Invent. math. (2018) 214:641–737
https://doi.org/10.1007/s00222-018-0811-3

CrossMark

The simplification of singularities of Lagrangian and Legendrian fronts

Daniel Álvarez-Gavela[1]

Received: 6 April 2017 / Accepted: 23 June 2018 / Published online: 3 August 2018

Abstract We establish a full h-principle (C^0-close, relative, parametric) for the simplification of singularities of Lagrangian and Legendrian fronts. More precisely, we prove that if there is no homotopy theoretic obstruction to simplifying the singularities of tangency of a Lagrangian or Legendrian submanifold with respect to an ambient foliation by Lagrangian or Legendrian leaves, then the simplification can be achieved by means of a Hamiltonian isotopy.

Contents

The author was partially supported by NSF Grant DMS-1505910.

✉ Daniel Álvarez-Gavela
dgavela@stanford.edu

1 Stanford University, Stanford, USA

1 Introduction and statement of results

1.1 Panoramic overview

In this paper we establish a general h-principle for the simplification of singularities of Lagrangian and Legendrian fronts. The precise formulation is given in Theorem 1.11 below. Here is a sample corollary of our results, where $\pi : T^*S^n \to S^n$ denotes the cotangent bundle of the standard n-dimensional sphere.

Corollary 1.1 *Let $S \subset T^*S^n$ be any embedded Lagrangian sphere. If n is even, then there exists a compactly supported Hamiltonian isotopy $\varphi_t : T^*S^n \to T^*S^n$ such that the singularities of the projection $\pi|_{\varphi_1(S)} : \varphi_1(S) \to S^n$ consist only of folds. An analogous result holds for even-dimensional Legendrian spheres in the 1-jet space $J^1(S^n, \mathbb{R}) = T^*S^n \times \mathbb{R}$.*

More generally, let $S \subset M$ be any embedded Lagrangian sphere, where (M^{2n}, ω) is a symplectic manifold equipped with a foliation $\mathcal{F}$ by Lagrangian leaves. Denote by $T\mathcal{F}$ the distribution of Lagrangian planes tangent to the foliation $\mathcal{F}$ and let V be the restriction of $T\mathcal{F}$ to S. It is easy to see that a necessary condition for S to be Hamiltonian isotopic to a Lagrangian sphere whose singularities of tangency with respect to $\mathcal{F}$ consist only of folds is that V is stably trivial as a real vector bundle over the sphere. When n is even, our h-principle implies the following converse.

Corollary 1.2 *Suppose that $V = T\mathcal{F}|_S$ is stably trivial as a real vector bundle over the sphere. If n is even, then there exists a compactly supported Hamiltonian isotopy $\varphi_t : M \to M$ such that the singularities of tangency of $\varphi_1(S)$ with respect to the foliation $\mathcal{F}$ consist only of folds. An analogous result holds for even-dimensional Legendrian spheres.*

Remark 1.3 As we will see, the assumption that V is stably trivial is automatically satisfied for all even n such that $n \not\equiv 2 \mod 8$. The simplest example in which more complicated singularities are necessary occurs when $n = 2$ and corresponds to the Hopf bundle on S^2, where in addition to the Σ^{10} folds we find that a Σ^{110} pleat is unavoidable. When n is odd the problem is not as straightforward due to the fact that $\pi_n(U_n) \neq 0$. Nevertheless, we will apply our h-principle to give a necessary and sufficient condition for the simplification of singularities to be possible in terms of the homotopy class of the distribution of Lagrangian planes V.

As another application of our h-principle, we establish that higher singularities are unnecessary for the homotopy theoretic study of the space of Legendrian knots in the standard contact $\mathbb{R}^3$. Before we can state our result we need to set some notation.

Recall that the front projection $\mathbb{R}^3 \to \mathbb{R}^2$ corresponds to the forgetful map $J^1(\mathbb{R}, \mathbb{R}) \to J^0(\mathbb{R}, \mathbb{R})$ where we identify $J^1(\mathbb{R}, \mathbb{R}) = \mathbb{R}^3$ and $J^0(\mathbb{R}, \mathbb{R}) = \mathbb{R}^2$. In coordinates, we have $\mathbb{R}^3 = \mathbb{R}(q) \times \mathbb{R}(p) \times \mathbb{R}(z)$, $\mathbb{R}^2 = \mathbb{R}(q) \times \mathbb{R}(z)$, $\xi_{std} = \ker(dz - pdq)$ and the front projection is the map $(q, p, z) \mapsto (q, z)$. The front of a Legendrian knot $f : S^1 \to \mathbb{R}^3$ is the composition of f with the front projection, which results in a map $S^1 \to \mathbb{R}^2$. Let $\mathcal{L}$ be the space of all (parametrized) Legendrian knots $f : S^1 \to \mathbb{R}^3$ and let $\mathcal{M} \subset \mathcal{L}$ be the subspace consisting of those Legendrian knots whose front only has mild singularities, namely cusps and embryos. A cusp of the front corresponds to a fold type singularity of tangency of f with respect to the foliation given by the fibres of the front projection. An embryo is the instance of birth/death of two cusps and corresponds to the familiar Reidemeister Type I move.

