

CONVEX INTEGRATION WITH AVOIDANCE AND HYPERBOLIC (4,6) DISTRIBUTIONS

JAVIER MARTÍNEZ-AGUINAGA AND ÁLVARO DEL PINO

ABSTRACT. This paper tackles the classification, up to homotopy, of tangent distributions satisfying various non-involutivity conditions. All of our results build on Gromov’s convex integration. For completeness, we first prove that the full h -principle holds for step-2 bracket-generating distributions. This follows from classic convex integration, no refinements of the theory are needed. The classification of (3, 5) and (3, 6) distributions follows as a particular case.

We then move on to our main example: A complete h -principle for hyperbolic (4, 6) distributions. Even though the associated differential relation fails to be ample along some principal subspaces, we implement an “avoidance trick” to ensure that these are avoided during convex integration. Using this trick we provide 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.

This so-called “avoidance trick” is part of a general avoidance framework, which is the main contribution of this article. Given any differential relation, the framework attempts to produce an associated object called an “avoidance template”. If this process is successful, we say that the relation is “ample up to avoidance” and we prove that convex integration applies. The example of hyperbolic (4, 6) distributions shows that our framework is capable of addressing differential relations beyond the applicability of classic convex integration.

CONTENTS

1. Introduction	1
2. Jets and relations	7
3. Ampleness and convex integration	11
4. Avoidance	17
5. Proof of the main Theorem	22
6. An example: Exact, linearly-independent differential forms	30
7. Relations involving functions	35
8. h -Principle for step-2 distributions	36
9. Maximal non-involutivity	39
10. h -Principle for hyperbolic (4, 6)-distributions	41
References	48

1. INTRODUCTION

1.1. **An overview of convex integration.** Convex integration appeared first in the work of J. Nash on C^1 isometric immersions/embeddings [29]. Broadly speaking, the idea is that a short immersion can be corrected, one codirection at a time, by introducing oscillations that increase its length. This

2020 *Mathematics Subject Classification.* Primary: 58A30. Secondary: 58A17, 57R57.

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 [21], 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:

1.1.1. *Errors.* 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.

1.1.2. *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).

1.1.3. *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.

1.1.4. *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 [17, p. 171], Y. Eliashberg and N. Mishachev posed the following question:

¹The crucial observation of Gromov is that the arguments presented in [29] and [33, 24] can be understood as integration processes. Recent work of M. Theilli ere [36, 37] 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.

Question 1.1. *Is there a (geometrically meaningful) differential relation \mathcal{R} that is ample in coordinate directions but not ample?*

The present article 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.

1.2. Statement of the main result. This article 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 4.2). Our main theorem reads:

Theorem 1.2. *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 5.1, Section 4. 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 1.1:

Corollary 1.3. *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 1.8 below, which is our main application. To get there and introduce the rest of our applications, we go into the theory of tangent distributions.

1.3. An overview of tangent distributions. A distribution ξ on a manifold M is a subbundle of TM . 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 a sequence of modules:

$$\Gamma^1(\xi) \subset \Gamma^2(\xi) \subset \Gamma^3(\xi) \subset \dots$$

$$\Gamma^1(\xi) := \Gamma(\xi), \quad \Gamma^{i+1}(\xi) := [\Gamma^1(\xi), \Gamma^i(\xi)],$$

that we call the *Lie flag*. 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 .

1.3.1. The nilpotentisation. One may then observe that the flag of distributions

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

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 say that ξ **bracket-generates** \mathcal{F} . When $\xi_{i_0} = TM$, we say that ξ is **bracket-generating** and we call the number i_0 the **step**.

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 we call 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 .

1.3.2. *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 .

1.3.3. *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 1.4. *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.

1.3.4. *Two results of Gromov.* We remind the reader that Gromov's method of flexible sheaves [20] 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. 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 1.4 for the case $T\mathcal{F} = TM$, where M is a closed manifold.

1.4. Applications in the study of distributions. We now review what is known about Question 1.4 for various choices of \mathcal{U} . We will explain what the contributions of this article are as we go along. Our main application is Theorem 1.8 below.

1.4.1. h -Principle for step 2. Let M be a smooth manifold. One expects the answer to Question 1.4 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 1.5. *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 [7]. The proof is presented in Section 8 and is, in fact, a routinary application of convex integration.

Remark 1.6. *Theorem 1.5 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, [16, 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.

1.4.2. Maximal non-involutivity. Theorem 1.5 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 [7], even though partial flexibility results do hold [15]. In the last few years we have seen spectacular progress in our understanding of higher-dimensional contact structures [8, 28].

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.

1.4.3. 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 1.4 was answered positively by McDuff [26], proving that they are (topologically) much more flexible than contact structures. However, interesting questions about them from a geometric perspective remain open [30].

1.4.4. Dimensions 3 and 4. In dimension 3, a regular 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 [40, 10, 12, 32] and the classification of their submanifolds [19, 31, 11] but a definite answer to Question 1.4 is still open.

1.4.5. *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 1.5, we have:

Theorem 1.7. *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 [9], which have been classified only in open manifolds [13]. 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.

1.4.6. *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 1.8. *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 10 and it is a consequence of our main result Theorem 1.2. We emphasise that this result requires ampleness up to avoidance and is beyond the scope of classic convex integration.

Remark 1.9. *We conjecture that the answer to Question 1.4 is negative for elliptic (4,6) distributions. In Corollary 10.5 we will show that ampleness fails badly for the differential relation that defines them.*

The remaining cases are: Corank-1 (which are even-contact structures), rank-3 (classified by Theorem 1.7) and rank-2 (the so-called (2,3,5,6) structures, for which nothing is known).

1.4.7. *Where is rigidity?* This article began as our attempt 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 previous discussion 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, 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 (classic) ampleness.
- In even rank $(4l, 4l + 2)$ we see a 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 avoidance.

We intend to address this conjecture in future work.

Remark 1.10. *Recently (and independently from us) J. Adachi has also tackled the classification of non-involutive distributions using (classic) convex integration [2].*

He looks at distributions that are of odd rank and satisfy a certain non-involutivity condition that we call Adachi non-integrability. This notion is weaker than standard maximal non-involutivity (as

defined in the first item above) and, in fact, it is not an intrinsic property of the distribution, but a property of a preferred coframe. In particular, our aforementioned conjecture does not follow from his statements.

On the other hand, Adachi non-integrability is stronger than being of step 2. The two notions coincide for (3,5) and (3,6). In particular, his claims should recover our step-2 result (Theorem 1.5) for odd rank. Theorem 1.7 would then be a particular case.

