \wedge \beta_1$ is colinear with a single element in L_2 (if $L_2 \neq L$) or with none (if $L_2 = L$). In the first case (Figure 5.8), the complement of $\mathcal{R}_{\lambda,F}$ is of codimension-2. In the second case (Figure 5.9), the complement of $\mathcal{R}_{\lambda,F}$ is the set $\{dF_2 = 0\}$, which is also of codimension-2. This completes the claim.




Figure 5.8: The second possibility in the proof of Lemma 5.5.2: None of the dF_i are contained in L .




Figure 5.9: The third possibility in the proof of Lemma 5.5.2: One of the dF_i is contained in L but the other one is not.

Recall the avoidance template notation from Subsection 5.3.2. We remark that the thinning process was defined as a pointwise process and Lemma 5.5.2 indeed applies to each point $p \in M$ individually. An immediate consequence is:

Corollary 5.5.3 *The following statements are equivalent:*

- $(F, \Xi) \in \mathcal{R} \times_M \text{H-Conf}(TM)$ belongs to $\text{Thin}(\mathcal{R})$.
- F belongs to $\mathcal{R} \setminus (\bigcup_{\lambda \in \Xi} \Sigma(\lambda))$.

5.5.3.2 A second step is not necessary

It turns out that the set we have removed from \mathcal{R} during the first step is itself thin:

Lemma 5.5.4 *Fix a point $p \in M$ and a codirection $\nu \in T_p^*M$. Then, $\Sigma(\nu) \subset J_p^1(T^*M \oplus T^*M)$ is a thin singularity.*

Proof. Fix a principal direction $\lambda \in T_p^*M$ and a formal datum $F \in J_p^1(T^*M \oplus T^*M) \setminus \Sigma(\nu)$. Recall the affine spaces $\text{Pr}_{\lambda, F}$, $d\text{Pr}_{\lambda, F}$, L and L_i from Lemma 5.5.2. We additionally consider the subspace L' of 2-forms spanned by ν . The symbol d maps the singularity $\Sigma(\nu) \cap \text{Pr}_{\lambda, F} \subset \text{Pr}_{\lambda, F}$ to the product $(L' \cap L_1) \times (L' \cap L_2) \subset d\text{Pr}_{\lambda, F}$. There are three options:

- λ and ν are proportional and $L_1 = L_2 = L = L'$. Then $\text{Pr}_{\lambda, F} = \Sigma(\nu)$, which contradicts the fact that F was in the complement.
- λ and ν are proportional, so $L = L'$, but at least one L_i , say L_1 , is distinct from L' . Then $L' \cap L_1$ is empty and so is $\Sigma(\nu) \cap \text{Pr}_{\lambda, F}$.
- λ and ν are not proportional. Then both intersections $L' \cap L_i$ are lines and thus $\Sigma(\nu)$ has codimension 2 in $\text{Pr}_{\lambda, F}$.

Only the two last possibilities can happen and the claim follows.

5.5.3.3 Completing the proof

The previous Lemmas allow us to conclude:

Proof (Proof of Proposition 5.5.1). We want to show that $\text{Thin}(\mathcal{R})$ is an thinning template. To do this, we must prove three properties. First, that for all $\Xi \in \text{H-Conf}(TM)$, the subset $\text{Thin}(\mathcal{R})(\Xi)$ has thin complement along each direction in Ξ . Second, that for all $F \in \mathcal{R}$, the subset $\text{Thin}(\mathcal{R})(F)$ is fibrewise dense in $\text{H-Conf}(TM)$. Lastly, that $\text{Thin}(\mathcal{R})$ is open.

According to Corollary 5.5.3, we have the following explicit description:

$$\text{Thin}(\mathcal{R})(\Xi) = \mathcal{R} \setminus \left(\bigcup_{\lambda \in \Xi} \Sigma(\lambda) \right).$$

Fix $\nu \in \Xi$. According to Lemma 5.5.2, each subspace $\mathcal{R}_{\nu,F}$ is contained in $\Sigma(\nu)$ (if $F \in \Sigma(\nu)$) or is disjoint from it and has thin complement. In the former case, $\text{Thin}(\mathcal{R})(\Xi)_{\nu,F}$ is empty. In the latter case, we use Lemma 5.5.4 to note that all other singularities $\Sigma(\lambda) \cap \mathcal{R}_{\nu,F}$ are thin singularities. This proves the first property.

For the second property, we fix $F \in \mathcal{R}$. A configuration Ξ is in the complement of $\text{Thin}(\mathcal{R})(F)$ if and only if there is $\lambda \in \Xi$ such that $F \in \Sigma(\lambda)$. Equivalently, if and only if $dF_1 \wedge \lambda = dF_2 \wedge \lambda = 0$. This is a non-trivial algebraic condition for Ξ , which proves the claim.

Lastly, note that $dF_1 \wedge \lambda = dF_2 \wedge \lambda = 0$ is an algebraic equality both on dF_i and λ . Therefore, its complement is open (and, in fact, open in each variable upon freezing the other one), proving the third property.

5.5.4 Linear algebra

As we stated in the proof of Lemma 5.5.2, the subspace of non-degenerate matrices is an ample subset of the space of all matrices. This shows that the relation \mathcal{R} studied in this Section was, in fact, ample. For completeness we provide a proof:

Proposition 5.5.5 *The subspace $\text{GL}(n)$ of non-singular matrices is an ample subset of the space of $(n \times n)$ -matrices $\mathcal{M}_{n \times n}$ if and only if $n \geq 2$.*

Proof. The two components $\text{GL}^+(n)$ and $\text{GL}^-(n)$ are connected. We have to show that each is individually ample.

First note that every $n \times n$ matrix can be expressed as the convex combination of two non-singular matrices, since

$$M = \frac{1}{2}(2M - 2\lambda \cdot Id) + \frac{1}{2}\lambda \cdot 2Id$$

and the right hand-side is the sum of two non-singular matrices for any choice of $\lambda \notin \text{Spec}(M) \setminus 0$. Therefore, it is enough to show that any matrix $M \in \text{GL}^+(n)$ can be expressed as a convex combination of matrices in $\text{GL}^-(n)$ (and viceversa). This readily follows by writing $M = (v_1, v_2, v_3, \dots, v_n)$ (expressed in column vectors) as $M = \frac{1}{2}M_1 + \frac{1}{2}M_2$, where

$$M_1 = (-v_1, 3v_2, v_3, \dots, v_n)$$

$$M_2 = (3v_1, -v_2, v_3, \dots, v_n).$$

Note that M_1 and M_2 do not belong to the same connected component of $\text{GL}(n)$ as M and, thus, the claim follows.

Proof (Alternate proof). Observe that $\mathcal{M}_{n \times n}$ is convexly spanned by those matrices with a single non-zero entry.

Then: Given a matrix M and a sufficiently large constant C , it holds that M is in the interior of the convex hull of the matrices $e_{i,j}^{\pm}$ whose single non-zero entry is (i,j) with value $\pm C$. The matrix $e_{i,j}^{\pm}$ has zero determinant so it may be perturbed to yield a matrix $\tilde{e}_{i,j}^{\pm}$ with positive (resp. negative) determinant. In doing so, the convex hull is perturbed as well. However, M will remain in the interior if the perturbations are small enough, by continuity. This proves the ampleness of $\text{GL}^+(n)$ (resp. $\text{GL}^-(n)$).

We conclude:

Corollary 5.5.6 *The subspace of non-singular matrices in $\mathcal{M}_{n \times m}$ is ample unless $n = m = 1$.*

Proof. We may assume $n \neq m$; otherwise the space of singular matrices has codimension greater than 1. Then the claim follows from the previous result.

Corollary 5.5.7 *The relation \mathcal{R} treated in this Section is ample in all principal directions and therefore abides by the h -principle.*

5.6 Relations involving functions

In this Section we restrict our attention to bundles $X \rightarrow M$ with 1-dimensional fibres. The typical example is the trivial \mathbb{R} bundle over M , whose sections are functions. Our claim is that convex integration (even if it includes avoidance) is ill-suited to address relations $\mathcal{R} \subset J^r(X)$.

In Subsection 5.6.1 we prove that the relation defining functions without critical points fails to be ample up to avoidance. In Subsection 5.6.2 we generalise this to arbitrary differential relations $\mathcal{R} \subset J^r(X)$ that satisfy a mild non-triviality condition.

Remark 5.6.1 *Let us compare these claims with [58, Remark 1.3.4]. Gromov states that a generic codimension-2 singularity Σ is thin. This is true if the fibres of X have dimension at least 2. Indeed, under genericity assumptions, Σ intersects every principal subspace in a (maybe empty) codimension-2 subset.*

However, if the fibres of X have dimension 1, Σ does not intersect all the principal subspaces, only a subset of codimension-1. Further, these intersections are necessarily of codimension-1. As such, thinness and ampleness (even up to avoidance) fail.

Observe further that a differential relation given by a concrete geometric problem is, by definition, not generic. We claim that our methods can be used to generalise Gromov's statement to such non-generic situations. Suppose Σ has codimension 2 and the fibres of X have dimension at least 2. Even if Σ intersects some principal subspaces in non-thin sets, it intersects most of them transversely, by Sard's theorem. One can then apply our methods to analyse the problematic subspaces.

Remark 5.6.2 *We observe that the applicability realms of Vassiliev's h -principle [102] and convex integration are, in some sense, complementary. The former is most interesting when \mathcal{R} is the complement of a singularity of large codimension (at least $\dim(M) + 2$) and X has 1-dimensional fibres. The latter is effective in the presence of much larger singularities but requires the fibres of X to have dimension at least 2.*

5.6.1 A non-example: Functions without critical points

Let M be a manifold and we let X be the trivial \mathbb{R} -bundle over M . As a differential relation in $J^1(X)$ we take

$$\mathcal{R} = \{F \in J^1(X) \mid dF \neq 0\},$$

i.e. the 1-jets of functions whose differential is non-zero. With the standard identification $J^1(X) \cong T^*M \times \mathbb{R}$ we see that \mathcal{R} is the complement of the singularity $M \times \mathbb{R}$.

Fix now a codirection λ and a formal datum $F \in J^1(X)$, both based at the same point $p \in M$. The principal subspace associated to them is one-dimensional and explicitly given by:

$$\text{Pr}_{\lambda,F} = \{F + (c\lambda, 0) \in T^*M \times \mathbb{R} \mid c \in \mathbb{R}\}.$$

We readily see that there are two possible situations:

- dF is proportional to λ . Then the complement of $\mathcal{R}_{\lambda,F}$ is a point, which is not thin.
- dF is not proportional to λ . Then $\mathcal{R}_{\lambda,F} = \text{Pr}_{\lambda,F}$ so ampleness holds trivially.

This shows that Lemma 5.3.14 applies to \mathcal{R} , allowing us to conclude that the standard pre-template $\text{Avoid}^\infty(\mathcal{R})$ is empty for all configurations of codirections that form a generating set. The same applies to the thinning pre-template $\text{Thin}^\infty(\mathcal{R})$. This was to be expected since, due to Morse theory, there is no h -principle for functions without critical points.

5.6.2 The general case

It follows immediately from Lemma 5.3.14 that:

Proposition 5.6.3 *Let $X \rightarrow M$ be a bundle with 1-dimensional fibres. Let $\mathcal{R} \subset J^r(X)$ be the complement of a singularity Σ that intersects non-trivially each fibre of $J^r(X) \rightarrow J^{r-1}(X)$.*

Let l_0 be the dimension of the fibres of $J^r(X) \rightarrow J^{r-1}(X)$. Then, $\text{Avoid}^{l_0}(\mathcal{R})$ is empty for all configurations of codirections of cardinality $l \geq l_0$ that form a principal basis.

In particular, \mathcal{R} is not ample nor ample up to avoidance.

Chapter 6

h -Principle for step-2 distributions

6.1 h -Principle for step-2 distributions

In this Chapter we will prove the h -principle for step-2 distributions (Theorem 1.10) and its corollary about the classification of (3,5) and (3,6) distributions (Theorem 1.11). The proof can be found in Subsection 6.1.4. We emphasise that the contents of this Section do not need avoidance and simply rely on classic convex integration.

Before we get to the results, and in order to set notation, we recall some of the basic theory of tangent distributions in Subsection 6.1.1. This will allow us, in Subsection 6.1.2, to translate our statements about distributions to statements about their annihilating forms.