The inclusion $\mathcal{M} \hookrightarrow \mathcal{L}$ is not a homotopy equivalence. Indeed, it is easy to see that $\pi_2(\mathcal{L}, \mathcal{M}) \neq 0$. However, by decorating the mild singularities of the Legendrian knots in $\mathcal{M}$ we define a space $\mathcal{D}$, equipped with a map $\mathcal{D} \to \mathcal{M}$ which forgets the decoration, such that the composition $\mathcal{D} \to \mathcal{M} \hookrightarrow \mathcal{L}$ is surjective on π_0 and restricts to a weak homotopy equivalence on each connected component. The precise definition of the space $\mathcal{D}$ is as follows.

For any $k \geq 0$, consider the unordered configuration space $\mathrm{C}_k(S^1)$ of k distinct points on the circle $S^1 = \mathbb{R}/\mathbb{Z}$. Define a space $\widetilde{\mathrm{C}}_k(S^1)$ fibered over $\mathrm{C}_k(S^1)$ such that the fibre over the configuration $\{t_1, \dots, t_k\} \subset S^1$ consists of all unordered collections of closed intervals $I_1, \dots, I_m \subset S^1$ which are disjoint from the points $t_1, \dots, t_k$ and such that $I_i \cap I_j \neq \varnothing$ implies either $I_i \subset \mathrm{int}(I_j)$ or $I_j \subset \mathrm{int}(I_i)$. In the degenerate case where the endpoints of an interval I_j coincide, the interval consists of a point and this is allowed. The topology is such that an interval I_j which contains no other intervals in its interior can continuously shrink to a point and disappear. Observe therefore that the fibre of the map $\widetilde{\mathrm{C}}_k(S^1) \to \mathrm{C}_k(S^1)$ is contractible. We give $\widetilde{\mathrm{C}}(S^1) = \bigsqcup_k \widetilde{\mathrm{C}}_k(S^1)$ the disjoint union topology, so that the points t_i are not allowed to collide. We will refer to the elements of $\widetilde{\mathrm{C}}(S^1)$ as decorations.

Let $D = \big(\{t_i\}, \{I_j\}\big) \in \widetilde{\mathrm{C}}(S^1)$ be any decoration. We say that a Legendrian knot $f : S^1 \to \mathbb{R}^3$ is compatible with D if its front has cusp singularities at each of the points t_j and if moreover for each interval I_j the following holds. If I_j is not degenerate, then we demand that the front has cusp singularities at each of the two endpoints of I_j and moreover we require that the two cusps have opposite Maslov co-orientations. If I_j is degenerate and thus consists of a single point, then we demand that the front of f has an embryo singularity at that point. At all other points of S^1 we demand that the front is regular.

Define $\mathcal{D}$ to be the space of all pairs (f, D) such that $f : S^1 \to \mathbb{R}^3$ is a Legendrian knot compatible with a decoration $D \in \widetilde{\mathrm{C}}(S^1)$. Note in particular that $f \in \mathcal{M}$. The composition of the forgetful map $\mathcal{D} \to \mathcal{M}$ given by $(f, D) \mapsto f$ with the inclusion $\mathcal{M} \hookrightarrow \mathcal{L}$ gives a map $\mathcal{D} \to \mathcal{L}$. It is easy to see that the induced map $\pi_0(\mathcal{D}) \to \pi_0(\mathcal{L})$ is surjective but not injective. The parametric version of our h-principle implies the following result.

Corollary 1.4 *The map $\mathcal{D} \to \mathcal{L}$ is a weak homotopy equivalence on each connected component.*

Given a family of Legendrian knots in $\mathbb{R}^3$ parametrized by a space of arbitrarily high dimension, Corollary 1.4 allows us to simplify the singularities of the corresponding family of fronts so that we end up having only cusps and embryos. Moreover we have a strong control on the structure of the singularity locus (in the source) given by the family of configurations decorating the mild singularities. Proofs of Corollaries 1.1, 1.2 and 1.4, as well as of the claims made in Remark 1.3 and elsewhere in the above overview will be given in Sect. 6.

The singularities of Lagrangian and Legendrian fronts, also known as caustics in the literature, were first extensively studied by Arnold and his collaborators. See [3] for an introduction to the theory. Today, caustics still play a central role in modern symplectic and contact topology, both rigid and flexible. In many situations it is desirable for a Lagrangian or Legendrian front to have singularities which are as simple as possible. For example the Reidemeister theorem for Legendrian knots in the standard contact $\mathbb{R}^3$ (of which Corollary 1.4 is a multi-parametric generalization) has allowed for the study of $\pi_0(\mathcal{L})$ using combinatorial tools. Another example is Ekholm's method of Morse flow-trees [7] for the computation of Legendrian contact homology, which can only be applied if the caustic of the Legendrian front consists only of cusps. A rather different situation in which the simplification of caustics is desirable occurs in the arborealization program for Lagrangian skeleta pioneered by Nadler in his papers [35] and [36]. Applications of our h-principle to the arborealization program have been hinted at in Starkston's recent paper [44] as well as in Eliashberg's review of Weinstein manifold topology [19] and are the subject of present research.