Our main result about hyperbolic (4,6) distributions (Theorem 1.8) falls outside of the scope of his paper.

1.5. Structure of the paper. In Sections 2 and 3 we review jet spaces and convex integration. Even though the material covered is classic, we hope readers will see it as a helpful starting point into the theory of convex integration. In particular, Section 3 provides a comparison of the various flavours of ampleness that have appeared in the literature.

Sections 4 and 5 introduce avoidance (pre-)templates (Definition 4.1), ampleness up to avoidance (Definition 4.2), and contain the proof of our main result Theorem 1.2.

In Section 6 we explain how avoidance may be used to prove the h -principle for exact, linearly-independent differential forms. We see this as a particularly easy example of how the theory is used and we include it for the sake of exposition. We note that classic convex integration without avoidance can already be used to prove such an h -principle.

In Section 7 we apply avoidance to relations involving functions into \mathbb{R} . We prove that, under mild assumptions, our theory produces empty avoidance pretemplates, so the relations are not ample up to avoidance. The concrete case of functions without critical points is spelled out in Subsection 7.1; note that this relation indeed does not abide by the h -principle, due to Morse theory. Other examples of this type can be found in Lemma 8.8 (contact structures) and Corollary 10.5 (elliptic (4,6) structures).

In Section 8 we study step-2 distributions and we prove Theorems 1.5 and 1.7.

In Section 9 we review the theory of rank-4 distributions in dimension 6. This sets the stage for our main application in Section 10: We show that the relation \mathcal{R}^{hyp} defining hyperbolic (4,6) distributions is ample up to avoidance. The h -principle for \mathcal{R}^{hyp} follows then from Theorem 1.2.

Acknowledgements: The authors are thankful to I. Zelenko for his interest in this project and his willingness to share his insights into the geometry of distributions. They are also grateful to F. Presas for his continued interest and support. During the development of this work the first author was 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”). Similarly, the second author was funded by the NWO grant 016.Veni.192.013; this grant also funded the visits of the first author to Utrecht.

2. JETS AND RELATIONS

We now introduce jet spaces (Subsection 2.1). We put particular emphasis in the geometry behind principal directions and subspaces (Subsection 2.2). This will allow us to discuss differential relations and over-relations and study them “one direction at a time” (Subsection 2.3). In Subsection 2.4 we introduce foliated analogues of these concepts; these will be used to phrase parametric statements in later Sections.

2.1. 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.

2.1.1. *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 $\text{Sym}^r(V, W)$.

2.2. **Principal subspaces.** The following notion formalises the idea of two r -jets that agree except along a pure derivative of order r :

Definition 2.1. *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 2.2. *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 $\perp(\lambda, r) := \perp(\ker(\lambda), r)$. 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 $\text{Pr}_{\lambda, j^r f}$.

2.2.1. *Principal subspaces in coordinates.* As in Subsubsection 2.1.1, we use vector spaces V and W as the local models for M and the fibre of $X \rightarrow M$, respectively.

Lemma 2.3. *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 2.3 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 2.4. *$\text{Sym}^r(V, W)$ 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 1.




FIGURE 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 $-\alpha^{\otimes 2}$.

2.2.2. *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 2.5. 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 2.4, 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 2.

2.3. **Over-relations.** We are interested in finding and classifying solutions of differential relations. More generally, we define:

Definition 2.6. 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.




FIGURE 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 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).

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 2.7. *More generally, Gromov defines open over-relations as those where $\iota_{\mathcal{R}}$ is a microfibration [22, p. 175] but not necessarily submersive. Such generality is unnecessary for our purposes.*

A trivial but key observation is the following:

Lemma 2.8. *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 2.9. *Given an over-relation $\iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^r(X)$ of order r , the projection*

$$\pi_{\mathcal{R}'}^r \circ \iota_{\mathcal{R}} : \mathcal{R} \rightarrow J^{r'}(X)$$

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 $\pi_{\mathcal{R}'}^r$ 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 .

2.4. 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 2.10. *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 $\mathcal{S}|_L \rightarrow J^r(Y|_L)$.

2.4.1. *Parametric lifts of over-relations.* The most important example of foliated over-relation is the following:

Definition 2.11. *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 2.12. *The parametric lift of an open over-relation is open.*

2.4.2. *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 2.13. *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 2.14. *The non-foliated preimage of an open, foliated over-relation is open.*

3. 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 [22]. We also borrow from Spring [34] and Eliashberg-Mishachev [17].

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 3.2. Ampleness along principal frames (often called ampleness in coordinate directions) is explained in Subsection 3.3. Ampleness in the sense of convex hull extensions appears in Subsection 3.4.

We then compare them in Subsection 3.5. This will clarify how ampleness up to avoidance (to appear in Section 4) fits within this greater context.

3.1. 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 3.1. *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 3.2. *A stratified subset $\Sigma \subset \mathbb{R}^n$ of codimension at least 2 is said to be thin. Its complement is ample. See Figure 3.*




FIGURE 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 3.3. *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 10.

3.2. 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 3.4. 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 $(\pi_{r'}^r \circ \iota_{\mathcal{R}})_{r'=1, \dots, r}$ 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 3.5. *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 [21, Corollary 1.3.2]. The statement for all orders appeared later in [22, Section 2.4, p. 180]. The first order case is treated as well in [17, Part 4].

3.2.1. 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 3.4.

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 3.6. *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 3.7. *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 3.6 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 3.5 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 $\mathcal{R} \times K$.

3.3. 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 3.8. *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 3.9. 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 C of principal frames of order r whose supports cover M .

A principal direction/subspace defined by a hyperplane in C/C will be called a C/C -principal direction/subspace.

The second flavour of ampleness reads:

Definition 3.10. 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 C** if \mathcal{R} is ample along all 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 $\{\tau_i = \ker(dx_i)\}$, 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 [21, Definition 1.2.6] and [17, p. 167].

Theorem 3.11. 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 [21, Theorem 1.3.1]; it appears in [17, p. 172] as well. For arbitrary order, it follows from [22, p. 179, Principal Stability Theorem C]. An alternate implementation for arbitrary order appeared in the master thesis [14]; it avoids convex hull extensions and adapts instead the idea from [21].

3.3.1. *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 3.10 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.

3.4. **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 3.12. 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 3.13. Definition 3.12 differs slightly from the definition of convex hull extension given in [22, 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 3.14. *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 2.3 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} . \square

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 3.15. *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 3.16. *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 $(F, \dots, j^T(\pi_0^r \circ \iota_{\mathcal{R}} \circ F))$.*

*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 3.17. *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 3.18. *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 [34, p. 123, Theorem 8.4]. The statement, without integrability, was already implicit in [22, p. 179, Principal Stability Theorem C]. The integrability hypothesis is restrictive; we discuss it further in the next Subsection.