In Subsection 6.1.5, for completeness, we look at convex integration in the setting of (even-)contact structures, following the work of McDuff [83].

6.1.1 The dual picture

Fix an n -dimensional ambient manifold M and a rank- k distribution ξ . Recall that the Lie flag and the nilpotentisation were already introduced in the introductory Subsection 1.5.

In practice, whenever we impose (natural) differential conditions on distributions, these can be read either using a frame of vector fields or a frame of the annihilator. In this thesis, the distributions we look at have greater rank than corank, so it is more convenient to use the annihilator viewpoint:

$$\xi^\perp := \{\alpha \in T^*M \mid \alpha(v) = 0, \forall v \in \xi\}.$$

In Subsection 1.5.0.3 we discussed the curvature of ξ . Its first entry is a 2-form with entries in ξ and image in TM/ξ . Upon passing to the wedge product, it is equivalent to a bundle morphism

$$\Omega^\xi : \xi \wedge \xi \longrightarrow TM/\xi.$$

We can then dualise it using the Cartan formula, yielding a bundle map

$$\begin{aligned} d^\xi : \xi^\perp &\longrightarrow \wedge^2 \xi^* \\ \alpha &\longmapsto -\alpha \circ \Omega^\xi = d\alpha|_\xi. \end{aligned}$$

We note that $d\alpha|_\xi$ only depends on the pointwise value of $\alpha \in \xi^\perp$. At the risk of overloading our notation, we will say that $d^\xi \alpha$ is the **curvature** associated to $\alpha \in \xi^\perp$.

In light of the Cartan formula, the kernel of d^ξ is ξ_2^\perp . In particular:

Lemma 6.1.1 ξ is of step-2 if and only if any of the following equivalent conditions holds:

- d_ξ is a monomorphism.
- Ω^ξ is an epimorphism.

In particular, if ξ is of step-2, the ambient dimension is at most $\text{rank}(\xi) + \binom{\text{rank}(\xi)}{2}$. Conversely, under this assumption on the dimension, generic distributions are of step-2 at a generic point (see [9]).

6.1.2 Step-2 as a differential relation

We see rank- k distributions as sections of the grassmannian bundle $\text{Gr}_k(TM)$. Being bracket-generating in two steps is then a differential relation $\mathcal{R}^{\text{step2}} \subset J^1(\text{Gr}_k(TM))$ of first order. More concretely, we observe that any element $F \in J^1(\text{Gr}_k(TM))$ defines a k -plane j^0F , as well as an associated curvature

$$\Omega^F : \wedge^2 j^0F \longrightarrow TM/j^0F,$$

simply because the curvature depends only on first order derivatives. Then:

Definition 6.1.2 The differential relation $\mathcal{R}^{\text{step2}} \subset J^1(\text{Gr}_k(TM))$ consists of those $F \in J^1(\text{Gr}_k(TM))$ such that Ω^F is an epimorphism.

We will write $\text{Dist}_{(k,n)}^f(M)$ for the space of formal solutions of $\mathcal{R}^{\text{step2}}$. The subspace of holonomic ones is denoted by $\text{Dist}_{(k,n)}(M)$. Being bracket-generating in $l+1$ steps is similarly a differential relation of order l .

Definition 6.1.2 is not very practical and it is best to pass to a description in terms of forms. Namely, we consider the bundle of tuples $\oplus^{n-k} T^*M$. Over the open set $\overline{\oplus^{n-k} T^*M}$ consisting of linearly-independent tuples, we have a quotient map:

$$\pi : \overline{\oplus^{n-k} T^*M} \longrightarrow \text{Gr}_k(TM),$$

which induces a map

$$j^r \pi : J^r(\overline{\oplus^{n-k} T^*M}) \longrightarrow J^r(\text{Gr}_k(TM))$$

between jet spaces. Then:

Definition 6.1.3 A jet $\tilde{F} \in J^1(\oplus^{n-k} T^*M)$ is **formally bracket-generating of step-2** if the following conditions hold:

- $\tilde{F} \in J^1(\overline{\oplus^{n-k} T^*M})$ and therefore it projects to an element $F \in J^1(\text{Gr}_k(TM))$. Denote $\xi = j^0F$.
- The 2-forms $d\tilde{F}|_\xi$ are linearly independent.

The subset of such \tilde{F} will be denoted by $\mathcal{S}^{\text{step2}} \subset J^1(\oplus^{n-k} T^*M)$.

The second condition is, according to Lemma 6.1.1, indeed equivalent to Ω^F being an epimorphism. It follows that:

Lemma 6.1.4 $\mathcal{S}^{\text{step2}}$ is the preimage of $\mathcal{R}^{\text{step2}}$ under $j^1\pi$. In particular, $\mathcal{S}^{\text{step2}}$ fibres affinely over $\mathcal{R}^{\text{step2}}$.

6.1.3 Localisation to a ball

As explained before, a formal solution in $F \in \text{Dist}_{(k,n)}^f(M)$ defines a k -plane field $j^0 F$. It may very well happen that the annihilator $(j^0 F)^\perp$ is not trivial as a bundle. This would imply that F cannot be lifted to $\tilde{F} \in J^1(\oplus^{n-k} T^*M)$. In particular, there may be no global lift of F to $\mathcal{S}^{\text{step}2}$.

Nonetheless, the h -principle for $\mathcal{R}^{\text{step}2}$ reduces to the h -principle for $\mathcal{S}^{\text{step}2}$. We could directly invoke that convex integration is local (i.e. that it is performed chart by chart). This would certainly be enough for our purposes in this Section, which rely on ampleness along all directions. However, it seems more delicate for ampleness up to avoidance.

In order to set the stage for later Sections, we follow a different approach. The following standard trick gets the job done (even in the presence of parameters and relatively): The manifold M (or the product $M \times K$, in the presence of a parameter space K) can be triangulated and holonomic approximation [39, Chapter 3] can be applied along the codimension-1 skeleton \mathcal{T} . This homotopes the formal solution F (resp. K -family of formal solutions) to a new formal solution $G \in \text{Dist}_{(k,n)}^f$ (resp. K -family of formal solutions) that is holonomic along \mathcal{T} and C^0 -close to F everywhere. See [5] for the general theory behind this.

The outcome is that now we can restrict our attention to the top dimensional cells, which are contractible. Over each ball, the annihilator of $j^0 G$ is now trivial, and thus a lift to $\mathcal{S}^{\text{step}2}$ exists. We conclude that:

Lemma 6.1.5 *In order to prove the h -principle for $\mathcal{R}^{\text{step}2}$, it is sufficient to prove it for $\mathcal{S}^{\text{step}2}$.*

The h -principle for $\mathcal{S}^{\text{step}2}$ will follow from convex integration, as we prove next.

6.1.4 h -Principle for step-2 distributions

Proof (of Theorem 1.10). According to Lemma 6.1.5, it is sufficient to check that $\mathcal{S}^{\text{step}2}$ is ample along all codirections. Fix a formal solution $F = (F_i)_{i=1}^{n-k} \in \mathcal{S}^{\text{step}2}$ based at some point $p \in M$. By assumption, the $j^0 F = (j^0 F_i)_{i=1}^{n-k}$ are linearly independent and thus annihilate a k -plane $\xi \subset T_p M$. Furthermore, the $dF_i|_\xi$ are linearly independent.

Fix a codirection $\lambda \in T_p^* M$. The principal space associated to λ and F reads:

$$\text{Pr}_{\lambda, F} = \{(F_i + (0, \lambda \otimes \beta_i))_{i=1}^{n-k} \mid \beta_i \in T_p^* M\}.$$

The differential of any $\tilde{F} \in \text{Pr}_{\lambda, F}$ reads $(d\tilde{F}_i = dF_i + \lambda \wedge \beta_i)_{i=1}^{n-k}$. A tuple \tilde{F} belongs to $\mathcal{S}^{\text{step}2}$ if and only if the tuple of two-forms $d\tilde{F}|_\xi$ is linearly independent. Suppose $\lambda \in \xi^\perp$. Then we have that $d\tilde{F}|_\xi = dF|_\xi$ and therefore ampleness holds (because all \tilde{F} are formal solutions, since F was).

Otherwise, we suppose that λ represents a non-trivial element in ξ^* . Then, as far as $d\tilde{F}|_\xi$ is concerned, only the restriction of β_i to $\xi \cap \text{Ker}(\lambda)$ is important. The forms $(dF_i|_{\text{Ker}(\lambda) \cap \xi})_{i=1}^{n-k}$ span a subspace, say, of dimension l . Up to a change of basis, we may then assume that

$$dF_i|_{\text{Ker}(\lambda) \cap \xi} = 0, \quad \text{for all } i = l+1, \dots, n-k.$$

Equivalently:

$$dF_i|_\xi = \lambda \wedge \nu_i, \quad \text{for all } i = l+1, \dots, n-k, \text{ for some } \nu_i \in T^* M.$$

Then, the tuple $d\tilde{F} \in \text{Pr}_{\lambda, F}$ is in $\mathcal{S}^{\text{step}2}$ if and only if the forms $\{(\lambda \wedge (\beta_i + \nu_i))|_\xi\}_{i=l+1}^{n-k}$ are linearly independent. Equivalently, if and only if the forms $\{(\beta_i + \nu_i)|_{\text{Ker}(\lambda) \cap \xi}\}_{i=l+1}^{n-k}$ are linearly independent.

This means that the ampleness of $\mathcal{S}^{\text{step2}}$ along $\text{Pr}_{\lambda,F}$ is equivalent to the ampleness of the subspace A of rank- $(n - k - l)$ matrices within $\mathcal{M}_{(n-k-l) \times (k-1)}$.

If $n - k - l > k - 1$ (equivalently $l < n - 2k + 1$), the subspace A is empty contradicting the fact that F was a formal solution. Otherwise, $n - k - l \leq k - 1$ holds and A is just the subspace of non-degenerate matrices. Due to our assumptions (step 2 and dimension at least 4), we have that $k \geq 3$ and thus we deduce $k - 1 \geq 2$. It follows that A is ample according to Lemma 5.5.6, concluding the proof.

Theorem 1.10 proves that step-2 distributions are flexible as long as we do not impose any further non-degeneracy constraints. Nonetheless, according to Theorem 1.11, there are two cases, (3,5) and (3,6), where the h -principle for maximally non-involutive distributions readily follows from the Theorem.

The following claims follow by inspection of the proof:

Remark 6.1.6 *When the dimension of $\wedge^2 \xi$ is exactly the corank of ξ (i.e. for distributions of maximal growth $(k, k + \binom{k}{2})$, including (3,6) and (4,10)), the singularity associated to $\mathcal{S}^{\text{step2}}$ has codimension-1 and is thus not thin. Indeed, in this situation, the singularity is either trivial or equivalent to the complement of GL inside of all square matrices.*

Further, the singularity may have codimension-1 even when the dimension of $\wedge^2 \xi$ is greater than the corank of ξ . This can be observed in the case of (3,5) distributions, where the singularity is thin along most principal subspaces, but not all (corresponding to the case of $l = 0$ in the previous proof). This is the same phenomenon observed for the relation studied in Section 5.5.

6.1.5 The contact and even-contact cases

As an appetiser for our study of maximally non-involutive distributions in Sections 7.1 and 7.2, we now revisit the contact and even-contact cases. We will show that the former fails to be ample (as was to be expected, since contact structures do not abide by the h -principle [11]), whereas the latter is thin along all principal directions (as proven by D. McDuff in [83]).

6.1.5.1 The Pfaffian

Let M be an n -dimensional manifold. Once again, it is convenient, since we are dealing with hyperplane fields, to work with forms. The bundle of interest for us will be the cotangent bundle T^*M . In order to measure non-involutivity, we introduce the *Pfaffian* map:

$$\begin{aligned} \Gamma: J^1(T^*M) &\longrightarrow T^*M \oplus \wedge^2 T^*M \longrightarrow \wedge^{2\lfloor \frac{n-1}{2} \rfloor + 1} T^*M \\ F &\longmapsto (j^0 F, dF) \longmapsto j^0 F \wedge (dF)^{\lfloor \frac{n-1}{2} \rfloor}, \end{aligned}$$

where the first arrow is the extended symbol of the exterior differential. The Pfaffian measures whether the formal curvature $dF|_{\text{Ker}(j^0 F)}$ has maximal rank.