The simplification of singularities of Lagrangian and Legendrian fronts is of course not always possible, since there exists a homotopy theoretic obstruction to removing higher singularities. The main point of this article is to prove that whenever this formal obstruction vanishes, the simplification can indeed be achieved by means of an ambient Hamiltonian isotopy. Our h-principle is *full* in the sense of [14] (C^0-close, relative parametric). See Sect. 1.6, where we state the result precisely, for further details. The key ingredients in the proof are (1) an explicit model for the local wrinkling of Lagrangian and Legendrian submanifolds and (2) our holonomic approximation lemma for $\perp$-holonomic sections from [1], which is a refinement of Eliashberg and Mishachev's holonomic approximation lemma [13].

Our work builds on Entov's paper [20], where the first h-principle for the simplification of caustics was proved. See Sect. 1.9 for a discussion of his results, which consist of an adaptation of Eliashberg's surgery of singularities [8,9] to the setting of Lagrangian and Legendrian fronts. Our paper instead follows the strategy employed by Eliashberg and Mishachev in the proof of their wrinkled embeddings theorem [15]. The main advantage of the wrinkled approach is the following. The surgery technique can only be applied to Σ^2-nonsingular fronts, which are fronts whose singularities have the lowest corank possible. This condition is not generic except in low dimensions. By contrast, the wrinkling technique can be applied to any front. By removing the Σ^2-nonsingularity restriction, we extend considerably the range of application of the h-principle.

Given any smooth manifold equipped with a smooth foliation, there is the analogous problem in geometric topology of simplifying of the singularities of tangency of a smooth submanifold with respect to the foliation by means of an ambient smooth isotopy. This problem also abides by an h-principle and has been studied by several authors. Gromov's method of continuous sheaves [24,26], as well as Eliashberg and Mishachev's holonomic approximation lemma [13,14] can be used to simplify the singularities of tangency when the submanifold is open. Gromov's theory of convex integration [25,26] also yields the same result. When the submanifold is closed, neither continuous sheaves nor holonomic approximation seem to work, but there are several other methods which do work. We have already mentioned two of them, namely Eliashberg's surgery of singularities [8,9] and the wrinkling embeddings theorem of Eliashberg and Mishachev [15]. Additionally, Spring showed in [42,43] that convex integration can be applied to the closed case. See also the approach of Rourke and Sanderson [38,39].

We should also mention that Corollary 1.4 can be thought of as a Legendrian analogue of Igusa's theorem [28] which states that higher singularities of smooth functions are unnecessary. The analogy becomes clearer from the viewpoint of generating functions. Closely related is another result of Igusa [29] on the high connectivity of the space of framed functions and Lurie's improvement in [32] which sketches a proof of the fact that the space of framed functions is contractible. Eliashberg and Mishachev generalized Igusa's original result in [12] and gave a proof of the contracibility of the space of framed functions in [16], in both cases using the wrinkling philosophy. There also exists a folklore approach for proving h-principles using a categorical delooping technique which was used by Galatius in unpublished work to obtain a different proof of the contractibility of the space of framed functions. The approach of Galatius inspired Kupers' recent paper [31], which provides an exposition to the delooping technique and includes yet another proof of the contractibility of the space of framed functions.

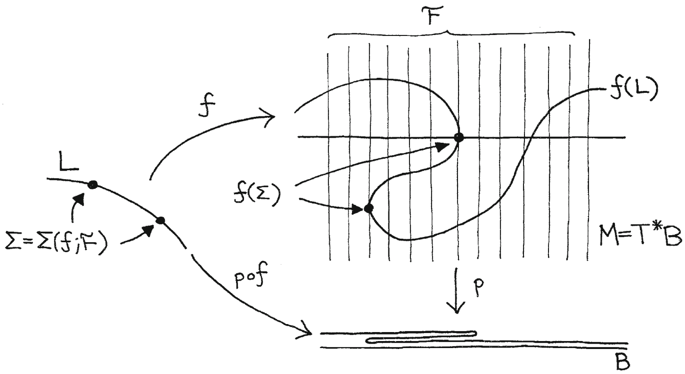

Fig. 1 The singularities of tangency of a Lagrangian embedding $f : L \to T^*B$

1.2 Singularities of tangency

Let $g : L \to B$ be any map between smooth manifolds, where we assume $\dim(L) \leq \dim(B)$ for simplicity. A point $q \in L$ is called a singularity of the map g if the differential $dg : T_qL \to T_{g(q)}B$ is not injective. The subset of L consisting of singular points is denoted by $\Sigma(g)$. Next, let $\pi : M \to B$ be a fibration of smooth manifolds and let $f : L \to M$ be a smooth embedding. The singularities of the composition $g = \pi \circ f : L \to B$ are precisely the singularities of tangency of the submanifold $f(L) \subset M$ with respect to the foliation $\mathcal{F}$ of M given by the fibres $\mathcal{F}_b = \pi^{-1}(b)$, $b \in B$. This latter notion makes sense for arbitrary foliations $\mathcal{F}$ not necessarily given by a globally defined fibration.

Definition 1.5 A singularity of tangency of an embedding $f : L \to M$ with respect to a foliation $\mathcal{F}$ of M is a point $q \in L$ such that $df(T_qL) \cap T_{f(q)}\mathcal{F} \neq 0$. The subset of L consisting of singular points is denoted by $\Sigma(f, \mathcal{F})$.