3.5. 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 3.11 implies Theorem 3.5. Furthermore, Theorem 3.18 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 3.18, without integrability, is contained in Gromov's text implicitly; it can be deduced from [22, p. 179, Principal Stability Theorem C]. In [34], 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 [22, 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 5.6) saying that there is a weak equivalence between the space of sections and the subspace of integrable ones. The authors of the present paper have not checked this claim in detail. In fact, it is not important for our results:

Remark 3.19. *None of the results from this paper rely on Definition 3.17 or Theorem 3.18. 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 3.11). Nonetheless, in Corollary 5.5 we prove that a relation that is ample up to avoidance is ample in the sense of integrable convex hull extensions.*

3.5.1. 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 3.20. *In [25], P. Massot and M. Theilli re 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 8. 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.

3.5.2. 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.

3.5.3. 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 4.2.1.

4. 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 1.2, 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 4.2 and 4.3. These definitions require us to introduce some auxiliary notation about configurations of hyperplanes; this is done in Subsection 4.1. In Subsection 4.5 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 4.6.

4.1. 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.

4.1.1. 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

$$\overline{\text{H-Conf}}(V) := \coprod_{a=0}^{\infty} \overline{\text{H-Conf}}_a(V), \quad \overline{\text{H-Conf}}^*(V) := \coprod_{a=0}^{\infty} \overline{\text{H-Conf}}_a^*(V),$$

where $\overline{\text{H-Conf}}_0(V)$ and $\overline{\text{H-Conf}}_0^*(V)$ are the singleton set $\{\emptyset\}$.

4.1.2. *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.

4.2. **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 4.1. *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 4.2. *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.

4.2.1. *Computability of avoidance.* We stated in Subsection 3.5 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 4.5. However, the reader should just think of these as rough guidelines. In practice (for instance, in the proof of Theorem 1.8), one needs to make adjustments in order to produce a template. Still, the adjustments that need to be made are somewhat standard; see Remark 10.8.

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 4.4. The parametric version of Theorem 1.2 will follow then from the non-parametric one.

4.3. Lifted avoidance templates. Definition 4.1 is intuitive conceptually but, in practice (see the proofs of Propositions 5.3 and 5.4), it is often more convenient to deal with the following notion:

Definition 4.3. *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 4.4. *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 $\overline{\text{H-Conf}}_m(T_pM)$.

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. \square

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.

4.4. Foliated templates. We will prove in Section 5 that the parametric analogue of Theorem 1.2 follows from Theorem 1.2 itself. Compare this to Theorem 3.5 and Corollary 3.7. 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 $\iota_S : \mathcal{S} \rightarrow J^r(Y; \mathcal{F})$. 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 4.5. *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 4.5:

Lemma 4.6. *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 4.7. *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 3.6. \square

4.5. Removing processes. The most straightforward way of producing templates consists of iteratively removing those principal subspaces along which the relation is not ample.

Definition 4.8. *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, $\text{Avoid}^l(\mathcal{R})$ 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 4.9. *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 4.10. *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. \square

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.

4.5.1. *Thinning.* In applications, the following more restrictive notion can also be useful.

Definition 4.11. *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 4.12. *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 4.13. *$\text{Thin}^l(\mathcal{R})$ 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 4.9 and 4.12. In fact, this will be necessary for our main application Theorem 1.8; see Section 10.

4.6. **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 4.14. *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 $\mathcal{R}_{\tau, z_0}^c$ 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 $\text{Avoid}^k(\mathcal{R})(\Xi)$. \square

A couple of concrete instances where Lemma 4.14 applies are the relation defining functions without critical points (Subsection 7.1) and the relation defining contact structures (Lemma 8.8).

Exactly the same reasoning shows:

Lemma 4.15. *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. PROOF OF THE MAIN THEOREM

In this Section we tackle the proof of Theorem 1.2. We restate it now in a slightly more general form that applies to over-relations:

Theorem 5.1. *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 5.3.
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.

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 5.2.

5.1. 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 5.2. *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 $\iota_{\mathcal{R}}$.

5.2. Proof of Theorem 1.2. Before we get to the proof we introduce two key auxiliary results. The first one states that avoidance relations are open and ample:

Proposition 5.3. *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. *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 5.3.

Proof of Proposition 5.3. Recall the three pointwise Properties in the definition of a template (Definition 4.1).

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 4.1, $\overline{\mathcal{A}}(C(p))$ 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. \square

Proof of Theorems 1.2 and 5.1 assuming Proposition 5.4. 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 2.4.1.

Since \mathcal{R} is ample up to avoidance we can apply Lemmas 4.6 and 4.7 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 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 3.11). 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 3.11, 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). \square

Corollary 5.5. *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 1.2 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 3.17 applies to \mathcal{R} . \square

5.3. Jiggling for principal covers. In this Subsection we prove Proposition 5.4, completing the proof of Theorems 1.2 and 5.1. 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 [38]. For completeness, we recall it in Subsection 5.4; 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 5.6. *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 [35, 6, 4, 5].

Recall the setup of Proposition 5.4: 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.

5.3.1. 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 $U \cap V$.

5.3.2. 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 .

5.3.3. *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 4.

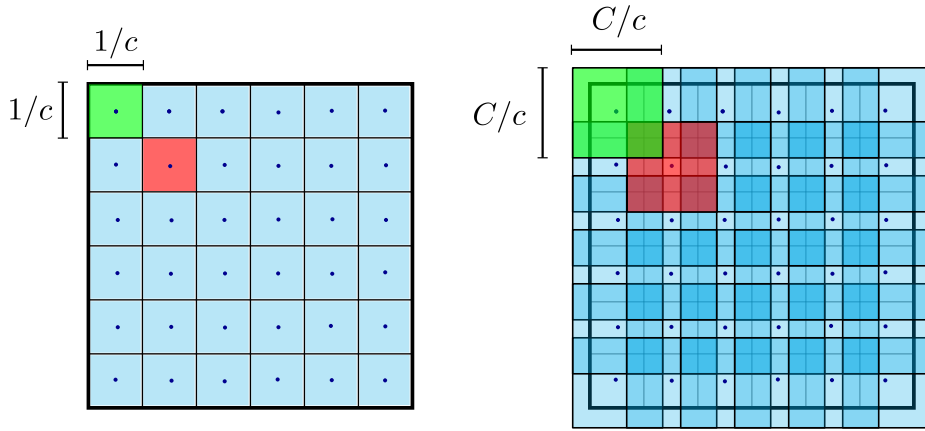


FIGURE 4. The big cube on the left represents the image $[-1, 1]^n$ of a chart $(U, \phi) \in \mathcal{U}$, before the C -dilation is introduced. We show how it is subdivided into smaller cubes, as well as the (image by ϕ of the) marked point of each smaller cube. Two of the children are marked in green 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.