Definition 6.1.7 *The (even)-contact differential relation for 1-forms is defined as:*

$$\mathcal{R}^{\text{cont}} := J^1(T^*M \setminus 0) \setminus \Gamma^{-1}(0) \subset J^1(T^*M).$$

Once again we emphasise that one can pass from distributions to forms locally (see Subsection 6.1.3), so the h -principle for (even-)contact structures is equivalent to the h -principle for $\mathcal{R}^{\text{cont}}$. We study its ampleness next.

6.1.5.2 Checking ampleness

Fix a coordinate direction λ and a formal solution $F \in \mathcal{R}^{\text{cont}}$, both based at a point $p \in M$. The two together define the principal subspace

$$\text{Pr}_{\lambda,F} := \{F + (0, \lambda \otimes \beta) \mid \beta \in T_p^*M\}$$

which maps, using the extended symbol of d , to:

$$d\text{Pr}_{\lambda,F} := \{(j^0F, dF + \lambda \wedge \beta) \mid \beta \in T_p^*M\} \subset T_p^*M \oplus \wedge^2 T_p^*M.$$

We write $\xi = \text{Ker}(j^0F)$.

A point $\tilde{F} = F + (0, \lambda \otimes \beta) \in \text{Pr}_{\lambda,F}$ is formally (even-)contact if and only if:

$$\Gamma(\tilde{F}) = n j^0F \wedge (dF)^{\lfloor \frac{n-1}{2} \rfloor - 1} \wedge (dF + \lambda \wedge \beta) \neq 0.$$

Since F was a formal solution, there are four possible situations:

1. λ is proportional to j^0F .
2. λ is not proportional to j^0F and n is odd. Then dF has a 1-dimensional kernel L when restricted to $\xi \cap \text{Ker}(\lambda)$.
- 3a. λ is not proportional to j^0F , n is even, and $\text{Ker}(\lambda)$ contains the 1-dimensional kernel of $dF|_{\xi}$. Then dF has a 2-dimensional kernel L' when restricted to $\xi \cap \text{Ker}(\lambda)$.
- 3b. λ is not proportional to j^0F , n is even, and $\text{Ker}(\lambda)$ is transverse to the 1-dimensional kernel of $dF|_{\xi}$. Then dF is non-degenerate when restricted to $\xi \cap \text{Ker}(\lambda)$.

Situation (1) means that $\Gamma(\tilde{F}) = \Gamma(F) \neq 0$, so ampleness holds trivially.

Situation (2) corresponds to the (non-trivial) contact case. Then, $\mathcal{R}^{\text{cont}} \cap \text{Pr}_{\lambda,F}$ corresponds to those choices of β that evaluate non-zero on L . The complement is then a hyperplane, proving that:

Lemma 6.1.8 *The differential relation describing contact structures is not ample. In fact, along any given principal subspace, $\mathcal{R}^{\text{cont}} \cap \text{Pr}_{\lambda,F}$ is either trivially ample or not ample.*

Situations (3a) and (3b) correspond to the (non-trivial) even-contact case. In (3a), $\mathcal{R}_{\lambda,F}^{\text{cont}}$ corresponds to those choices of β that evaluate non-zero on L . Its complement is codimension-2. In (3b), ampleness holds trivially since $\mathcal{R}_{\lambda,F}^{\text{cont}} = \text{Pr}_{\lambda,F}$. We have shown:

Lemma 6.1.9 *The differential relation describing even-contact structures has thin complement.*

The h -principle for even-contact structures follows then from classic convex integration Theorem [5.1](#).

Chapter 7

h –Principle for hyperbolic $(4, 6)$ -distributions.

7.1 Maximal non-involutivity

Fix two positive integers $k < n$. We want to define maximal non-involutivity for step-2 distributions of rank k in dimension n . Just like in Subsection 6.1.5, we will use the Pfaffian map to capture this in an algebraic manner (Subsection 7.1.1). In Subsection 7.1.2 we will particularise the discussion to the rank-4 case. Much of what we explain in this Section we learnt from the book by R. Montgomery [82].

7.1.1 The Pfaffian and degenerate differential forms

Let ξ be a distribution of rank k in an n -dimensional manifold M . We measure the non-involutivity of ξ using the curvature Ω^ξ . Recall that Ω^ξ is a 2-form with entries in ξ and values in TM/ξ . Dually, we think of it as a ξ^\perp -family of 2-forms. As in Subsection 6.1.5, we can take the highest potentially-non-trivial power of the curvatures using the map:

$$\begin{aligned} p : \wedge^2 \xi^* &\longrightarrow \wedge^{2\lfloor \frac{k}{2} \rfloor} \xi \\ \omega &\longmapsto \omega^{\lfloor \frac{k}{2} \rfloor}. \end{aligned}$$

That is, a curvature gets mapped to a top-form when the rank k is even, and to a codimension-1 form when k is odd.

Definition 7.1.1 A 2-form $\omega \in \wedge^2 \xi^*$ is **degenerate** if $p(\omega) = 0$.

The map p is algebraic and of degree $\lfloor \frac{k}{2} \rfloor$. Its zero level set \mathcal{C} consists of the degenerate 2-forms. It has codimension 1 inside $\wedge^2 \xi^*$ if k is even and codimension k if k is odd. Then we define:

Definition 7.1.2 The composition $p \circ d^\xi$ is called the **Pfaffian**:

$$\begin{aligned} \text{Pf} : \xi^\perp &\longrightarrow \wedge^2 \xi^* \longrightarrow \wedge^{2\lfloor \frac{k}{2} \rfloor} \xi^* \\ \alpha &\longmapsto d\alpha|_\xi \longmapsto (d\alpha|_\xi)^{\lfloor \frac{k}{2} \rfloor}. \end{aligned}$$

We are interested in the Diff-invariant non-degeneracy condition:

Definition 7.1.3 Let k be even. A step-2 distribution ξ is **maximally non-involutive** if Pf intersects \mathcal{C} transversely.

Maximal non-involutivity may be described similarly (but differently) for k odd. This is unnecessary for the purposes of this thesis. We now focus on rank-4 distributions.

7.1.2 4-distributions

Since $k = 4$ is even, the target of p is a 1-dimensional line bundle; namely, the determinant of ξ . Assuming orientability of ξ , which we can do locally, we can fix a volume form on ξ to trivialise it. This allows us to see p as a quadratic form on $\wedge^2 \xi^*$ and study its signature. The signature does not depend on the choice of volume form:

Lemma 7.1.4 *The real quadratic form p has signature $(3, 3)$.*

Proof. Take a local frame $\xi^* = \langle \beta_1, \beta_2, \beta_3, \beta_4 \rangle$ compatible with the chosen orientation. Define now the space of self-dual forms $\wedge^+(\xi^*)$ and the space of anti self-dual forms $\wedge^-(\xi^*)$ as follows:

$$\wedge^+(\xi^*) = \langle a_1 = \beta_1 \wedge \beta_2 + \beta_3 \wedge \beta_4, a_2 = \beta_1 \wedge \beta_3 + \beta_4 \wedge \beta_2, a_3 = \beta_1 \wedge \beta_4 + \beta_2 \wedge \beta_3 \rangle \subset \wedge^2 \xi^*.$$

$$\wedge^-(\xi^*) = \langle b_1 = \beta_1 \wedge \beta_2 - \beta_3 \wedge \beta_4, b_2 = \beta_1 \wedge \beta_3 - \beta_4 \wedge \beta_2, b_3 = \beta_1 \wedge \beta_4 - \beta_2 \wedge \beta_3 \rangle \subset \wedge^2 \xi^*.$$

A straightforward computation shows that the matrix associated to the bilinear form p with respect to the basis $\langle a_1, a_2, a_3, b_1, b_2, b_3 \rangle$ consists of an upper-left $\text{Id}_{3 \times 3}$ identity block and another $-\text{Id}_{3 \times 3}$ in the right-down corner. I.e. p diagonalises with the claimed signature.

Since ξ is bracket-generating, the exterior differential d^ξ maps ξ^\perp injectively into $\wedge^2 \xi^*$. We can then talk about the signature of p restricted to the image. This is equivalent to:

Definition 7.1.5 *The signature of a distribution ξ is the signature of the quadratic form $\text{Pf} : \xi^\perp \rightarrow \wedge^2(\xi^*) \cong \mathbb{R}$.*

Two remarks are in order. First: the signature is well-defined only once a volume form on ξ has been chosen. Otherwise, we cannot distinguish the signatures (i, j, k) and (j, i, k) . Furthermore, the signature of ξ may vary from point to point.

We focus on the case of $(4, 6)$ -distributions:

Definition 7.1.6 *A bracket-generating 4-distribution in a 6-dimensional manifold M is said to be **maximally non-involutive** if, at all points, any of the following equivalent conditions holds:*

- *The map d^ξ is transverse to the locus of degenerate 2-forms.*
- *The Pfaffian is transverse to zero.*
- *The Pfaffian is non-degenerate as a quadratic form.*

Furthermore, we can distinguish two different types of $(4, 6)$ -distributions:

Definition 7.1.7 *A maximally non-involutive $(4, 6)$ -distribution is*

- i. **elliptic** or **fat** if the signature is definite.*
- ii. **hyperbolic** if the signature is mixed.*

Since p has signature $(3, 3)$ and ξ^\perp has dimension 2, there are, up to changing the orientation, four possible signatures for ξ : $(0, 0, 2)$, $(1, 0, 1)$, $(1, 1, 0)$ and $(2, 0, 0)$. Only the last two cases are maximally non-involutive. They correspond, respectively, to the hyperbolic and elliptic cases.

7.1.3 Formal maximally non-involutive 4-distributions

As in Section 6.1, it is more convenient not to work with the distribution itself but with its annihilating forms. We define:

Definition 7.1.8 *A formal datum $F = (F_i)_{i=1,2} \in J^1(T^*M \oplus T^*M)$ is said to be **formally elliptic** if:*

- *the 1-forms $j^0 F_i$ are linearly independent and thus span a 4-plane ξ .*
- *the 2-forms $dF_i|_\xi$ are linearly independent and span a 2-plane of definite signature.*

*The formal datum F is **formally hyperbolic** if, instead:*

- *the 1-forms $j^0 F_i$ are linearly independent and thus span a 4-plane ξ .*
- *the 2-forms $dF_i|_\xi$ are linearly independent and span a 2-plane of mixed signature.*

A first jet of distribution is elliptic/hyperbolic if and only if any pair of forms $(F_i)_{i=1,2}$ representing it is elliptic/hyperbolic.

We write \mathcal{R}^{ell} for the differential relation defining elliptic $(4,6)$ distributions. Similarly, \mathcal{R}^{hyp} denotes the differential relation consisting of formal hyperbolic $(4,6)$ -distributions. Their counterparts at the level of forms are denoted by \mathcal{S}^{ell} and \mathcal{S}^{hyp} , respectively. As in Subsection 6.1.3, one can use holonomic approximation to reduce to the case of a ball, proving that:

Lemma 7.1.9 *The full C^0 -close h -principle for \mathcal{R}^{hyp} reduces to the full C^0 -close h -principle for \mathcal{S}^{hyp} .*

The upcoming final Section of the chapter deals with the construction of an avoidance template for \mathcal{S}^{hyp} .

7.2 h -Principle for hyperbolic $(4,6)$ -distributions

In this Section we tackle the proof of Theorem 1.12. According to Lemma 7.1.9, the h -principle for \mathcal{R}^{hyp} will follow from the h -principle for formally hyperbolic pairs of 1-forms $\mathcal{S}^{\text{hyp}} \subset J^1(T^*M \oplus T^*M)$. Applying Theorem 1.8 we see that we just need to construct an avoidance template for \mathcal{S}^{hyp} .

7.2.1 First avoidance step

As we advanced in Subsection 1.6.1, \mathcal{S}^{hyp} (and thus \mathcal{R}^{hyp}) does not intersect all principal subspaces in ample sets, so avoidance will act non-trivially. Before we provide a precise statement, we need to introduce some notation.