We will be interested in the special case in which (M, ω) is a symplectic $2n$-dimensional manifold and $\mathcal{F}$ is a foliation of M by Lagrangian leaves. Such a setup could arise from a Lagrangian fibration $\pi : M \to B$, where B is any n-dimensional manifold. A good example to keep in mind is the cotangent bundle $M = T^*B$ with $\pi : T^*B \to B$ the standard projection (Fig. 1).

We will also consider the analogous notion in contact topology. Here (M, ξ) is a $(2n + 1)$-dimensional contact manifold and $\mathcal{F}$ is a foliation of M by Legendrian leaves. Such a setup could arise from a Legendrian fibration $\pi : M \to B$, where B is an $(n + 1)$-dimensional manifold. A good example to keep in mind is the 1-jet space $M = J^1(E, \mathbb{R})$, where E is any n-dimensional manifold, $B = J^0(E, \mathbb{R})$ and $\pi : J^1(E, \mathbb{R}) \to J^0(E, \mathbb{R})$ is the forgetful map (which in the literature is usually referred to as the front projection). We remark for future reference that $J^1(E, \mathbb{R}) = T^*E \times \mathbb{R}$, $J^0(E, \mathbb{R}) = E \times \mathbb{R}$

and that the front projection $T^*E \times \mathbb{R} \to E \times \mathbb{R}$ is the product of the cotangent bundle projection $T^*E \to E$ and the identity map $\mathbb{R} \to \mathbb{R}$.

Suppose that $\mathcal{F}$ is induced by a Lagrangian or Legendrian fibration $\pi : M \to B$, so that the singularities of tangency $\Sigma(f, \mathcal{F}) = \{q \in L : df_q(T_qL) \cap T_{f(q)}\mathcal{F} \neq 0\}$ coincide with the singularity locus $\Sigma(p \circ f) = \{q \in L : \ker\big(d(p \circ f)_q\big) \neq 0\}$ of the smooth map $p \circ f : L \to B$. Then the composition $p \circ f$ is called the Lagrangian or Legendrian front associated to f. The image of the singularity locus $p \circ f(\Sigma) \subset B$ is called the caustic of the front.

1.3 The Thom–Boardman hierarchy

To state our results precisely, we first need to recall some notions from the Thom–Boardman hierarchy of singularities. We do not intend to be thorough and only discuss the basic facts which are necessary to frame our discussion. For a detailed exposition to the theory of singularities we refer the reader to the original papers, including those of Thom [45], Boardman [6] and Morin [33], as well as to the books [4,5] by Arnold, Gusein-Zade and Varchenko.

Suppose first that $g : L \to B$ is any smooth map between smooth manifolds, where $\dim(L) = n$ and $\dim(B) = m$. The singularity locus $\Sigma = \Sigma(g) \subset L$ of g can be stratified in the following way.

$$\Sigma = \Sigma^1 \cup \Sigma^2 \cup \dots \cup \Sigma^n, \qquad \Sigma^k = \left\{q \in L : \ \dim\left(\ker(dg_q)\right) = k\right\}.$$

The Thom transversality theorem implies that generically Σ^k is a smooth submanifold of L, whose codimension equals $k\big(m - n + k\big)$. In fact, to any non-increasing sequence I of non-negative integers $i_1 \geq i_2 \geq \dots \geq i_k$ we can associate a singularity locus $\Sigma^I \subset L$. Provided that g is generic enough so that its k-jet extension $j^k(g)$ satisfies a certain transversality condition, Σ^I is a smooth submanifold whose codimension is given by an explicit combinatorial formula. For such g, the locus Σ^I is determined inductively by $\Sigma^I = \Sigma^{i_k}\big(g|_{\Sigma^{I'}} : \Sigma^{I'} \to B\big)$, where I' denotes the truncated sequence $i_1 \geq i_2 \geq \dots \geq i_{k-1}$. In particular, $\Sigma^I \subset \Sigma^{I'}$.

We will mainly be interested in the flag of submanifolds $\Sigma^1 \supset \Sigma^{11} \supset \dots \supset \Sigma^{1^n}$, where we denote a string of $1'$s of length k by 1^k. Generically, Σ^{1^k} is a smooth codimension k submanifold of L, so that $\dim(\Sigma^{1^k}) = n - k$. To understand this flag geometrically it is useful to think of the line field $l = \ker(dg)|_{\Sigma^1} \subset TL$, which is defined along Σ^1. Inside Σ^1 we have the secondary singularity $\Sigma^{11} = \Sigma^1\big(g|_{\Sigma^1} : \Sigma^1 \to B\big)$, which consists of the set of points $q \in \Sigma^1$ where l is tangent to Σ^1. Points in the complement $\Sigma^{10} = \Sigma^1 \setminus \Sigma^{11}$, where l is transverse to Σ^1, are called fold points. Similarly, the singularity Σ^{111} consists of the set of points $q \in \Sigma^{11}$ where l is tangent