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.

5.3.4. *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 5.7. *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. \square

5.3.5. *Colouring.*

Lemma 5.8. *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} . \square

5.3.6. *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 $\overline{\mathbb{H}\text{-Conf}}(TU)$ as $U \times \overline{\mathbb{H}\text{-Conf}}(T_uU)$. This produces again a projection

$$\overline{\mathbb{H}\text{-Conf}}(TU) \rightarrow \overline{\mathbb{H}\text{-Conf}}(T_uU).$$

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 $\overline{\mathbb{H}\text{-Conf}}$, we have that:

Lemma 5.9. *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 $\overline{\mathbb{H}\text{-Conf}}_j(T_uU)$ and $\overline{\mathbb{H}\text{-Conf}}_j(T_xU)$.

5.3.7. *The diameter of a hyperplane field.* Given $(U, \phi) \in \mathcal{V}(c)$, we look at the principal directions coming from neighbourhouring 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 5.10. *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. \square

5.3.8. *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 $U \rightarrow \overline{\text{H-Conf}}_j(T_uU)$.

Lemma 5.11. *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 5.10 due to the finiteness of j . \square

5.3.9. *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:

$$\Sigma_U := \pi_U(\overline{\mathcal{A}}(F)^c \cap \overline{\text{H-Conf}}(TU)) \subset \overline{\text{H-Conf}}(T_uU),$$

where the superscript c denotes taking complement.

Lemma 5.12. *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 $\overline{\text{H-Conf}}_j(T_uU)$ 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 $\overline{\text{H-Conf}}_j(T_xM)$ contains a δ_x -ball fully contained in $\overline{\mathcal{A}}(F)$; see Figure 5. Indeed: suppose B is a $\varepsilon/2$ -ball in $\overline{\text{H-Conf}}_j(T_xM)$. Since $\overline{\mathcal{A}}(F)$ is fibrewise dense, there exists $\Xi \in \overline{\mathcal{A}}(F) \cap \overline{\text{H-Conf}}_j(T_xM)$. By openness of $\overline{\mathcal{A}}(F)$ there is a δ_Ξ -ball $D \subset \overline{\mathcal{A}}(F) \cap \overline{\text{H-Conf}}_j(T_xM)$ centered at Ξ and contained in B . We now use the compactness of $\overline{\text{H-Conf}}_j(T_xM)$ 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 $\overline{\text{H-Conf}}_j(T_xM)$ 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 $\overline{\text{H-Conf}}_j(TP) = P \times \overline{\text{H-Conf}}_j(T_pP)$, 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 $\overline{\text{H-Conf}}_j(T_pP)$ (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 $\{D_i^p\}$ 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 5.9.

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 $\overline{\text{H-Conf}}_j(T_uM)$, with u the marked point of U , is contained in a $2\rho \times \varepsilon$ -polydisc $B_i^{p_l^P}$ centered at $\overline{\text{H-Conf}}_j(T_{p_l^P}M)$. Then, upon




FIGURE 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.

projecting with π_U , the corresponding $2\rho \times \delta$ -polydisc $D_i^{p_i}$ provides the claimed δ -ball disjoint from Σ_U . See Figure 6. \square

5.3.10. *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 5.12. 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:

$$(1) \quad \varepsilon_i > 2.A \sum_{j>i} \varepsilon_j.$$

The successive applications of Lemma 5.12 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 5.3.8. According to Lemma 5.11, 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 5.12 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 $\overline{\mathcal{A}}(F)$.

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 5.12 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 .




FIGURE 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 .

- 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 1, is not destroyed in later steps thanks to the first item. After d_2 steps, the proofs of Proposition 5.4, Theorem 1.2, and Theorem 5.1 are complete. \square

5.4. Detour: Thurston's jiggling. We invite the reader to compare the upcoming discussion to the proof of Proposition 5.4. 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 [38]. 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 5.13. *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. 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. \square

6. AN EXAMPLE: EXACT, LINEARLY-INDEPENDENT DIFFERENTIAL FORMS

In this section we introduce a differential relation \mathcal{R} for which the thinning process (Subsection 4.5.1) 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 6.1 we recall some notation, involving the exterior differential, that will also be helpful in later sections. In Subsection 6.2 we formalise our claim about the thinning of \mathcal{R} , which we then prove in Subsection 6.3.

In Subsection 6.4 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.

6.1. Intermezzo: The symbol of the exterior differential. Let N be a manifold. We focus on the bundle $\wedge^k T^*N$ of k -forms. The corresponding space of 1-jets $J^1(\wedge^k T^*N)$ is an affine fibration

$$\pi_0 : J^1(\wedge^k T^*N) \longrightarrow \wedge^k T^*N$$

whose model vector bundle is $\text{Hom}(TN, \wedge^k T^*N) \cong T^*N \otimes \wedge^k T^*N$.

It follows that, given a formal datum $F \in J^1(\wedge^k T^*N)$ 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:

$$\text{Pr}_{\lambda, F} := \{F + (0, \lambda \otimes \beta) \mid \beta \in \wedge^k T_p^*N\}.$$

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:

$$d : J^1(\wedge^k T^*N) \longrightarrow \wedge^{k+1} T^*N.$$

Given $F \in J^1(\wedge^k T^*N)$, 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**:

$$\begin{aligned} \text{id} \oplus d : J^1(\wedge^k T^*N) &\longrightarrow \wedge^k T^*N \oplus \wedge^{k+1} T^*N, \\ F &\mapsto (\pi_0(F), dF). \end{aligned}$$

6.2. 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 6.1. *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 1.2, the h -principle holds for \mathcal{R} .

The proof of Proposition 6.1 requires an analysis of the structure of \mathcal{R} along principal subspaces (Subsection 6.3.1). This will then allow us to describe $\text{Thin}(\mathcal{R})$ explicitly and deduce that it is an avoidance template (Subsection 6.3.3).

6.3. 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 6.4 below).