7.2.1.1 The singularity associated to non-ameness

We define $\Sigma^{(1)} \subset \mathcal{S}^{\text{hyp}} \times_M T^*M$ as the subspace of pairs (F, λ) such that

$$j^0 F_1 \wedge j^0 F_2 \wedge \lambda \wedge dF_1 \quad \text{and} \quad j^0 F_1 \wedge j^0 F_2 \wedge \lambda \wedge dF_2 \quad \text{are linearly dependent.} \quad (7.1)$$

We write $\Sigma^{(1)}(\lambda) \subset \mathcal{S}^{\text{hyp}}$ for the subset of those F such that $(F, \lambda) \in \Sigma^{(1)}$. Similarly, $\Sigma^{(1)}(F) \subset T^*M$ denotes those λ such that $(F, \lambda) \in \Sigma^{(1)}$.

Lemma 7.2.1 *The following statements hold:*

- $\Sigma^{(1)}$ is a closed subset of $\mathcal{S}^{\text{hyp}} \times_M T^*M$.
- All the fibres $(\Sigma_p^{(1)})_{p \in M}$ are isomorphic algebraic subvarieties.
- Fix $F \in \mathcal{S}^{\text{hyp}}$ lying over $p \in M$. The subspace $\Sigma^{(1)}(F)$ has positive codimension in T_p^*M .

Proof. \mathcal{S}^{hyp} is Diff-invariant, and therefore all its fibres are isomorphic to one another. Furthermore, the expressions in Equation 7.1 are algebraic on their entries and linear dependence is itself a closed algebraic condition. These statements prove the first two claims.

For the last claim, we observe that fixing F yields still an algebraic equality for λ that is non-trivial as long as $j^0 F_1 \wedge j^0 F_2 \wedge dF_1$ and $j^0 F_1 \wedge j^0 F_2 \wedge dF_2$ are linearly independent. This is indeed the case if $F \in \mathcal{S}^{\text{hyp}}$.

Write $\xi \subset T_p M$ for the 4-plane given as the kernel of $j^0 F$. We note that the following are equivalent:

- $(F, \lambda) \in \Sigma^{(1)}$.
- The 3-forms $(\lambda \wedge dF)|_\xi$ are linearly-dependent.
- $\lambda|_\xi$ is zero or the 2-forms $dF|_{\xi \cap \text{Ker}(\lambda)}$ are linearly-dependent.

7.2.1.2 Main statement

We claim that $\Sigma^{(1)}$ is precisely the set to be removed in order to carry out the first avoidance step.

Proposition 7.2.2 *Let $F \in \mathcal{S}^{\text{hyp}}$ and $\lambda \in T_p^*M$, both based at the same point $p \in M$. Write $\xi \subset T_p M$ for the 4-plane defined by $j^0 F$. Then:*

- $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is non-trivially ample if and only if $(F, \lambda) \notin \Sigma^{(1)}$.
- $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is trivially ample if and only if $\lambda|_\xi = 0$. In particular, $(F, \lambda) \in \Sigma^{(1)}$.
- $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is not ample otherwise. I.e. if $(F, \lambda) \in \Sigma^{(1)}$ but $\lambda|_\xi \neq 0$.

Let us provide some geometric insight before we get into the proof. An element \tilde{F} in the principal subspace $\text{Pr}_{\lambda, F}$ maps under the exterior differential to a pair $(d\tilde{F}_i = dF_i + \lambda \wedge \beta_i)_{i=1,2}$, where the β_i range over T_p^*M .

According to Definition 7.1.8, the pair $dF|_\xi$ spans a plane of 2-forms L . The restriction $p|_L$ is a bilinear form of mixed signature, due to hyperbolicity. Now consider the subspace

$$K := \{(\lambda \wedge \beta)|_\xi \mid \beta \in T_p^*M\} \subset \wedge^2 \xi^*.$$

By definition, given any other element $\tilde{F} \in \text{Pr}_{\lambda, F}$, the pair $d\tilde{F}|_\xi$ is obtained from $dF|_\xi$ by shifting each form $dF_i|_\xi$ along K . As such, when K and L are transverse, the pair $d\tilde{F}|_\xi$ will span a plane \tilde{L} that is a graph over L in the direction of K . If transversality fails, it may very well happen that the pair $d\tilde{F}|_\xi$ is linearly dependent; however, we still think of its span \tilde{L} as a degenerate plane.

We further note that being a graph over L in the direction of K is an intrinsic characterisation of the planes \tilde{L} associated to elements $\tilde{F} \in \text{Pr}_{\lambda,F}$. I.e. the set of all such planes does not depend on the concrete basis dF of L . We furthermore note that $\tilde{F} \in \mathcal{S}_{\lambda,F}$ if and only if $p|_{\tilde{L}}$ is non-degenerate of mixed signature. This means that all relevant properties of \tilde{F} can be read from \tilde{L} . We conclude that we are allowed to choose a convenient basis of L in order to simplify our computations.

Proof (Proof of Proposition 7.2.2). Consider $\tilde{F} \in \text{Pr}_{\lambda,F}$ and restrict p to its (possibly degenerate) span \tilde{L} . This restriction can be represented by the 2-by-2 matrix

$$\begin{pmatrix} g_{11} & g_{12} \\ g_{12} & g_{22} \end{pmatrix} \quad (7.2)$$

whose coefficients read:

$$\begin{aligned} g_{11}(\beta_1, \beta_2) &= (dF_1 + \lambda \wedge \beta_1)^2 = dF_1^2 + 2\lambda \wedge \beta_1 \wedge dF_1, \\ g_{22}(\beta_1, \beta_2) &= (dF_2 + \lambda \wedge \beta_2)^2 = dF_2^2 + 2\lambda \wedge \beta_2 \wedge dF_2, \\ g_{12}(\beta_1, \beta_2) &= (dF_1 + \lambda \wedge \beta_1) \wedge (dF_2 + \lambda \wedge \beta_2) = dF_1 \wedge dF_2 + \lambda \wedge (\beta_2 \wedge dF_1 + \beta_1 \wedge dF_2). \end{aligned}$$

Each of these expressions can be identified with a real function by fixing a volume form on ξ . We fix such a volume; all our upcoming computations do not depend on this auxiliary choice.

We have effectively defined an affine map that takes values in the space of symmetric 2-by-2 matrices, which we think of as \mathbb{R}^3 :

$$\begin{aligned} \Psi: \xi^* \oplus \xi^* &\longrightarrow \mathbb{R}^3 \\ (\beta_1, \beta_2) &\longmapsto (g_{11}(\beta_1, \beta_2), g_{22}(\beta_1, \beta_2), g_{12}(\beta_1, \beta_2)) \end{aligned}$$

It is now convenient to introduce the determinant, which we see as a quadratic form in the space of symmetric 2-by-2 matrices:

$$\begin{aligned} \det: \mathbb{R}^3 &\longrightarrow \mathbb{R} \\ (x, y, z) &\longmapsto xy - z^2. \end{aligned}$$

We saw back in Example 5.2.3 that the signature of \det was $(1, 2, 0)$, so its zero set \mathcal{C} is a cone. See Figure 7.1. The cone \mathcal{C} divides the space in 3 components: The two positive ones we called \mathcal{H}^+ ; they are convex and thus not ample. The third component \mathcal{H}^- is the exterior of the cone; it corresponds to the matrices with negative determinant and it is ample. In particular, hyperbolicity is equivalent to $\det \circ \Psi(-) < 0$ and thus equivalent to $\Psi(-) \in \mathcal{H}^-$. Similarly, ellipticity is equivalent to $\det \circ \Psi(-) > 0$ and thus equivalent to $\Psi(-) \in \mathcal{H}^+$. We will come back to ellipticity in Lemma 7.2.4 below; for now we focus on hyperbolicity. We want to check ampleness; there are various possibilities, depending on what the image of Ψ is.

Suppose that $\lambda|_{\xi}$ is zero. Then the subspace K defined before the proof is zero as well. In particular, Ψ is constant and its image must be in \mathcal{H}^- , since F is a formal solution. It follows that $\mathcal{S}_{\lambda,F}^{\text{hyp}} = \text{Pr}_{\lambda,F}$, so ampleness holds trivially. Situation (ii.) holds. We henceforth assume $\lambda|_{\xi} \neq 0$ and thus $K \neq 0$.

Suppose that the forms $dF|_{\xi \cap \text{Ker}(\lambda)}$ are both zero. This is equivalent to $L \subset K$. This means that both $dF_i|_{\xi}$ are proportional to λ . However, this readily implies that $dF_i \wedge dF_i = dF_1 \wedge dF_2 = 0$, contradicting the fact that F was a formal solution.

Suppose that the forms $(dF_i|_{\xi \cap \text{Ker}(\lambda)})_{i=1,2}$ are linearly independent. This amounts to transversality of L and K . It follows that Ψ is surjective. Then, ampleness of \mathcal{S}^{hyp} along $\text{Pr}_{\lambda,F}$ is equivalent to the ampleness of the subspace \mathcal{H}^- of symmetric matrices with negative determinant (and thus, of mixed signature). This space is indeed ample, but not trivially. Situation (i.) holds.

Lastly, suppose that the $dF|_{\xi \cap \text{Ker}(\lambda)}$ are linearly dependent but not identically zero; i.e. $L \cap K$ is 1-dimensional. Up to changing basis we may assume that $dF_1|_{\xi \cap \text{Ker}(\lambda)} = 0$; i.e. dF_1 spans the line




Figure 7.1: The space of 2-by-2 symmetric matrices. In blue, the cone of degenerate matrices. The outside of the cone \mathcal{H}^- corresponds to matrices with negative determinant. It has a single ample component. The subspace of matrices with positive determinant \mathcal{H}^+ has two non-ample components. The image of Ψ can be a single point, the whole space, or a vertical plane A , shown in green. In the latter case, the intersection of A with the cone is a single line, cutting A in two non-ample components.

$L \cap K$. In this case, $dF_2|_{\xi \cap \text{Ker}(\lambda)} \neq 0$. Then $dF_1 \wedge dF_1 = 0$, so the image of Ψ is a 2-dimensional plane A through the origin, tangent to the cone. The restriction $\mathcal{H}^- \cap A$ consists of two half-spaces, separated by the line $A \cap \mathcal{C}$. Since this complement $A \cap \mathcal{C}$ is linear and of codimension-1, it is not a thin singularity. Situation (iii.) holds.

7.2.1.3 Conclusion of the first avoidance step

We now describe $\text{Avoid}(\mathcal{S}^{\text{hyp}})$ using Proposition 7.2.2.

Corollary 7.2.3 *$\text{Avoid}(\mathcal{S}^{\text{hyp}})$ consists of those pairs $(F, \Xi) \in \mathcal{S}^{\text{hyp}} \times_M \text{H-Conf}(TM)$ such that*

$$(F, \lambda) \notin \Sigma^{(1)} \quad \text{for every codirection } \lambda \in \Xi.$$

Proof. Write ξ for the kernel of $j^0 F$. Recall Situations (i.), (ii.) and (iii.) from Proposition 7.2.2. We define $\Delta_3 \subset \mathcal{S}^{\text{hyp}} \times_M \text{H-Conf}(TM)$ as the subspace of pairs (F, Ξ) for which Situation (iii.) holds for F and some $\lambda \in \Xi$. We define $\Delta_2 \subset \mathcal{S}^{\text{hyp}} \times_M \text{H-Conf}(TM)$ as the subspace of those $(F, \Xi) \notin \Delta_3$ such that Situation (ii.) holds for some $\lambda \in \Xi$.

According to Proposition 7.2.2, Δ_3 consists exactly of those pairs (F, Ξ) such that $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is not ample, for some $\lambda \in \Xi$. By definition, it follows that:

$$\text{Avoid}(\mathcal{S}^{\text{hyp}}) = \mathcal{S}^{\text{hyp}} \times_M \text{H-Conf}(TM) \setminus \overline{\Delta_3}.$$

We claim that the closure $\overline{\Delta_3}$ is exactly $\Delta_2 \cup \Delta_3$.

We first observe that the closure is contained in the union. Indeed, the pairs (F, λ) satisfying Situation (ii.) or (iii.) are precisely the elements of $\Sigma^{(1)}$. Lemma 7.2.1 states that this is a closed subset.

We then prove $\Delta_2 \subset \overline{\Delta_3}$. For fixed F , the set of $\nu \in T^*M$ such that Situation (iii.) holds for (F, ν) is non-empty, invariant under scalings of ν , and depends only on the restriction $\nu|_\xi$. Take $\lambda \in \Xi$ such that $\lambda|_\xi = 0$. Then there is a neighbourhood of λ in T^*M that submerses onto a neighbourhood of the zero section in ξ^* . This implies that any neighbourhood of λ contains codirections ν such that Situation (iii.) holds for (F, ν) . This proves the claim and concludes the proof.