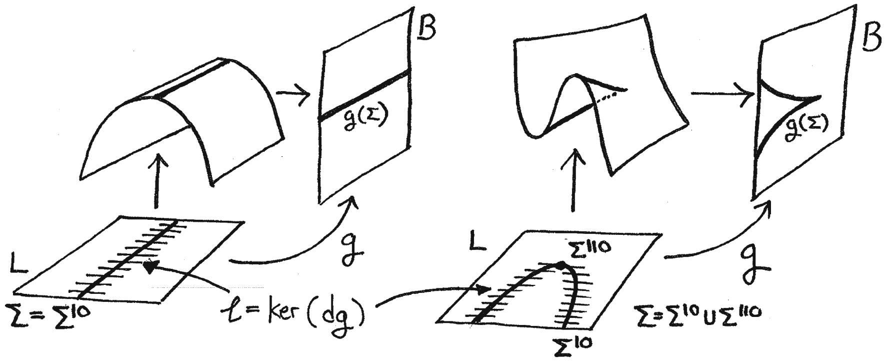

Fig. 2 The singularities Σ^{10} and Σ^{110}

to Σ^{11}. Points in the complement $\Sigma^{110} = \Sigma^{11} \setminus \Sigma^{111}$, where l is transverse to Σ^{11} inside Σ^1, are called pleats. And so on. See Fig. 2 for an illustration of Σ^{10} and Σ^{110}. Each of the singularities $\Sigma^{1^k 0} = \Sigma^{1^k} \setminus \Sigma^{1^{k+1}}$ has a unique local model and is easy to understand explicitly. We call them Σ^1-type singularities.

Singularities of type Σ^k, $k > 1$ are much more complicated than Σ^1-type singularities. In particular, there is no finite list of possible local models for the generic Σ^k singularity when $k > 1$. The situation is in fact much worse: except in simple cases where the source and target manifolds have low dimension, the generic singularities of smooth maps have moduli. Furthermore, when the dimension is sufficiently high the number of moduli is infinite. Whence the desire to simplify these complicated singularities into singularities which are at least of type Σ^1 and ideally consisting only of Σ^{10} folds.

We now return to the setting where $f : L \to M$ is a Lagrangian or Legendrian embedding into a symplectic or contact manifold M equipped with an ambient foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. We will assume that the foliation is given by the fibres of a Lagrangian or Legendrian fibration $\pi : M \to B$, indeed there is no harm in doing so since this is always the case locally. Hence the singularities of tangency $\Sigma(f; \mathcal{F})$ of f with respect to $\mathcal{F}$ are the same as the singularities $\Sigma(g)$ of the smooth mapping $g = \pi \circ f$, the Lagrangian or Legendrian front of f. Since the map g is constrained by the condition of being a Lagrangian or Legendrian front, we cannot hope for its k-jet extension $j^k(g)$ to generically satisfy the transversality condition mentioned in the definition of the loci Σ^I. For example, the generic codimension of $\Sigma^k(f; \mathcal{F})$ in L is $k(k+1)/2$, which differs from the formula given above for the singularities of smooth maps. This point is better understood from the viewpoint of generating functions, which remove the Lagrangian or Legendrian condition in exchange of increasing the jet order by one. However, we will not pause to discuss this subtlety any further since transversality can be generically achieved at the level of fronts for the singularities that we will be

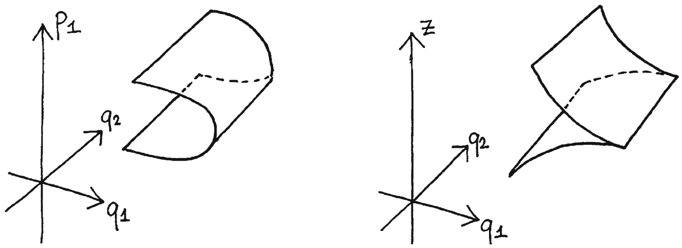

Fig. 3 The standard Σ^{10} fold. The Lagrangian submanifold on the left corresponds to the Legendrian front on the right (in that $p_j = \partial z/\partial q_j$). The former is the trivial product of a parabola $q_1 = p_1^2$ with $\mathbb{R}^{n-1}$ and the latter is the trivial product of a semi-cubical cusp $q_1^3 = z^2$ with $\mathbb{R}^{n-1}$

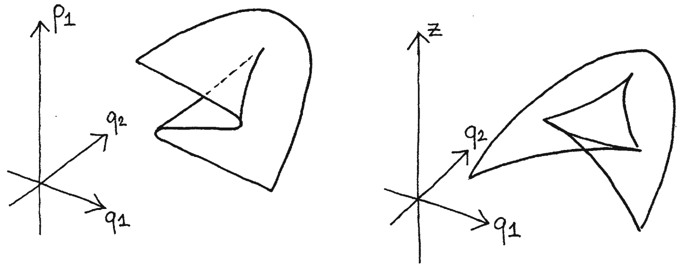

Fig. 4 The standard Σ^{110} pleat. The Lagrangian submanifold on the left corresponds to the Legendrian front on the right (in that $p_j = \partial z/\partial q_j$). The former is the birth/death of two parabolas and the latter is the birth/death of two semi-cubical cusps

interested in: the Σ^1-type singularities. In particular, the generic codimension of $\Sigma^{1^k}(f;\mathcal{F})$ in L is k, just like in the case of smooth mappings.

Figure 3 illustrates the Σ^{10} fold and Figure 4 illustrates the Σ^{110} pleat, both in their Lagrangian and Legendrian realizations. Here and below we use the standard coordinates $(q, p) \in \mathbb{R}^n \times \mathbb{R}^n = T^*\mathbb{R}^n$ and $(q, p, z) \in T^*\mathbb{R}^n \times \mathbb{R} = J^1(\mathbb{R}^n, \mathbb{R})$, where the symplectic form on $T^*\mathbb{R}^n$ is $dp \wedge dq$ and the contact form on $J^1(\mathbb{R}^n, \mathbb{R})$ is $dz - pdq$.