6.3.1. The thinning step. The following criterion will allow us to compute $\text{Thin}(\mathcal{R})$:

Lemma 6.2. *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 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:

$$\begin{aligned} \tau : \quad d\text{Pr}_{\lambda,F} &\longrightarrow \ker(\lambda)^* \times \ker(\lambda)^* \cong \mathcal{M}_{2 \times 2} \\ (dF_i + \lambda \wedge \beta_i)_{i=1,2} &\mapsto ((\nu_i + \beta_i)|_{\ker(\lambda)})_{i=1,2}. \end{aligned}$$

The subset $\mathcal{R}_{\lambda,F}$ maps precisely to the linearly independent pairs $((\nu_i + \beta_i)|_{\ker(\lambda)})_{i=1,2}$. 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 not thin, because the complement of $\text{GL}(2)$ (the zero set of the determinant, which has codimension 1) is not. We will see in Subsection 6.4 that $\text{GL}(2)$ is nonetheless ample, and so is $\mathcal{R}_{\lambda,F}$.




FIGURE 7. The first possibility in the proof of Lemma 6.2: Both dF_i are contained in L .

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 8), the complement of $\mathcal{R}_{\lambda,F}$ is of codimension-2. In the second case (Figure 9), the complement of $\mathcal{R}_{\lambda,F}$ is the set $\{dF_2 = 0\}$, which is also of codimension-2. This completes the claim. \square

Recall the avoidance template notation from Subsection 4.2. We remark that the thinning process was defined as a pointwise process and Lemma 6.2 indeed applies to each point $p \in M$ individually. An immediate consequence is:

Corollary 6.3. *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 $\mathcal{R} \setminus (\bigcup_{\lambda \in \Xi} \Sigma(\lambda))$.

6.3.2. *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 6.4. *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 6.2. 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.




FIGURE 8. The second possibility in the proof of Lemma 6.2: None of the dF_i are contained in L .




FIGURE 9. The third possibility in the proof of Lemma 6.2: One of the dF_i is contained in L but the other one is not.

- 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. \square

6.3.3. *Completing the proof.* The previous Lemmas allow us to conclude:

Proof of Proposition 6.1. We want to show that $\text{Thin}(\mathcal{R})$ is a 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 6.3, 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 6.2, 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 6.4 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. \square

6.4. **Linear algebra.** As we stated in the proof of Lemma 6.2, 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 6.5. *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. \square

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. $\text{GL}^-(n)$). \square

We conclude:

Corollary 6.6. *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. \square

Corollary 6.7. *The relation \mathcal{R} treated in this Section is ample in all principal directions and therefore abides by the h -principle.*

7. 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 7.1 we prove that the relation defining functions without critical points fails to be ample up to avoidance. In Subsection 7.2 we generalise this to arbitrary differential relations $\mathcal{R} \subset J^r(X)$ that satisfy a mild non-triviality condition.

Remark 7.1. *Let us compare these claims with [21, 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 7.2. *We observe that the applicability realms of Vassiliev's h -principle [39] 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.*

7.1. 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 4.14 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.

7.2. The general case. It follows immediately from Lemma 4.14 that:

Proposition 7.3. *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.

8. h -PRINCIPLE FOR STEP-2 DISTRIBUTIONS

In this Section we will prove the h -principle for step-2 distributions (Theorem 1.5) and its corollary about the classification of (3,5) and (3,6) distributions (Theorem 1.7). The proof can be found in Subsection 8.4. 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 8.1. This will allow us, in Subsection 8.2, to translate our statements about distributions to statements about their annihilating forms.

In Subsection 8.5, for completeness, we look at convex integration in the setting of (even-)contact structures, following the work of McDuff [26].

8.1. 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 1.3.

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 paper, 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 1.3.1 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 8.1. *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.

8.2. 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 $\mathcal{R}^{\text{step}2} \subset J^1(\text{Gr}_k(TM))$ 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 8.2. *The differential relation $\mathcal{R}^{\text{step}2} \subset J^1(\text{Gr}_k(TM))$ 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 $\mathcal{R}^{\text{step}2}$. 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 8.2 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 8.3. *A jet $\tilde{F} \in J^1(\overline{\oplus^{n-k} T^*M})$ 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 $\mathcal{S}^{\text{step}2} \subset J^1(\overline{\oplus^{n-k} T^*M})$.

The second condition is, according to Lemma 8.1, indeed equivalent to Ω^F being an epimorphism. It follows that:

Lemma 8.4. $\mathcal{S}^{\text{step}2}$ is the preimage of $\mathcal{R}^{\text{step}2}$ under $j^1\pi$. In particular, $\mathcal{S}^{\text{step}2}$ fibres affinely over $\mathcal{R}^{\text{step}2}$.

8.3. 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^0F . It may very well happen that the annihilator $(j^0F)^\perp$ is not trivial as a bundle. This would imply that F cannot be lifted to $\tilde{F} \in J^1(\overline{\oplus^{n-k} T^*M})$. 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 [17, 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 [3] 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^0G is now trivial, and thus a lift to $\mathcal{S}^{\text{step}2}$ exists. We conclude that:

Lemma 8.5. *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.

8.4. h -Principle for step-2 distributions.

Proof of Theorem 1.5. According to Lemma 8.5, 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 $\xi \cap \ker(\lambda)$ is important. The forms $(dF_i|_{\ker(\lambda) \cap \xi})_{i=1}^{n-k}$ span a subspace, say, of dimension l . Up to a change of basis, we may then assume that

$$dF_i|_{\ker(\lambda) \cap \xi} = 0, \quad \text{for all } i = l + 1, \dots, n - k.$$

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 $\{(\beta_i + \nu_i)|_{\ker(\lambda) \cap \xi}\}_{i=l+1}^{n-k}$ are linearly independent. This means that the ampleness of $\mathcal{S}^{\text{step}2}$ 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 6.6, concluding the proof. \square

Theorem 1.5 proves that step-2 distributions are flexible as long as we do not impose any further non-degeneracy constraints. Nonetheless, according to Theorem 1.7, 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 8.6. *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 $\mathcal{S}^{\text{step}2}$ 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 6.

8.5. The contact and even-contact cases. As an appetiser for our study of maximally non-involutive distributions in Sections 9 and 10, 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 [7]), whereas the latter is thin along all principal directions (as proven by D. McDuff in [26]).

8.5.1. *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:

$$\begin{aligned} \Gamma: \quad J^1(T^*M) &\longrightarrow T^*M \oplus \wedge^2 T^*M &\longrightarrow \wedge^{2\lfloor \frac{n-1}{2} \rfloor + 1} T^*M \\ F &\longmapsto (j^0 F, dF) &\longmapsto j^0 F \wedge (dF)^{\lfloor \frac{n-1}{2} \rfloor}, \end{aligned}$$

where the first arrow is the extended symbol of the exterior differential. The Pfaffian measures whether the formal curvature $dF|_{\ker(j^0 F)}$ has maximal rank.