7.2.1.4 Ampleness fails for (4,6) elliptic distributions

For completeness, we observe:

Lemma 7.2.4 *Let $F \in \mathcal{S}^{\text{ell}}$ and $\lambda \in T_p^*M$, both based at the same point. Write $\xi \subset T_pM$ for the 4-plane defined by j^0F . Then:*

- $\mathcal{S}_{\lambda, F}^{\text{ell}}$ is trivially ample if and only if $\lambda|_\xi = 0$.
- $\mathcal{S}_{\lambda, F}^{\text{ell}}$ is not ample otherwise.

*In particular, for fixed F , the set of $\lambda \in T_p^*M$ such that $\mathcal{S}_{\lambda, F}^{\text{ell}}$ is not ample is open and dense.*

Proof. We reason as in the proof of Proposition 7.2.2. Using the same notation as there, we have that the image of Ψ is either a point, a plane through the origin tangent to the cone \mathcal{C} , or the whole of \mathbb{R}^3 . The first case corresponds to trivial ampleness and to $\lambda|_\xi = 0$. The second case cannot happen, as the plane would be disjoint from \mathcal{H}^+ . The last case, corresponding to $(F, \lambda) \notin \Sigma^{(1)}$, is not ample as \mathcal{H}^+ consists of two convex components. Density follows from the fact that $\lambda|_\xi = 0$ cuts out a linear subspace.

Reasoning as in Corollary 7.2.3 implies that:

Corollary 7.2.5 *Avoid(\mathcal{S}^{ell}) is empty.*

Which is exactly the same situation as in the contact case (Lemma 6.1.8). This leads us to conjecture that there is no full h -principle for elliptic (4,6) distributions.

7.2.2 Second avoidance step

We will prove now that $\text{Avoid}(\mathcal{S}^{\text{hyp}})$ is not ample along all principal subspaces, so further elements have to be removed. This is different from the example in Section 5.5, where a single thinning step was sufficient to produce a thinning template.

Applying standard avoidance would lead us to study $\text{Avoid}^2(\mathcal{S}^{\text{hyp}})$. It turns out that it is difficult to determine whether $\text{Avoid}^2(\mathcal{S}^{\text{hyp}})$ is an avoidance template. The reason is that we do not have a explicit description of the elements removed from $\text{Avoid}(\mathcal{S}^{\text{hyp}})$ to yield $\text{Avoid}^2(\mathcal{S}^{\text{hyp}})$.

Due to this, it is more fruitful to ignore $\text{Avoid}^2(\mathcal{S}^{\text{hyp}})$ altogether and instead construct an avoidance template $\mathcal{A} \subset \text{Avoid}(\mathcal{S}^{\text{hyp}})$ with more transparent properties.

7.2.2.1 The singularity of interest

We define a new singularity

$$\Sigma^{(2)} \subset \mathcal{S}^{\text{hyp}} \times_M T^*M \times_M T^*M$$

as the subspace of pairs $(F, \lambda_1, \lambda_2)$ such that

$$j^0 F_1 \wedge j^0 F_2 \wedge \lambda_1 \wedge \lambda_2 \wedge dF_1 = j^0 F_1 \wedge j^0 F_2 \wedge \lambda_1 \wedge \lambda_2 \wedge dF_2 = 0. \quad (7.3)$$

It is also convenient to denote $\Sigma^{(2)}(\lambda_1, \lambda_2) \subset \mathcal{S}^{\text{hyp}}$ for the subset of elements F such that $(F, \lambda_1, \lambda_2) \in \Sigma^{(2)}$. Similarly we define $\Sigma^{(2)}(F)$.

We note again that the expressions in Equation 7.3 are algebraic on their entries. Furthermore, these expressions are non-trivial on the λ_i as long as $j^0 F_1 \wedge j^0 F_2 \wedge dF \neq 0$. This proves:

Lemma 7.2.6 *The following statements hold:*

- $\Sigma^{(2)}$ is a closed subset of $\mathcal{S}^{\text{hyp}} \times_M T^*M \times_M T^*M$.
- All the fibres $(\Sigma_p^{(2)})_{p \in M}$ are isomorphic algebraic subvarieties.
- Fix $F \in \mathcal{S}^{\text{hyp}}$ lying over p . The subspace $\Sigma^{(2)}(F)$ has positive codimension in $T_p^*M \oplus T_p^*M$.

Write $\xi \subset T_p M$ for the 4-plane given as the kernel of $j^0 F$. We note that $(F, \lambda_1, \lambda_2) \in \Sigma^{(2)}$ if and only if at least one of the following conditions holds:

- $(\lambda_1 \wedge \lambda_2)|_\xi$ is zero.
- Both 2-forms $(dF_i)|_{\xi \cap \text{Ker}(\lambda_1) \cap \text{Ker}(\lambda_2)}$ are zero. Equivalently, both $(dF_i)|_\xi$ are proportional to $(\lambda_1 \wedge \lambda_2)|_\xi$.

7.2.2.2 Main statement

We now study the ampleness of $\text{Avoid}(\mathcal{S}^{\text{hyp}})$. This amounts to the following: Given a collection of codirections Ξ and a codirection $\lambda \in \Xi$, we try to determine whether

$$\mathcal{S}_{\lambda, F}^{\text{hyp}} \setminus \left(\bigcup_{\nu \in \Xi} \Sigma^{(1)}(\nu) \right)$$

is an ample subset of $\text{Pr}_{\lambda, F}$, for each formal solution F .

Proposition 7.2.7 *Fix codirections $\lambda, \nu \in T_p^*M$ and a formal datum $F \in \mathcal{S}^{\text{hyp}}$, based also at p . Write $\xi \subset T_p M$ for the 4-plane cut out by $j^0 F$.*

The following statements hold:

- Suppose $F \notin \Sigma^{(2)}(\lambda, \nu)$. Then $\text{Pr}_{\lambda, F} \cap \Sigma^{(1)}(\nu)$ is a thin singularity.
- Suppose $F \in \Sigma^{(2)}(\lambda, \nu)$ but $\lambda \wedge \nu|_\xi \neq 0$. Then $\text{Pr}_{\lambda, F} \setminus \Sigma^{(1)}(\nu)$ is ample but its complement is of codimension 1.

Proof. In both situations we are assuming that the forms $\lambda|_\xi$ and $\nu|_\xi$ are linearly independent. This allows us to define the restriction map

$$\Phi : \text{Pr}_{\lambda, F} \longrightarrow \wedge^2(\xi \cap \text{Ker}(\nu))^* \oplus \wedge^2(\xi \cap \text{Ker}(\lambda))^*$$

that sends $\tilde{F} \in \text{Pr}_{\lambda, F}$ to $d\tilde{F}|_{\xi \cap \text{Ker}(\nu)}$. Recall that $\tilde{F} \in \Sigma^{(1)}(\nu)$ if and only if the pair $d\tilde{F}|_{\xi \cap \text{Ker}(\nu)}$ is linearly dependent. The upcoming argument follows the proof of Lemma 5.5.2 (Section 5.5), since the linear dependence problems under consideration are exactly the same.

Write $L \subset \wedge^2(\xi \cap \text{Ker}(\nu))^*$ for the subspace of 2-forms proportional to $\lambda|_{\xi \cap \text{Ker}(\nu)}$. Due to the linear independence of $\lambda|_{\xi}$ and $\nu|_{\xi}$, L is a 2-dimensional plane. We write L_i for the plane passing through dF_i parallel to L . Then, the image of $\text{Pr}_{\lambda, F}$ under Φ is the sum $L_1 \oplus L_2$, which is 4-dimensional.

The condition $F \in \Sigma^{(2)}(\lambda, \nu)$ is equivalent to $\Phi(F) = 0$, which in turn is equivalent to $L_1 = L_2 = L$. Ampleness of $\text{Pr}_{\lambda, F} \setminus \Sigma^{(1)}(\nu)$ is thus equivalent to ampleness of $\text{GL}_2 \subset M_{2 \times 2}$, proving the second claim.

Similarly, $F \notin \Sigma^{(2)}(\lambda, \nu)$ means that $\Phi(F) \neq 0$. I.e. at least one L_i is different from L and therefore does not pass through the origin. We deduce the singularity $\Sigma^{(1)}(\nu)$ has codimension 2.

Remark 7.2.8 Consider the following elementary facts:

- Removing a thin singularity from an ample set still yields an ample set.
- The intersection of two ample sets need not be ample. For instance, the union of two codimension-1 singularities, both having ample complement individually, may separate the space into non-ample pieces. See Figure 7.2.

Our claim is that these statements largely determine the steps to be taken during avoidance.

Indeed, suppose in our current example that $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is ample, for some λ and F . Then, the condition $F \notin \Sigma^{(2)}(\lambda, \nu)$ implies that $\mathcal{S}_{\lambda, F}^{\text{hyp}} \setminus \Sigma^{(1)}(\nu)$ is ample, according to the first fact. This may still be true even if $F \in \Sigma^{(2)}(\lambda, \nu)$, but the second fact tells us that it need not be. This is even more delicate when there are several codirections involved (which will always be the case in the construction of a template). Avoiding the uncertainty of the second situation effectively forces us to consider the singularity $\Sigma^{(2)}$, and thus prescribes what the second avoidance step must be.

Our claim is that this type of analysis, which is algorithmic in nature, is not specific to \mathcal{S}^{hyp} . Indeed, it must guide the avoidance process of any given differential relation.




Figure 7.2: On the left and middle, two cubics in \mathbb{R}^2 . Their complements (as well as themselves) are ample subsets. On the right, we intersect the complements, yielding four components, none of which is ample. The intersection of the two cubics is a point, also not ample.

7.2.2.3 Conclusion of the second step

Using Proposition 7.2.7 we now define:

$$\mathcal{A} := \{(F, \Xi) \in \text{Avoid}(\mathcal{S}^{\text{hyp}}) \mid (F, \lambda_1, \lambda_2) \notin \Sigma^{(2)} \text{ for all } \lambda_1 \neq \lambda_2 \in \Xi\}.$$

Lemma 7.2.9 \mathcal{A} is a pretemplate.

Proof. As we noted in Lemma 7.2.6, the condition $(F, \lambda_1, \lambda_2) \in \Sigma^{(2)}$ is closed, smooth in its entries, and algebraic over each given point $p \in M$. Furthermore, for a given F formal solution of \mathcal{S}^{hyp} , the condition is non-trivial on λ_1 and λ_2 . This proves the claim.

7.2.3 End of the proof

The proof of Theorem 1.12 will be complete once we show that:

Theorem 7.1. \mathcal{A} is a template for \mathcal{S}^{hyp} .

This statement requires the following auxiliary result:

Proposition 7.2.10 Fix codirections $\lambda, \nu_1, \nu_2 \in T_p^*M$ and a formal datum $F \notin \Sigma^{(2)}(\nu_1, \nu_2)$, based also at p . Write $\xi \subset T_pM$ for the 4-plane cut out by j^0F .

Then $\text{Pr}_{\lambda, F} \cap \Sigma^{(2)}(\nu_1, \nu_2)$ is a thin singularity.

Proof. Recall that $\Sigma^{(2)}$ is defined by Equation 7.3. The condition $F \notin \Sigma^{(2)}(\nu_1, \nu_2)$ implies that the forms $\nu_i|_{\xi}$ are linearly independent. As such, Equation 7.3 reads:

$$(dF_1)|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)} = (dF_2)|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)} = 0.$$

Suppose first that λ is in the span $\langle \nu_1|_{\xi}, \nu_2|_{\xi} \rangle$. Any given element $\tilde{F} \in \text{Pr}_{\lambda, F}$ then satisfies

$$d\tilde{F}|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)} = dF|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)},$$

proving that $\text{Pr}_{\lambda, F} \cap \Sigma^{(2)}(\nu_1, \nu_2)$ must be empty.

Suppose instead that λ is linearly independent from the other two forms. We can write out $d\tilde{F}|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)}$ as a pair of 1-forms

$$(dF_i + \lambda \wedge \beta_i)|_{\xi \cap \text{Ker}(\nu_1) \cap \text{Ker}(\nu_2)}.$$

Equation 7.3 provides then two independent, affine, codimension-1 constraints, one for each β_i . The locus cut out by both has then codimension-2, proving thinness.