Example 1.6 A Lagrangian or Legendrian front has the following unique local model in a neighborhood of any fold point $q \in \Sigma^{10}$.

- In the symplectic setting where the Lagrangian fibration is $\pi : T^*B \to B$, the front $\pi \circ f : L^n \to B^n$ is locally equivalent near the point q to the map $(q_1, q_2 \dots, q_n) \mapsto (q_1^2, q_2, \dots, q_n)$ near the origin.
- In the contact setting where the Legendrian fibration is $\pi : J^1(E, \mathbb{R}) \to J^0(E, \mathbb{R})$, the front $\pi \circ f : L^n \to E^n \times \mathbb{R}$ is locally equivalent near the point q to the map $(q_1, \dots, q_n) \mapsto (q_1^2, q_2, \dots, q_n, q_1^3)$ near the origin.

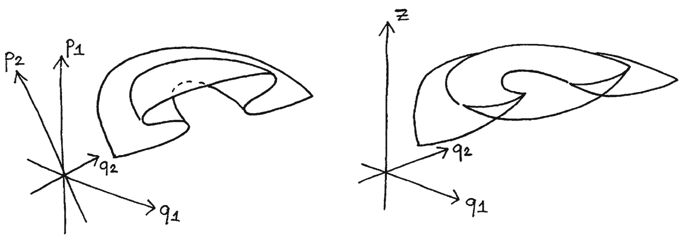

Fig. 5 One half of a double fold. The Lagrangian submanifold on the left corresponds to the Legendrian front on the right

1.4 The double fold

An example of a singularity locus which will be particularly relevant to our discussion is the so-called double fold, which we now describe. For an illustration, see Fig. 5. Before we give the definition, observe that near a fold point $q \in \Sigma^{10}$, the Lagrangian or Legendrian submanifold $f(L) \subset M$ could be turning in one of two possible directions with respect to $\mathcal{F}$. This direction can be specified by a co-orientation of the $(n-1)$-dimensional submanifold Σ^1 inside L, which is called the Maslov co-orientation and was implicitly introduced in [2]. Informally, we can view $df(T_qL)$ as a quadratic form over $T_{f(q)}\mathcal{F}$ whose signature changes by one as q crosses Σ^{10} transversely. The Maslov co-orientation specifies the direction in which the signature is increasing. This is the same Maslov co-orientation which appears in Entov's work [20].

Definition 1.7 A double fold is a pair of topologically trivial $(n-1)$-spheres S_1 and S_2 in the fold locus Σ^{10} which have opposite Maslov co-orientations and such that $S_1 \cup S_2$ is the boundary of an embedded annulus $A \subset L$.

By a topologically trivial sphere we mean a sphere which bounds an embedded n-ball in L. We say that a pair of double folds $F = S_1 \cup S_2$ and $\widetilde{F} = \widetilde{S}_1 \cup \widetilde{S}_2$ bounding annuli A and $\widetilde{A}$ in L are nested if one annulus is contained inside the other, say $A \subset \widetilde{A}$, and furthermore A bounds an n-ball $B \subset L$ which is completely contained in $\widetilde{A}$. See Fig. 6 for an illustration.

1.5 Tangential rotations

The Lagrangian Grassmannian of a symplectic manifold (M^{2n}, ω) is a fibre bundle $\Pi : \Lambda(M) \to M$ whose fibre $\Pi^{-1}(x)$ over a point $x \in M$ consists of all linear Lagrangian subspaces of the symplectic vector space (T_xM, ω_x). To each Lagrangian embedding $f : L \to M$ we can associate its Gauss map $G(df) : L \to \Lambda(M)$, given by $G(df)(q) = df(T_qL) \subset T_{f(q)}M$. Observe that $\Pi \circ G(df) = f$, in other words, $G(df)$ covers f.

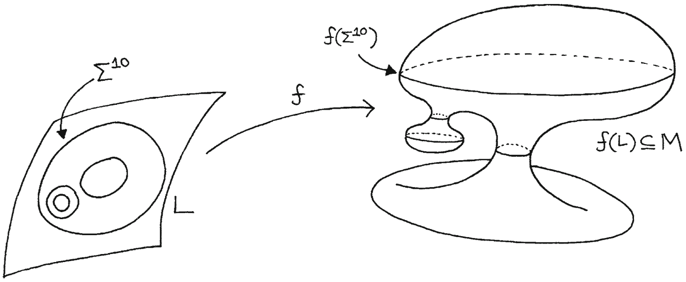

Fig. 6 A nested double fold

Similarly, given a contact manifold (M^{2n+1}, ξ), where locally $\xi = \ker(\alpha)$ for some 1-form α such that $d\alpha$ is non-degenerate on ξ, the Lagrangian Grassmannian is a fibre bundle $\Pi : \Lambda(M) \to M$ whose fibre $\Pi^{-1}(x)$ over a point $x \in M$ consists of all linear Lagrangian subspaces of the symplectic vector space $(\xi_x, d\alpha_x)$. To each Legendrian embedding $f : L \to M$ we associate its Gauss map $G(df) : L \to \Lambda(M)$, given as before by $G(df)(q) = df(T_q L) \subset \xi_{f(q)}$.