Definition 8.7. *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 8.3), so the h -principle for (even-)contact structures is equivalent to the h -principle for $\mathcal{R}^{\text{cont}}$. We study its ampleness next.

8.5.2. *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:

$$d\text{Pr}_{\lambda, F} := \{(j^0 F, dF + \lambda \wedge \beta) \mid \beta \in T_p^*M\} \subset T_p^*M \oplus \wedge^2 T_p^*M.$$

We write $\xi = \ker(j^0 F)$.

A point $\tilde{F} = F + (0, \lambda \otimes \beta) \in \text{Pr}_{\lambda, F}$ is formally (even-)contact if and only if:

$$\Gamma(\tilde{F}) = n j^0 F \wedge (dF)^{\lfloor \frac{n-1}{2} \rfloor - 1} \wedge (dF + \lambda \wedge \beta) \neq 0.$$

Since F was a formal solution, there are four possible situations:

1. λ is proportional to $j^0 F$.
2. λ is not proportional to $j^0 F$ and n is odd. Then dF has a 1-dimensional kernel L when restricted to $\xi \cap \ker(\lambda)$.
- 3a. λ is not proportional to $j^0 F$, n is even, and $\ker(\lambda)$ contains the 1-dimensional kernel of $dF|_{\xi}$. Then dF has a 2-dimensional kernel L' when restricted to $\xi \cap \ker(\lambda)$.
- 3b. λ is not proportional to $j^0 F$, n is even, and $\ker(\lambda)$ is transverse to the 1-dimensional kernel of $dF|_{\xi}$. Then dF is non-degenerate when restricted to $\xi \cap \ker(\lambda)$.

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 8.8. *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 8.9. *The differential relation describing even-contact structures has thin complement.*

The h -principle for even-contact structures follows then from classic convex integration Theorem 3.5.

9. 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 8.5, we will use the Pfaffian map to capture this in an algebraic manner (Subsection 9.1). In Subsection 9.2 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 [27].

9.1. 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 8.5, we can take the highest potentially-non-trivial power of the curvatures using the map:

$$\begin{aligned} p : \Lambda^2 \xi^* &\longrightarrow \Lambda^{2\lfloor \frac{k}{2} \rfloor} \xi \\ \omega &\longmapsto \omega^{\lfloor \frac{k}{2} \rfloor}. \end{aligned}$$

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 9.1. A 2-form $\omega \in \Lambda^2 \xi^*$ 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 $\Lambda^2 \xi^*$ if k is even and codimension k if k is odd. Then we define:

Definition 9.2. The composition $p \circ d^\xi$ is called the **Pfaffian**:

$$\begin{aligned} \text{Pf} : \xi^\perp &\longrightarrow \Lambda^2 \xi^* \longrightarrow \Lambda^{2\lfloor \frac{k}{2} \rfloor} \xi^* \\ \alpha &\longmapsto d\alpha|_\xi \longmapsto (d\alpha|_\xi)^{\lfloor \frac{k}{2} \rfloor}. \end{aligned}$$

We are interested in the Diff-invariant non-degeneracy condition:

Definition 9.3. 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 article. We now focus on rank-4 distributions.

9.2. 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 $\Lambda^2 \xi^*$ and study its signature. The signature does not depend on the choice of volume form:

Lemma 9.4. 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 $\Lambda^+(\xi^*)$ and the space of anti self-dual forms $\Lambda^-(\xi^*)$ as follows:

$$\begin{aligned} \Lambda^+(\xi^*) &= \langle a_1 = \beta_1 \wedge \beta_2 + \beta_3 \wedge \beta_4, a_2 = \beta_1 \wedge \beta_3 + \beta_4 \wedge \beta_2, a_3 = \beta_1 \wedge \beta_4 + \beta_2 \wedge \beta_3 \rangle \subset \Lambda^2 \xi^*. \\ \Lambda^-(\xi^*) &= \langle b_1 = \beta_1 \wedge \beta_2 - \beta_3 \wedge \beta_4, b_2 = \beta_1 \wedge \beta_3 - \beta_4 \wedge \beta_2, b_3 = \beta_1 \wedge \beta_4 - \beta_2 \wedge \beta_3 \rangle \subset \Lambda^2 \xi^*. \end{aligned}$$

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. \square

Since ξ is bracket-generating, the exterior differential d^ξ maps ξ^\perp injectively into $\Lambda^2 \xi^*$. We can then talk about the signature of p restricted to the image. This is equivalent to:

Definition 9.5. The signature of a distribution ξ is the signature of the quadratic form $\text{Pf} : \xi^\perp \rightarrow \Lambda^2(\xi^*) \cong \mathbb{R}$.

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 9.6. 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 $d\xi$ 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 9.7. A maximally non-involutive (4,6)-distribution is

- elliptic or fat* if the signature is definite.
- 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.

9.3. Formal maximally non-involutive 4-distributions. As in Section 8, it is more convenient not to work with the distribution itself but with its annihilating forms. We define:

Definition 9.8. 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^0F_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^0F_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 8.3, one can use holonomic approximation to reduce to the case of a ball, proving that:

Lemma 9.9. 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 paper deals with the construction of an avoidance template for \mathcal{S}^{hyp} .

10. h -PRINCIPLE FOR HYPERBOLIC (4,6)-DISTRIBUTIONS

In this Section we tackle the proof of Theorem 1.8. According to Lemma 9.9, 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 1.2 we see that we just need to construct an avoidance template for \mathcal{S}^{hyp} .

10.1. First avoidance step. As we advanced in Subsection 1.2, \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.

10.1.1. 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

$$(2) \quad j^0F_1 \wedge j^0F_2 \wedge \lambda \wedge dF_1 \quad \text{and} \quad j^0F_1 \wedge j^0F_2 \wedge \lambda \wedge dF_2 \quad \text{are linearly dependent.}$$

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 10.1. 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 2 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 $F \in \mathcal{S}^{\text{hyp}}$. \square

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 $dF|_{\xi \cap \ker(\lambda)}$ are linearly-dependent.

10.1.2. *Main statement.* We claim that $\Sigma^{(1)}$ is precisely the set to be removed in order to carry out the first avoidance step.