Proof (Proof of Theorem 7.1). First we prove that \mathcal{A} satisfies Property (II) in the definition of template. Fix $(F, \Xi) \in \mathcal{A}$ and $\lambda \in \Xi$. By construction, such a pair is characterised by the following properties:

- $F \in \mathcal{S}^{\text{hyp}}$.
- $F \notin \Sigma^{(1)}(\nu)$ for all $\nu \in \Xi$.
- $F \notin \Sigma^{(2)}(\nu_1, \nu_2)$ for all pairs $\nu_1 \neq \nu_2 \in \Xi$.

We need to show that $\mathcal{A}(\Xi)$ is ample along $\text{Pr}_{\lambda, F}$. An explicit description reads:

$$\mathcal{A}(\Xi)_{\lambda, F} = \mathcal{S}_{\lambda, F}^{\text{hyp}} \setminus \left[\bigcup_{\nu \in \Xi} \Sigma^{(1)}(\nu) \cup \bigcup_{\nu_1 \neq \nu_2 \in \Xi} \Sigma^{(2)}(\nu_1, \nu_2) \right].$$

We then observe:

- Using condition $F \notin \Sigma^{(1)}(\lambda)$, we apply Proposition 7.2.2 to deduce that $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is ample.
- Using condition $F \notin \Sigma^{(2)}(\lambda, \nu)$, we apply Proposition 7.2.7 and deduce that each $\Sigma^{(1)}(\nu)$ is thin.
- According to Proposition 7.2.10, the singularities $\Sigma^{(2)}(\nu_1, \nu_2)$ are all thin.

The claim follows because ampleness is preserved upon removal of thin singularities.

Secondly, we prove that \mathcal{A} satisfies Property (III) in the definition of template. We need to show that, for each $F \in \mathcal{S}^{\text{hyp}}$ lying over $p \in M$, the subspace $\mathcal{A}(F)$ is dense in $\text{H-Conf}(T_p M)$. Its complement can be explicitly written down as

$$\mathcal{A}(F)^c \cap \text{H-Conf}(T_p M) = \{\Xi \in \text{H-Conf}(T_p M) \mid F \in \Sigma^{(1)}(\nu_1) \cup \Sigma^{(2)}(\nu_1, \nu_2) \text{ for some } \nu_1 \neq \nu_2 \in \Xi\}.$$

The conditions $\nu_1 \in \Sigma^{(1)}(F)$ or $(\nu_1, \nu_2) \in \Sigma^{(2)}(F)$ are given, respectively, by Equations 7.1 and 7.3. Both of them are non-trivial on their entries, since F is in \mathcal{S}^{hyp} ; see Lemmas 7.2.1 and 7.2.6. We deduce that $\mathcal{A}(F)^c \cap \text{H-Conf}(T_p M)$ is a finite union of positive-codimension subvarieties, proving the claim.

Proof (Proof of Theorem 1.12). According to Lemma 7.1.9, the h -principle for \mathcal{R}^{hyp} reduces to the h -principle for \mathcal{S}^{hyp} . Theorem 1.8 says that the h -principle for \mathcal{S}^{hyp} follows from the existence of an avoidance template. Theorem 7.1 yields such a template \mathcal{A} .

Part IV

Future work

Chapter 8

Future work

In this last chapter we discuss some possible research lines to explore in the short-mid term taking this PhD thesis as starting point. We state some of the expected/conjectured results based on our current understanding of the global topology of bracket-generating distributions.

8.1 Transverse submanifolds.

In this thesis we proved an h -principle for 1-dimensional embeddings transverse to bracket-generating distributions. We intend to address, in future work, the analogous problem for higher dimensional transverse submanifolds.

We expect the h -principle to hold for n -dimensional submanifolds transverse to corank k distributions for certain values of $n \leq k$ (recall that closed n -dimensional submanifolds transverse to corank k distributions abide by all forms of the h -principle if $k > n$, see [39, 4.6.2]).

The role that higher dimensional transverse submanifolds can play in showing flexibility for distributions becomes apparent in the paper [92], where del Pino and Vogel construct transverse 2-tori to Engel distributions and manipulate them in order to prove a complete h -principle for overtwisted Engel structures.

Therefore, we expect this type of submanifolds to be key in showing flexibility for broader families of bracket-generating distributions.

8.2 Spaces of distributions.

8.2.1 Towards an h -principle for maximal-growth distributions of higher step.

We have shown that step-2 distributions of maximal growth abide by a complete h -principle but, what about higher step distributions? We expect the argument to generalise. In other words, we expect an h -principle for distributions \mathcal{D} of step- r ($r \geq 2$) to hold in the following case:

- $\text{rank}(\mathcal{D}) > 2$ and
- the growth vector of \mathcal{D} is of maximal growth.

This is ongoing work in progress and, in order to give a better understanding of the situation, the following examples aim to illustrate the bracket generating condition when regarded along

horizontal line fields. This is a potential strategy to tackle the problem through convex integration, since this description becomes relevant when studying the associated differential relation along principal subspaces.

Let's start with two known examples: contact distributions and Engel structures.

8.2.1.1 Contact distributions as immersed curves in \mathbb{S}^1 .

Consider a 2- distribution \mathcal{D} in \mathbb{R}^3 and a vector field $Y \in \mathcal{D}$. Consider now a flowbox chart for Y (where Y is seen as a coordinate direction ∂_t) and take a germ of a transverse disk $\mathbb{D}^2(x, y)$ transverse to ∂_t around $p = (x_0, y_0, t_0)$. For

Let X_t be vector field tangent to the leaves of the foliation $\sqcup_t \mathbb{D}^2 \times \{t\}$. It can be regarded as a 1-parametric family of vector fields such that $\mathcal{D}_{(x,y,t)} = \langle \partial_t, X_t \rangle$ and $\|X_t\| = 1$ (we can thus regard the 1-parametric family X_t as a curve $\gamma : \mathbb{R} \rightarrow \mathbb{S}^1 \subset T\mathbb{D}^2$). Denoting $\dot{X}_t = [\partial_t, X_t]$, the contact condition then reads as

$$\dim \langle \partial_t, X_t, \dot{X}_t \rangle = 3$$

By construction, since $\dot{X}_t = [\partial_t, X_t] = \frac{\partial X_t}{\partial t}$, X_t and \dot{X}_t are both linearly independent with respect to ∂_t and the contact condition thus reads as $\dim \langle X_t, \dot{X}_t \rangle = 2$. But, from the condition $\|X_t\| = 1$ we get that $\frac{d}{dt} \langle X_t, X_t \rangle = \frac{1}{2} \langle X_t, \dot{X}_t \rangle = 0$ and, thus, X_t and \dot{X}_t are orthogonal. Therefore, the contact condition reads as X_t (when regarded as a curve in \mathbb{S}^1) being an immersion. Another way of interpreting it is noticing that the vector X_t must be turning (without stop) along the direction ∂_t .




Figure 8.1: The contact condition for $\mathcal{D} = \langle \partial_t, X_t \rangle$ locally reads as X_t being an immersed curve when regarded as a curve into \mathbb{S}^1 . In terms of vector fields, the Legendrian vector X_t must turn positively along the direction ∂_t .

8.2.1.2 Engel distributions as convex curves in \mathbb{S}^2 .

(First described in [29] (Sec. 2.2), and subsequently treated in [91], [87], [25], [92]). Consider a rank-2 distribution $(\mathbb{R}^4, \mathcal{D}^2)$ and a vector field $Y \in \mathcal{D}$ which, after taking a flowbox chart, becomes the coordinate direction ∂_t . Take a germ of a 3-disk $\mathbb{D}^3(x, y, z)$ transverse to ∂_t around a point $p = (x_0, y_0, z_0, t_0)$.

Again, take X_t a vector field tangent to the leaves of the foliation $\sqcup_t \mathbb{D}^3 \times \{t\}$ (as before, it can be regarded as a 1-parametric family of vector fields such that $\mathcal{D}_{(x,y,z,t)} = \langle \partial_t, X_t \rangle$). Also, we can impose that $\|X_t\| = 1$ and we can thus regard the 1-parametric family X_t as a curve $\gamma : \mathbb{R} \rightarrow \mathbb{S}^2$. Denoting $\dot{X}_t = [\partial_t, X_t]$ and $\ddot{X}_t = [\partial_t, \dot{X}_t]$, the Engel condition then reads as:

- $\dim \langle \partial_t, X_t, \dot{X}_t \rangle = 3$ and
- $\dim \langle \partial_t, X_t, \dot{X}_t, [X_t, \dot{X}_t], \ddot{X}_t \rangle = 4$

By the same argument as in the previous example, a sufficient condition for a distribution \mathcal{D}_t to be Engel is given by $\dot{X}_t \neq 0$ and $\ddot{X}_t \neq 0$ which, in terms of curves into \mathbb{S}^2 , translates as the curve described by X_t being convex.




Figure 8.2: In view of the aforementioned Engel conditions for $\mathcal{D} = \langle \partial_t, X_t \rangle$, a sufficient condition for being Engel is X_t being convex when regarded as a curve into the 2-sphere.

Observe that the other sufficient condition for \mathcal{D}_t to be Engel is given by $\dot{X}_t \neq 0$ and $[X_t, \dot{X}_t] \neq 0$. This tantamounts to $\xi_t := \langle X_t, \dot{X}_t \rangle$ locally defining a contact structure in the transverse 3-disk \mathbb{D}^3 .

8.2.1.3 Bracket-generating distributions of higher step

Consider a manifold equipped with a k -rank distribution (M, \mathcal{D}) and associated Lie flag

$$\mathcal{D} \subset \mathcal{D}_2 \subset \cdots \subset \mathcal{D}_r \subset TM$$

Proceeding in the same fashion as in the previous examples, we take a vector field $Y \in \mathcal{D}$ and locally regard it (after taking a flowbox chart) as the coordinate direction ∂_t . We consider a germ of an $n - 1$ -disk \mathbb{D}^{n-1} transverse to ∂_t around a point $p = (x_0, y_0, z_0, t_0)$.

Take now a set $\mathfrak{B}\mathfrak{r}_t^1$ of $k - 1$ different 1-parametric vector fields $\mathfrak{B}\mathfrak{r}_t^1 := \{X_t^1, \dots, X_t^{k-1}\}$ such that $\mathcal{D} = \langle X_t^1, \dots, X_t^{k-1} \rangle$. Denote $X^0 := \partial_t$. For $i > 1$, let $\mathfrak{B}\mathfrak{r}_t^i$ denote the set of brackets of vector fields $\partial_t, X^0, X_t^1, \dots, X_t^{k-1}$ (possibly with repetitions) of length less or equal than i where ∂_t does not appear $i - 1$ times in each bracket-expression:

$$\mathfrak{B}\mathfrak{r}_t^i := \{[X^{\ell_q}, \dots, [X^{\ell_2}, [X^{\ell_1}, X^{\ell_0}]] \dots] : q < i, \sigma \circ (X_{\ell_0}, \dots, X_{\ell_q}) \neq (\partial_t, \dots, \partial_t, X_j) \forall \sigma \in \Sigma_i, j \leq k - 1\}$$

The condition for \mathcal{D} having growth vector $\nu_{\mathcal{D}} = (n_1, n_2, \dots, n_i, \dots)$ then translates as

$$\text{For each } i \geq 1, \dim \left(\langle \partial_t, \mathfrak{B}\mathfrak{r}_t^1, \mathfrak{B}\mathfrak{r}_t^i, \partial_t^{i-1}(X_t^1), \dots, \partial_t^{i-1}(X_t^{k-1}) \rangle \right) = n_i. \quad (8.1)$$

Here $\partial_t^0(X_t^j)$ just means X_t^j . There are two cases where condition (8.1) will become relevant:

C1) In the case of growth vectors of free Lie algebras; i.e. $n_i - n_{i-1}$ is the dimension of a free Lie algebra. Then *every generated bracket* must contribute to producing new directions. Thus, a necessary condition for condition (8.1) to be satisfied is that for each $i \geq 1$, the vectors $\partial_t^i(X_t^1), \dots, \partial_t^i(X_t^{k-1})$ must be linearly independent in the quotient space $\mathbb{R}^{\dim(M)} / \langle \partial_t, \mathfrak{B}\mathfrak{r}_t^1, \mathfrak{B}\mathfrak{r}_t^i \rangle$. This will be the case for all elements \mathcal{D}_i in the flag of a maximal growth distribution except for, possibly, $\mathcal{D}_{\text{step}(\mathcal{D})-1}$.