The formal analogue of the Gauss map is obtained by decoupling a Lagrangian or Legendrian embedding from its tangential information.

Definition 1.8 A tangential rotation of a Lagrangian or Legendrian embedding $f : L \to M$ is a compactly supported deformation $G_t : L \to \Lambda(M)$, $t \in [0, 1]$, of $G_0 = G(df)$ such that $\Pi \circ G_t = f$.

Example 1.9 In the previous section we introduced the double fold as an example of a singularity locus. Observe that any double fold is homotopically trivial in the following sense. If f has a double fold on the annulus $A \subset L$, then we can always construct a tangential rotation G_t of f supported in a neighborhood of A such that at time $t = 1$ we have $G_1 \pitchfork \mathcal{F}$ in that same neighborhood. In other words, there is no formal obstruction to removing a double fold.

The formal analogue of the condition $\Sigma^k(f; \mathcal{F}) = \varnothing$ is the following.

Definition 1.10 A map $G : L \to \Lambda(M)$ is called Σ^k-nonsingular with respect to the foliation $\mathcal{F}$ if $\dim(G(q) \cap T_{g(q)}\mathcal{F}) < k$ for all $q \in L$, where $g = \Pi \circ G$. When $k = 1$ we simply say that G is nonsingular, or transverse to $\mathcal{F}$, and write $G \pitchfork \mathcal{F}$.

Accordingly, we say that a Lagrangian or Legendrian embedding f is Σ^k-nonsingular with respect to $\mathcal{F}$ when $G(df)$ is Σ^k-nonsingular with respect to $\mathcal{F}$. When the foliation is clear from the context we will simply say Σ^k-nonsingular and omit the reference to $\mathcal{F}$. It is easy to see that a necessary condition for f to be Hamiltonian isotopic to a Σ^k-nonsingular embedding is the existence of a tangential rotation G_t such that G_1 is Σ^k-nonsingular. Indeed, if we denote the Hamiltonian isotopy by φ_t and we choose a family of symplectic bundle isomorphisms $\Phi_t : TM|_{f(L)} \to TM|_{\varphi_t \circ f(L)}$ such that $\Phi_0 = id$ and such that $\Phi_t\big(T\mathcal{F}|_{f(L)}\big) = T\mathcal{F}|_{\varphi_t \circ f(L)}$, then we can set $G_t = \Phi_t^{-1} \cdot G\big(d(\varphi_t \circ f)\big)$. The family Φ_t exists by the homotopy lifting property of a Serre fibration. Note that in the contact case we must replace the symplectic bundle (TM, ω) by the symplectic bundle $(\xi, d\alpha)$, but the argument is the same.

The results we state in the next section assert that this necessary condition is also sufficient when $k = 2$ and is almost sufficient when $k = 1$. The 'almost' part comes from the necessity of double folds and will be discussed below.

1.6 Main results

We are now ready to state the h-principle. Recall that M is a symplectic or contact manifold equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. By the singularities of a Lagrangian or Legendrian embedding we mean its singularities of tangency with respect to $\mathcal{F}$.

Theorem 1.11 *Suppose that there exists a tangential rotation $G_t : L \to \Lambda(M)$ of a Lagrangian or Legendrian embedding $f : L \to M$ such that $G_1 \pitchfork \mathcal{F}$. Then there exists a compactly supported Hamiltonian isotopy $\varphi_t : M \to M$ such that the singularities of $\varphi_1 \circ f$ consist of a union of nested double folds.*

Remark 1.12 In particular, $\varphi_1 \circ f$ is Σ^2-nonsingular. Indeed all of its singularities are of the simplest possible type, namely the Σ^{10} fold.

Theorem 1.11 is a full h-principle in the sense of [14]. More precisely, the following C^0-close, relative and parametric versions of the statement hold.

- (C^0-close) We can choose the Hamiltonian isotopy φ_t to be arbitrarily C^0-close to the identity. Moreover, we can arrange it so that $\varphi_t = id_M$ outside of an arbitrarily small neighborhood of $f(L)$ in M.
- (relative) Suppose that $G_t = G(df)$ on $Op(A) \subset L$ for some closed subset $A \subset L$, where here and below we use Gromov's notation $Op(A)$ for an arbitrarily small but unspecified neighborhood of A. Then we can arrange it so that $\varphi_t = id_M$ on $Op\big(f(A)\big) \subset M$.
- (parametric) An analogous result holds for families of Lagrangian or Legendrian embeddings parametrized by a compact manifold of any dimension. The statement also holds relative to a closed subset of the parameter space. For example, it holds for the pair (D^n, S^{n-1}) formed by the unit disk and its boundary sphere. For details see Sect. 6.

For singularities of type Σ^2 we have the following h-principle, in which we don't have to worry about the presence of double folds since they are singularities of type Σ^1.

Theorem 1.13 *Suppose that there exists a tangential rotation $G_t : L \to \Lambda(M)$ of a Lagrangian or Legendrian embedding $f : L \to M$ such that G_1 is Σ^2-nonsingular with respect to the foliation $\mathcal{F}$. Then there exists a compactly supported Hamiltonian isotopy $\varphi_t : M \to M$ such that $\varphi_1 \circ f$ is Σ^2-nonsingular.*

In fact, we prove a much stronger version of Theorem 1.13 which allows for the prescription of any homotopically allowable Σ^1-type singularity locus. The precise statement is given in Theorem 1.17 below, after we discuss Entov's results on the surgery of Lagrangian and Legendrian singularities.