Proposition 10.2. *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 9.8, 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 of Proposition 10.2. 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

$$(3) \quad \begin{pmatrix} g_{11} & g_{12} \\ g_{12} & g_{22} \end{pmatrix}$$

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 :

$$\begin{aligned} \Psi: \quad \xi^* \oplus \xi^* &\longrightarrow \mathbb{R}^3 \\ (\beta_1, \beta_2) &\longmapsto (g_{11}(\beta_1, \beta_2), g_{22}(\beta_1, \beta_2), g_{12}(\beta_1, \beta_2)) \end{aligned}$$

It is now convenient to introduce the determinant, which we see as a quadratic form in the space of symmetric 2-by-2 matrices:

$$\begin{aligned} \det: \quad \mathbb{R}^3 &\longrightarrow \mathbb{R} \\ (x, y, z) &\longmapsto xy - z^2. \end{aligned}$$

We saw back in Example 3.3 that the signature of \det was $(1, 2, 0)$, so its zero set \mathcal{C} is a cone. See Figure 10. 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 10.4 below; for now we focus on hyperbolicity. We want to check ampleness; there are various possibilities, depending on what the image of Ψ is.




FIGURE 10. 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.

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 $dF|_{\xi \cap \ker(\lambda)}$ 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 $(dF_i|_{\xi \cap \ker(\lambda)})_{i=1,2}$ 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 $dF|_{\xi \cap \ker(\lambda)}$ are linearly dependent but not identically zero; i.e. $L \cap K$ is 1-dimensional. Up to changing basis we may assume that $dF_1|_{\xi \cap \ker(\lambda)} = 0$; i.e. dF_1 spans the line $L \cap K$. In this case, $dF_2|_{\xi \cap \ker(\lambda)} \neq 0$. Then $dF_1 \wedge dF_2 = 0$, 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. \square

10.1.3. *Conclusion of the first avoidance step.* We now describe $\text{Avoid}(\mathcal{S}^{\text{hyp}})$ using Proposition 10.2.

Corollary 10.3. *$\text{Avoid}(\mathcal{S}^{\text{hyp}})$ 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^0F . Recall Situations (i.), (ii.) and (iii.) from Proposition 10.2. 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 10.2, 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 10.1 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. \square

10.1.4. *Ampleness fails for (4,6) elliptic distributions.* For completeness, we observe:

Lemma 10.4. *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 10.2. 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. \square

Reasoning as in Corollary 10.3 implies that:

Corollary 10.5. *$\text{Avoid}(\mathcal{S}^{\text{ell}})$ is empty.*

Which is exactly the same situation as in the contact case (Lemma 8.8). This leads us to conjecture that there is no full h -principle for elliptic (4,6) distributions.

10.2. 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 6, 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 an 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.

10.2.1. *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

$$(4) \quad j^0 F_1 \wedge j^0 F_2 \wedge \lambda_1 \wedge \lambda_2 \wedge dF_1 = j^0 F_1 \wedge j^0 F_2 \wedge \lambda_1 \wedge \lambda_2 \wedge dF_2 = 0.$$

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 4 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 10.6. *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 $(dF_i)|_{\xi \cap \ker(\lambda_i) \cap \ker(\lambda_2)}$ are zero. Equivalently, both $(dF_i)|_\xi$ are proportional to $(\lambda_1 \wedge \lambda_2)|_\xi$.

10.2.2. *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 10.7. *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

$$\Phi : \text{Pr}_{\lambda, F} \longrightarrow \wedge^2(\xi \cap \ker(\nu))^* \oplus \wedge^2(\xi \cap \ker(\lambda))^*$$

that sends $\tilde{F} \in \text{Pr}_{\lambda, F}$ to $d\tilde{F}|_{\xi \cap \ker(\nu)}$. Recall that $\tilde{F} \in \Sigma^{(1)}(\nu)$ if and only if the pair $d\tilde{F}|_{\xi \cap \ker(\nu)}$ is linearly dependent. The upcoming argument follows the proof of Lemma 6.2 (Section 6), since the linear dependence problems under consideration are exactly the same.

Write $L \subset \wedge^2(\xi \cap \ker(\nu))^*$ for the subspace of 2-forms proportional to $\lambda|_{\xi \cap \ker(\nu)}$. 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. \square

Remark 10.8. *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 11.*

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 11. 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.

10.2.3. *Conclusion of the second step.* Using Proposition 10.7 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 10.9. *\mathcal{A} is a pretemplate.*

Proof. As we noted in Lemma 10.6, 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. \square

10.3. **End of the proof.** The proof of Theorem 1.8 will be complete once we show that:

Theorem 10.10. \mathcal{A} is a template for \mathcal{S}^{hyp} .

This statement requires the following auxiliary result:

Proposition 10.11. 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 4. The condition $F \notin \Sigma^{(2)}(\nu_1, \nu_2)$ implies that the forms $\nu_i|_{\xi}$ are linearly independent. As such, Equation 4 reads:

$$(dF_1)|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)} = (dF_2)|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)} = 0.$$

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

$$d\tilde{F}|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)} = dF|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)},$$

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 $d\tilde{F}|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)}$ as a pair of 1-forms

$$(dF_i + \lambda \wedge \beta_i)|_{\xi \cap \ker(\nu_1) \cap \ker(\nu_2)}.$$

Equation 4 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. \square

Proof of Theorem 10.10. 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 10.2 to deduce that $\mathcal{S}_{\lambda, F}^{\text{hyp}}$ is ample.
- Using condition $F \notin \Sigma^{(2)}(\lambda, \nu)$, we apply Proposition 10.7 and deduce that each $\Sigma^{(1)}(\nu)$ is thin.
- According to Proposition 10.11, 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 $\text{H-Conf}(T_pM)$. Its complement can be explicitly written down as

$$\mathcal{A}(F)^c \cap \text{H-Conf}(T_pM) = \{\Xi \in \text{H-Conf}(T_pM) \mid F \in \Sigma^{(1)}(\nu_1) \cup \Sigma^{(2)}(\nu_1, \nu_2) \text{ for some } \nu_1 \neq \nu_2 \in \Xi\}.$$

The conditions $\nu_1 \in \Sigma^{(1)}(F)$ or $(\nu_1, \nu_2) \in \Sigma^{(2)}(F)$ are given, respectively, by Equations 2 and 4. Both of them are non-trivial on their entries, since F is in \mathcal{S}^{hyp} ; see Lemmas 10.1 and 10.6. We deduce that $\mathcal{A}(F)^c \cap \text{H-Conf}(T_pM)$ is a finite union of positive-codimension subvarieties, proving the claim. \square