C2) Let n_i be the first entry in the growth vector for which $\dim(n_i) = \dim(M)$ holds. Whenever the inequality $\dim(\langle \partial_t, \mathfrak{B}\mathfrak{r}_t^1, \mathfrak{B}\mathfrak{r}_t^{i-1} \rangle) \geq n_i - n_{i-1} - (k-1)$ holds, a direct computation shows that the condition of the vectors $\partial_t^i(X_t^1), \dots, \partial_t^i(X_t^{k-1})$ being linearly independent in the quotient space $\mathbb{R}^{\dim(M)} / \langle \partial_t, \mathfrak{B}\mathfrak{r}_t^1, \mathfrak{B}\mathfrak{r}_t^i \rangle$ is also sufficient for \mathcal{D} being bracket-generating of maximal growth. By normalizing the vectors X_t^α , this generalises the immersion condition for contact distributions (8.2.1.1) and the convexity condition for Engel structures (8.2.1.2).

After generalising to the setting of general bracket-generating distributions (Subsection 8.2.1.3) the result in [29] where Engel structures along horizontal line fields are regarded as curves in \mathbb{S}^2 (Example 8.2.1.2), we intend to fit this geometric description into the general scheme of M. Gromov's higher order convex integration [59]. In fact, describing this geometric condition along horizontal line fields can be understood as describing the associated differential relation along principal subspaces which lie (locally) entirely in the distribution. We claim that the rest of directions do not contribute in terms of ampleness and, thus, these are the only relevant directions to check.

So, a natural question is why we need $\text{rank}(\mathcal{D}) > 2$. First, note that the contact case does not abide by a complete h -principle. The reason is that when one tries to check if ampleness holds by observing the conditions described in C1) and C2), it turns out that this will only follow if $\text{rank}(\mathcal{D}) \geq 3$.

The geometric intuition is as follows: once we fix a principal direction in \mathcal{D} (for $\text{rank}(\mathcal{D}) = 2$), we have only one degree of freedom (by choosing another vector along the principal direction so as to complete a frame of \mathcal{D}) in order to introduce oscillations to our initial solution. When checking ampleness, this does not give raise to ample sets. Instead, if $\text{rank}(\mathcal{D}) \geq 3$, we have at least two degrees of freedom and ampleness will follow.

This different behaviour comes from the fact that we can reduce the problem to checking if ampleness holds for $GL(r)$ inside the space of all matrices for certain values of $r > 0$. And, as we saw, $GL(r)$ is ample in the space of all matrices if and only if $r \geq 2$, which translates by our reduction into $\text{rank}(\mathcal{D}) \geq 3$.

8.2.2 Spaces of maximally non-integrable corank 2 distributions.

Part of the goal of Part III in this thesis was to pinpoint those pairs (k, n) , consisting of a rank k and a dimension n , for which the relation defining maximally non-involutive distributions is not ample. Our goal was to narrow down the (open and Diff-invariant) classes of distributions that may display rigid behaviours (as contact structures do). Our discussion in the Introduction (see 1.6.2) provides a list of candidates in dimensions up to 6. Proving rigidity and/or constructing suitable overtwisted classes in each case are interesting open questions.

Beyond dimension 6, we propose (jointly with I. Zelenko and Álvaro del Pino), the following concrete conjecture. Consider maximally non-involutive distributions of corank-2. These are always of step 2, with the exception of dimension 4 (the Engel case, which we leave out of the discussion). Then:

- In odd rank $(2l + 1, 2l + 3)$, maximal non-involutivity means that the top-wedge of the pencil of curvatures has maximal linear span (i.e. of dimension $l + 1$). The differential relation is then the complement of a singularity of codimension l . We expect flexibility to hold due to (classical) ampleness.
- In even rank $(4l, 4l + 2)$ we see an elliptic versus hyperbolic dichotomy (just like for $(4, 6)$). We expect flexibility to hold in the hyperbolic case, but avoidance to be necessary to prove it. Elliptic distributions are good candidates for rigidity.
- In even rank $(4l + 2, 4l + 4)$ maximal non-involutivity is equivalent to hyperbolicity. We expect flexibility based on the developed new technique convex integration up to avoidance.

We intend to address this conjecture in future work.

Acknowledgements

During the development of this PhD thesis the author has been supported by the **Programa Predoctoral de Formación de Personal Investigador No Doctor** scheme funded by the **Basque department of education** (“**Departamento de Educación del Gobierno Vasco**”).

Durante el desarrollo de esta Tesis Doctoral, el autor ha sido financiado por el **Programa Predoctoral de Formación de Personal Investigador No Doctor del Departamento de Educación del Gobierno Vasco**.

First of all, I would like to show my gratitude to my thesis advisors Francisco Presas and Álvaro del Pino for their exceptional guidance and for being a source of inspiration and support during the development of this thesis. Fran, your many valuable insights as well as your generosity with time and ideas are immensely appreciated. Álvaro, thank you for generously sharing your knowledge with me and for your warm hospitality during my many visits to Utrecht. You both have immensely contributed, through your wisdom and encouragement, to my mathematical and personal development.

Thank very much, Fabio Gironella, for your thorough and serious work reading my thesis. You have definitely contributed to greatly improve this work with your many comments and fine observations. Thank you very much to Thomas Vogel as well for his comments and good questions regarding the contents of my thesis.

I would also like to thank all the members of Fran’s group for their inspiration and for the many hours of mathematical discussions: Eduardo, Guillermo, Samuel, Dayhana, José Luis, Álvaro, Luis, Sergio... Among them, I would specially want to thank my friend and coauthor Eduardo Fernández. Thank you very much for the many insights and hours of discussions and, more importantly, for your kindness and friendship along these years. Also, I would like to thank Luis Hernández Corbato and José Luis Pérez for the very nice introductory lectures about Algebraic Geometry and Contact Topology during the year of my Master’s. Thank you, Daniel Peralta, for your inspiration, kindness and for accepting me in your project during my first year.

I am thankful to the Geometry Group of Marius Crainic in Utrecht for their warmth and kindness during my visits to Álvaro in Utrecht. Thank you Aldo, Dussan, Francesco, Luran, Luca, Marius and Marteen. You definitely made me feel like home.

Also, I would like to thank all the members from the Symplectic Geometry group of Eva Miranda in Barcelona. We have shared many nice hours during the many conferences where we have been together. Thank you Eva, Robert, Cédric, Anastasia, Joaquim, Pablo, Pau... .

Thank you very much to Ryan Budney for his kindness and hospitality during my visit, together with Eduardo, to Victoria (Canada). It was a huge pleasure learning from you about the fascinating world of embedding spaces.

During my teaching duties at Complutense University I have learned a lot from the more experienced Professors who have generously helped me during the three different courses I have taught. Thank you very much for your help and kindness, Inés Gómez Chacón, Luis Giraldo and José Manuel Rodríguez Sanjurjo.

Special thanks to Vicente Muñoz and Marco Castrillón for being my formal advisors at Complutense University. You both have always been willing to help me. Thank you very much.

Thank you to the many good teachers and professors that have contributed to my mathematical education from high school to University. I am thankful to the High School Problem Seminar of La Rioja University, wich inspired me to study Mathematics; with special mention to Emilio Fernández. Also, I had very good University Professors who greatly contributed to my mathematical education; among others, Alejandro Melle, José Manual Gamboa, Fernando Cobos, Socorro Ponte, José María Martínez Ansemil, Manuel Alonso Morón, Enrique Arrondo, Ignacio Sols and many others. Thank you, Alejandro Melle, for your warmth and inspiration during your advising work of my Bachelor Thesis.

During the developement of this PhD thesis I have learned a lot from many other mathematicians who have greatly inspired me with their ideas and knowledge. Thank you, Emmy Murphy, Daniel Álvarez-Gavela, Marián Poppr, Ángel González, David Jesús. Special thanks to Igor Zelenko for his generosity and his many valuable insights about bracket-generating distributions. I am thankful to many people that have supported me during the developement of this thesis. Special thanks to Víctor, Álex, David, Nejat, Jaime, Cristina, María, Hernán, Pedro, Miguel G., Miguel M., Dominique and all my close friends.

Muchas gracias a mis padres por todo el apoyo, la dedicación y por su trabajo, al cual debo mi formación. Muchas gracias también a mis abuelos y al resto de mi familia por todo su cariño y apoyo.

References

1. J. Adachi. *Classification of horizontal loops in standard Engel space*. Int. Math. Res. Not. 2007:3 (2007).
2. J. Adachi. *Existence and classification of maximally non-integrable distributions of derived length one*. arXiv:2111.01403.
3. C. C. Adams. *The knot book. An elementary introduction to the mathematical theory of knots*. American Mathematical Society, Providence, RI, 2004. xiv+307 pp.
4. A.A. Agrachev, M. Caponigro. *Controllability on the Group of Diffeomorphisms*. Annales de l'I.H.P. Analyse non linéaire 26:6 (2009), 2503–2509.
5. D. Álvarez-Gavela. *Refinements of the holonomic approximation lemma*. Algebraic & Geometric Topology 18:4 (2018), 2265–2303.
6. D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization II: Geomorphology of Lagrangian ridges*. arXiv:1912.03439
7. D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization III: Positive arborealization of polarized Weinstein manifolds*. arXiv:2011.08962
8. D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization I: Stability of arboreal models*. arXiv:2101.04272
9. V.I. Arnol'd, S.P. Novikov. *Dynamical systems VII*. Enc. of Math. Sci., vol. 16, Springer-Verlag, Berlin Heidelberg, 1994.
10. V. I. Arnold, S. M. Gusein-Zade, A. N. Varchenko. *Singularities of Differentiable Maps. Volume I*. Birkhäuser, Boston, Basel, Stuttgart. 1985.
11. D. Bennequin, *Entrelacements et equations de Pfaff*. Asterisque 107–108 (1983), 87–161.
12. M.S. Borman, Y. Eliashberg, E. Murphy. *Existence and classification of overtwisted contact structures in all dimensions*. Acta Math. 215 (2015), no. 2, 281–361.
13. Bleiler, S.A. *Knots prime on many strings*. Transactions of the American Mathematical Society 282 (1984), 385–401.
14. R.L. Bryant, L. Hsu. *Rigidity of integral curves of rank 2 distributions*. Invent. Math. 114:2 (1993), 435–461.
15. J.S. Birman, W.W. Menasco. *Stabilization in the braid groups. II. Transversal simplicity of knots*. Geom. Topol. 10 (2006), 1425–1452.
16. A. Bhowmick. *The h-Principle for Maps Transverse to Bracket-Generating Distributions*. Preprint. ArXiv:205.04323.
17. F. Boarotto, A. Lerario. *Homotopy properties of horizontal path spaces and a theorem of Serre in sub-Riemannian geometry*. Communications in Analysis and Geometry 25:2 (2017), 269–301.
18. R. Budney. *A family of embedding spaces*. Geometry & Topology 13 (2008): 41–83.
19. R. Budney, Ryan. *Topology of knot spaces in dimension 3*. Proc. Lond. Math. Soc. (3) 101 (2010), no. 2, 477–496.