Proof of Theorem 1.8. According to Lemma 9.9, the h -principle for \mathcal{R}^{hyp} reduces to the h -principle for \mathcal{S}^{hyp} . Theorem 1.2 says that the h -principle for \mathcal{S}^{hyp} follows from the existence of an avoidance template. Theorem 10.10 yields such a template \mathcal{A} . \square

REFERENCES

- [1] J. Adachi. *Classification of horizontal loops in standard Engel space*. Int. Math. Res. Not. 2007 no. 3.
- [2] J. Adachi. *Existence and classification of maximally non-integrable distributions of derived length one*. arXiv:2111.01403
- [3] D. Álvarez-Gavela. *Refinements of the holonomic approximation lemma*. Algebraic & Geometric Topology 18:4 (2018), 2265–2303.
- [4] D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization II: Geomorphology of Lagrangian ridges*. arXiv:1912.03439
- [5] D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization III: Positive arborealization of polarized Weinstein manifolds*. arXiv:2011.08962
- [6] D. Álvarez-Gavela, Y. Eliashberg, D. Nadler. *Arborealization I: Stability of arboreal models*. arXiv:2101.04272
- [7] D. Bennequin. *Entrelacements et equations de Pfaff*. Asterisque 107-108 (1983), pp. 87–161.
- [8] 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.
- [9] E. Cartan. *Les systemes de Pfaff a cinq variables et les equations sur derivees partielles du second ordre*. Ann. Ec. Norm. 27 (1910).
- [10] R. Casals, J.L. Pérez, Á. del Pino, F. Presas. *Existence h-principle for Engel structures*. Invent. Math. 210 (2017), 417-451
- [11] R. Casals, Á. del Pino. *Classification of Engel knots*. Math. Ann. 371(1-2) (2018), 391–404.
- [12] R. Casals, Á. del Pino, F. Presas. *Loose Engel structures*. Compos. Math. 156(2) (2020), 412–434.
- [13] S. Dave, S. Haller. *On 5-manifolds admitting rank two distributions of Cartan type*. Trans. Amer. Math. Soc. 371 (2019), 4911–4929.
- [14] J. Denkers. *Semifree and Totally Non-Parallel Immersions*. Master thesis. Utrecht University.
- [15] Y. Eliashberg. *Classification of overtwisted contact structures on 3-manifolds*. Invent. Math. 98 (1989), no. 3, 623–637.
- [16] Y. Eliashberg, D. Kotschick, E. Murphy, T. Vogel. *Engel structures*. AIM Workshop report (2017).
- [17] Y. Eliashberg, N. Mishachev. *Introduction to the h-principle*. Graduate Studies in Mathematics, 48. American Mathematical Society, Providence, RI, 2002.
- [18] G. Frobenius, *Über das Pfaffsche Problem*. J. für Reine und Angew. Math., 82 (1877) 230–315.
- [19] H. Geiges. *Horizontal loops in Engel space*. Math. Ann. Volume 342, Issue 2, (2008) 291–296.
- [20] 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.
- [21] M. Gromov. *Convex integration of partial differential relations*. I, Izv. Akad. Nauk SSSR Ser. Mat., 37(1973), 329–343.
- [22] M. Gromov. *Partial differential relations*. Ergebnisse Math. und ihrer Gren. , Springer-Verlag, Berlin, 1986.
- [23] M. Gromov, *Carnot-Carathéodory spaces seen from within*. Sub-Riemannian geometry, Progr. Math., vol. 144, Birkhäuser, Basel, 1996.
- [24] M. Hirsch. *Immersions of manifolds*. Trans. Amer. Math. Soc., 93 (1959), 242–276.
- [25] P. Massot, M. Theillière. *Holonomic approximation through convex integration*, arXiv:2107.04344
- [26] D. McDuff. *Applications of convex integration to symplectic and contact geometry*. Ann. Inst. Fourier 37 (1987), 107–133.
- [27] R. Montgomery. *A tour of subriemannian geometries, their geodesics and applications*. Vol. 91. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2002.
- [28] E. Murphy. *Loose Legendrian Embeddings in high dimensional contact manifolds*, arXiv:1201.2245.
- [29] J. Nash. *C^1 isometric imbeddings*. Annals of Mathematics. Second Series 60 (1954), 383–396.
- [30] N. Pia. *On the dynamics of some vector fields tangent to non-integrable plane fields*. Journal of Symplectic Geometry 19(2) (2021), 399–412.
- [31] Á. del Pino, F. Presas. *Flexibility for tangent and transverse immersions in Engel manifolds*. Rev. Mat. Comp 32(1) (2019), 215–238.
- [32] Á. del Pino, T. Vogel. *The Engel-Lutz twist and overtwisted Engel structures*. Geometry & Topology 24(5) (2020), 2471–2546.
- [33] S. Smale. *A classification of immersions of the 2-sphere*. Bulletin Amer. Math. Soc., 63(3) (1957), 196.
- [34] D. Spring. *Convex Integration Theory*. Birkhauser, 1998.
- [35] L. Starkston. *Arboreal Singularities in Weinstein Skeleta*. Selecta Mathematica, 24 (2018), 4105–4140.
- [36] M. Theillière. *Corrugation Process and ε -isometric maps*. Actes du séminaire de Théorie spectrale et géométrie 35 (2017-2019), 245–264.
- [37] M. Theillière. *Convex Integration Theory without Integration*. Mathematische Zeitschrift (2021).
- [38] W. Thurston. *The theory of foliations of codimension greater than one*. Comm. Math. Helv. 87 (1974), 214–231.
- [39] V. A. Vassiliev. *Complements of discriminants of smooth maps: topology and applications*. Translations of Mathematical Monographs, vol. 98, American Mathematical Society, Providence, RI, 1992.
- [40] T. Vogel. *Existence of Engel structures*. Ann. of Math. (2) 169 (2009), no. 1, 79–137.

UNIVERSIDAD COMPLUTENSE DE MADRID, DEPARTAMENTO DE ÁLGEBRA, GEOMETRÍA Y TOPOLOGÍA, FACULTAD DE MATEMÁTICAS, AND INSTITUTO DE CIENCIAS MATEMÁTICAS CSIC-UAM-UC3M-UCM, C. NICOLÁS CABRERA, 13-15, 28049 MADRID, SPAIN

Email address: `frmart02@ucm.es`

UTRECHT UNIVERSITY, DEPARTMENT OF MATHEMATICS, BUDAPESTLAAN 6, 3584 UTRECHT, THE NETHERLANDS

Email address: `a.delpinogomez@uu.nl`