20. M. Caponigro. *Controllability on the Group of Diffeomorphisms*. PhD Thesis, 2009. SISSA, Trieste.
21. E. Cartan. *Sur quelques quadratures dont l'élément différentiel contient des fonctions arbitraires*. Bull. Soc. Math. France 29 (1901), 118–130.
22. E. Cartan. *Les systemes de Pfaff a cinq variables et les equations sur derivees partielles du second ordre*. Ann. Ec. Norm. 27 (1910).
23. R. Casals, O. Spáčil. *Chern-Weil theory and the group of strict contactomorphisms*. J. Topol. Anal. 8 (2016), no. 1, 59–87.
24. R. Casals, Á. del Pino. *Classification of Engel knots*. Math. Ann. 371(1-2) (2018), 391–404.
25. R. Casals, Á. del Pino, F. Presas. *Loose Engel structures*. Compos. Math. 156(2) (2020), 412–434.
26. Y. Chekanov. *Differential algebra of Legendrian links*. Invent. Math. 150 (2002), no. 3, 441–483.
27. W. L. Chow *Über Systeme von linearen partiellen Differentialgleichung erster Ordnung* Math. Ann. 117, (1939). 98–105.
28. J. H. Conway. *An enumeration of knots and links and some of their algebraic properties*. Computational Problems in Abstract Algebra. Pergamon Press. Oxford and New York. 1969. pp. 329–358.
29. R. Casals, J.L. Pérez, Á. del Pino, F. Presas. *Existence h -principle for Engel structures*. Invent. Math. 210 (2017), 417–451.
30. L. Dahinden, Á. del Pino. *Introducing sub-Riemannian and sub-Finsler Billiards*. Discr. Cont. Dyn. Sys. (2022).
31. S. Dave, S. Haller. *On 5-manifolds admitting rank two distributions of Cartan type*. Trans. Amer. Math. Soc. 371 (2019), 4911–4929. bibitemDen J. Denkers. *Semifree and Totally Non-Parallel Immersions*. Master thesis. Utrecht University.
32. F. Ding, H. Geiges. *The diffeotopy group of $\mathbb{S}^1 \times \mathbb{S}^2$ via contact topology*. Compos. Math. 146 (2010), no. 4, 1096–1112.
33. T. Ekholm, K. Honda, and T. Kálmán, *Legendrian knots and exact Lagrangian cobordisms*, J. Eur. Math. Soc. (JEMS) 18 (2016), no. 11, 2627–2689.
34. Y. Eliashberg. *Classification of overtwisted contact structures on 3-manifolds*. Invent. Math. 98 (1989), no. 3, 623–637.
35. Y. Eliashberg. *Contact 3-manifolds twenty years since J. Martinet's work*. Ann. Inst. Fourier 46 (1992), 165–192.
36. Y. Eliashberg. *Classification of contact structures on \mathbb{R}^3* . Internat. Math. Res. Notices 1993, no. 3, 87–91.
37. Y. Eliashberg, M. Fraser. *Topologically trivial Legendrian knots*. J. Symplectic Geom. 7 (2009), no. 2, 77–127.
38. Y. Eliashberg, D. Kotschick, E. Murphy, T. Vogel. *Engel structures*. AIM Workshop report (2017).
39. Y. Eliashberg, N. Mishachev. *Introduction to the h -principle*. Graduate Studies in Mathematics, 48. American Mathematical Society, Providence, RI, 2002.
40. Y. Eliashberg, N. Mishachev. *Wrinkled embeddings*. Foliations, geometry, and topology, 207–232, Contemp. Math., 498, Amer. Math. Soc., Providence, RI, 2009.
41. J. Etnyre. *Legendrian and transversal knots*. Handbook of knot theory, 105–185, Elsevier B. V., Amsterdam, 2005.
42. J. Etnyre and K. Honda. *Knots and contact geometry I: Torus knots and the figure eight knot*. J. Symplectic Geom. 1 (2001), 63–120.
43. J.B. Etnyre, K. Honda. *On connected sums and Legendrian knots*. Adv. Math. 179 (2003), no. 1, 59–74.
44. E. Fernández, J. Martínez-Aguinaga, F. Presas. *Fundamental groups of formal legendrian and horizontal embedding spaces*. Algebraic & Geometric Topology 20 (2020), 3219–3312.

45. E. Fernández, J. Martínez-Aguinaga, F. Presas. *Loops of Legendrians in contact 3-manifolds*. Classical and Quantum Physics. 60 Years Alberto Ibort Fest Geometry, Dynamics and Control. Springer Proceedings in Physics 229, pp 361–372. (2019).
46. E. Fernández, J. Martínez-Aguinaga, F. Presas. *The homotopy type of the contactomorphism groups of tight contact 3-manifolds, part I*. arXiv:2012.14948.
47. E. Fernández, J. Martínez-Aguinaga, F. Presas. *Parametric connected sums in the space of legendrian embeddings*. In preparation.
48. G. Frobenius, *Über das Pfaffsche Problem*. J. für Reine und Angew. Math., 82 (1877), 230–315.
49. D. Fuchs, S. Tabachnikov, Invariants of Legendrian and transverse knots in the standard contact space, *Topology*, 5 (1997), 1025–1053.
50. M. Gardner, *Knotted Doughnuts and Other Mathematical Entertainments*. New York: W. H. Freeman, 1986.
51. Z. Ge. *Horizontal path spaces and Carnot-Carathéodory metrics*. Pacific J. Math. Volume 161:2 (1993), 255–286.
52. H. Geiges. *Horizontal loops in Engel space*. Math. Ann. Volume 342:2 (2008), 291–296.
53. H. Geiges. *An Introduction to Contact Topology*. Cambr. Studies in Adv. Math. 109. Cambr. Univ. Press 2008.
54. E. Giroux *Sur les transformations de contact au-dessus des surfaces*. Essays on geometry and related topics, Vol. 1, 2, 329–350, Monogr. Enseign. Math., 38, Enseignement Math., Geneva, 2001.
55. T.G. Goodwillie, M.S. Weiss. *Embeddings from the point of view of immersion theory II*. Geometry and Topology 3 (1999), 103–118.
56. A. Gramain. *Sur le groupe fondamental de l'espace des noeuds*. Ann. Inst. Fourier. 27 (1977), 29–44.
57. M. Gromov. *A topological technique for the construction of solutions of differential equations and inequalities*. Actes du Congrès International des Mathématiciens (Nice, 1970), Tome 2. 1971, pp. 221–225.
58. M. Gromov. *Convex integration of partial differential relations*. I, Izv. Akad. Nauk SSSR Ser. Mat., 37(1973), 329–343.
59. M. Gromov. *Partial differential relations*. Ergebnisse Math. und ihrer Gren. , Springer-Verlag, Berlin, 1986.
60. M. Gromov, *Carnot-Carathéodory spaces seen from within*. Sub-Riemannian geometry, Progr. Math., vol. 144, Birkhäuser, Basel, 1996.
61. A. M. Vershik, V. Y. Gershkovich, *Determination of the functional dimension of the space of orbits of germs of generic distributions*. Mat. Zametki, 44:5 (1988), 596–603; Math. Notes, 44:5 (1988), 806–810.
62. R. S. Hamilton. *The inverse function theorem of Nash and Moser*. Bull. Amer. Math. Soc. (N.S.) 7 (1982), no. 1, 65–222.
63. A. Hatcher. *A proof of the Smale conjecture, $\text{Diff}(\mathbb{S}^3) \simeq O(4)$* . Ann. of Math. (2) 117 (1983), no. 3, 553–607.
64. A. Hatcher, *Algebraic topology*. Cambridge University Press, Cambridge, 2002. xii+544 pp. ISBN: 0-521-79160-X; 0-521-79540-0
65. A. Hatcher. *Spaces of knots*. arXiv:math/9909095.
66. A. Hatcher, Notes on Basic 3-Manifold Topology, [<https://pi.math.cornell.edu/hatcher/3M/3M.pdf>].
67. M. Hirsch. *Immersions of manifolds*. Trans. Amer. Math. Soc., 93 (1959), 242–276.
68. L.Hsu, *Calculus of Variations via the Griffiths formalism*, J. Diff. Geom. 36 (1992), 551–589.
69. K. Igusa. *Higher singularities of smooth functions are unnecessary*. Ann. of Math. (2) 119 (1984), no. 1, 1–58.
70. D. Jänichen. *The homotopy type of the space of contact structures over the 3-sphere*. PhD. Thesis, Mathematisches Institut Justus-Liebig-Universität Gießen (2018)

71. T. Kálmán. *Contact homology and one parameter families of Legendrian knots*. Geometry & Topology 9 (2005): 2013–2078.
72. T. Kálmán. *Braid-positive Legendrian links*. Int. Math. Res. Not. 2006, Art. ID 14874.
73. S. Lang, Fundamentals of differential geometry. Graduate Texts in Mathematics, 191. Springer-Verlag, New York, 1999. xviii+535 pp. ISBN: 0-387-98593-X 53-01 (58-01).
74. J. Martínez-Aguinaga. *Existence and classification of maximal growth distributions*. In preparation.
75. L. Ng. *Computable Legendrian invariants*. Topology 42 (2003), no. 1, 55–82.
76. J. Martínez-Aguinaga, A. del Pino. *Convex integration with avoidance and hyperbolic $(4,6)$ distributions*. arXiv:2112.14632.
77. J. Martínez-Aguinaga, A. del Pino. *Classification of tangent and transverse knots in bracket-generating distributions*. arXiv:2210.00582
78. P. Massot, M. Theillière. *Holonomic approximation through convex integration*, arXiv:2107.04344
79. M. Mauhart, P. W. Michor. *Commutators of flows and fields*. Archivum Mathematicum (Brno) 28 (1992), 228–236.
80. J. Milnor. Morse theory. Based on lecture notes by M. Spivak and R. Wells. Annals of Mathematics Studies, No. 51 Princeton University Press, Princeton, N.J. 1963 vi+153 pp. 57.50 (53.72).
81. R. Montgomery. *Engel deformations and contact structures*. Northern California Symplectic Geometry Seminar, 103–117, Amer. Math. Soc. Transl. Ser. 2, 196, Amer. Math. Soc., Providence, RI, 1999.
82. R. Montgomery. *A tour of subriemannian geometries, their geodesics and applications*. Vol. 91. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2002.
83. D. McDuff. *Applications of convex integration to symplectic and contact geometry*. Ann. Inst. Fourier 37 (1987), 107–133.
84. E. Murphy. *Loose Legendrian Embeddings in high dimensional contact manifolds*, arXiv:1201.2245.
85. J. Nash. C^1 isometric imbeddings. Annals of Mathematics. Second Series 60 (1954), 383–396.
86. N. Pia. *On the dynamics of some vector fields tangent to non-integrable plane fields*. Journal of Symplectic Geometry 19(2) (2021), 399–412.
87. Á. del Pino, *On the classification of prolongations up to Engel homotopy*. Proc. Amer. Math. Soc. 146 (2018) 891–907.
88. Á. del Pino, *Topological aspects in the study of tangent distributions*. Textos de Matemática. Série B [Texts in Mathematics. Series B], 48. Universidade de Coimbra, Departamento de Matemática, Coimbra, 2019.
89. Á. del Pino, T. Shin *Microflexibility and local integrability of horizontal curves*, arXiv:2009.14518.
90. Á. del Pino, F. Presas. *Flexibility for tangent and transverse immersions in Engel manifolds*. Rev. Mat. Comp 32(1) (2019), 215–238.
91. Presas, F. *Non-integrable distributions and the h -principle*. Eur. Math. Soc. Newsl. 99, 18–26 (2016).
92. Á. del Pino, T. Vogel. *The Engel-Lutz twist and overtwisted Engel structures*. Geometry & Topology 24(5) (2020), 2471–2546.
93. J. M. Sabloff, M. G. Sullivan. *Families of Legendrian submanifolds via generating families*. Quantum Topol. 7 (2016), no. 4, 639–668.
94. S. Smale. *A classification of immersions of the 2-sphere*. Bulletin Amer. Math. Soc., 63(3) (1957), 196.
95. M.W. Hirsch. *Immersions of manifolds*. Trans. Amer. Math. Soc. 93 1959 242–276.
96. D. Spring. *Convex Integration Theory*. Birkhauser, 1998.
97. L. Starkston. *Arboreal Singularities in Weinstein Skeleta*. Selecta Mathematica, 24 (2018), 4105–4140.
98. P. Ozsváth, Z. Szabó, D. Thurston. *Legendrian knots, transverse knots and combinatorial Floer homology*. Geometry & Topology 12 (2008), no. 2, 941–980.

- 99. M. Theillière. *Corrugation Process and ε -isometric maps*. Actes du séminaire de Théorie spectrale et géométrie 35 (2017-2019), 245–264.
- 100. M. Theillière. *Convex Integration Theory without Integration*. Mathematische Zeitschrift (2021).
- 101. W. Thurston, *The theory of foliations of codimension greater than one* Comment. Math. Helv. 49 (1974).
- 102. V. A. Vassiliev. *Complements of discriminants of smooth maps: topology and applications*. Translations of Mathematical Monographs, vol. 98, American Mathematical Society, Providence, RI, 1992.
- 103. T. Vogel. *Existence of Engel structures*. Ann. of Math. (2) 169:1 (2009) 79–137.
- 104. M.S. Weiss. *Embeddings from the point of view of immersion theory I*. Geometry and Topology 3 (1999), 67–101.