1.7 The homotopical obstruction

Consider the subset $\Sigma(M, \mathcal{F}) \subset \Lambda(M)$ which over each point $x \in M$ consists of all planes $P_x \in \Lambda(M)_x$ such that $P_x \cap T_x\mathcal{F} \neq 0$. We have a stratification $\Sigma(M, \mathcal{F}) = \bigcup_k \Sigma^k(M, \mathcal{F})$, where $\Sigma^k(M, \mathcal{F}) = \{P_x : \dim(P_x \cap T_x\mathcal{F}) = k\}$. The formal obstruction to Σ^k-nonsingularity can be understood as follows: *is it possible to smoothly homotope the map $G(df) : L \to \Lambda(M)$ through maps G_t covering f so that its image becomes completely disjoint from the subset $\Sigma^k(M, \mathcal{F}) \subset \Lambda(M)$?* This is a purely topological question.

The most obvious cohomological obstruction is given by the higher Maslov classes. To define them, observe that $\Sigma^k(M, \mathcal{F}) = \{P_x \in \Lambda(M)_x : \dim(P_x \cap T_x\mathcal{F}) = k\}$ is a stratified subset of codimension $k(k+1)/2$ inside the Grassmannian $\Lambda(M)$, whose boundary $\partial\Sigma^k(M, \mathcal{F}) = \bigcup_{l>k} \Sigma^l(M, \mathcal{F})$ has dimension strictly less than $\dim\big(\Sigma^k(M, \mathcal{F})\big) - 1$. We can therefore define $\mu_k = G(df)^* m_k \in H^{k(k+1)/2}(L; \mathbb{Z}/2)$, where $m_k \in H^{k(k+1)/2}\big(\Lambda(M); \mathbb{Z}/2\big)$ is Poincaré dual to the cycle $\big[\Sigma^k(\mathcal{F})\big]$. The class μ_k is an obstruction to removing the singularity Σ^k.

Remark 1.14 Even for orientable M and $\mathcal{F}$, the characteristic class μ_k lifts over $\mathbb{Z}$ only for k odd (for example, the lift of μ_1 in $H^1(L; \mathbb{Z})$ is the familiar Maslov class). This can be seen by the following simple argument, which the author learnt from Givental (private communication). In the homogeneous space $\Lambda_n = U(n)/O(n)$, consider the subset $\Sigma_k = \{[A] : A \in U(n),\ \dim(A\mathbb{R}^n \cap \mathbb{R}^n) = k\} \subset \Lambda_n$. Then it is easy to check that the normal space to Σ_k in Λ_n at $[A] \in \Sigma_k$ can be identified with the space of quadratic forms on the intersection $A\mathbb{R}^n \cap \mathbb{R}^n$. Hence the normal bundle to Σ_k in Λ_n is isomorphic to the pullback by the projection $\Sigma_k \to \mathrm{Gr}_{n,k}$, $[A] \mapsto A\mathbb{R}^n \cap \mathbb{R}^n$

of the second symmetric power of the tautological bundle on the (n, k) Grassmannian $\mathrm{Gr}_{n,k}$. When k is even, the generator of $\pi_1(\mathrm{Gr}_{n,k})$ (which lies in the image of $\pi_1(\Sigma_k)$ by the above projection) can be easily checked to induce a change of orientation on this bundle, hence the cycle Σ_k is not orientable.

More generally, to each multi-index $I = (i_1 \geq i_2 \geq \cdots \geq i_k)$ there exists a cohomology class μ_I which obstructs the removal of Σ^I and which is the pullback of a universal class in the appropriate jet space. In addition to these cohomological obstructions there exist subtler homotopical obstructions to the simplification of singularities.

In certain situations the obstruction to the simplification of singularities can be straightforwardly seen to vanish. In Sect. 6 we explore a couple of such cases and are thus able to deduce concrete applications of our h-principle. However, in general this homotopical problem can be nontrivial. For instance, consider the setup of the nearby Lagrangian conjecture, so that $f : L \to T^*B$ is an exact Lagrangian embedding of a connected closed manifold L into the cotangent bundle of a connected closed manifold B. Abouzaid and Kragh showed in [30] that the first Maslov class μ_1 always vanishes. However, to the extent of the author's knowledge it is not known whether the higher Maslov classes μ_k must also vanish.

1.8 Strategy of the proof and outline of the paper

The strategy of proof of our main result Theorem 1.11 is an adaptation to the symplectic and contact setting of the strategy employed in Eliashberg and Mishachev's wrinkled embeddings paper [15]. Wrinkled embeddings are topological embeddings of smooth manifolds which are smooth embeddings away from a finite union of spheres of codimension 1, called wrinkles, where the mapping has cusps (together with their birth/deaths on the equator of each sphere). The rank of the differential falls by one on the wrinkling locus, hence the map fails to be a smooth embedding near the wrinkles. However, there is a well-defined tangent plane at every point of the image and so wrinkled embeddings have Gauss maps just like smooth embeddings. In this paper we define wrinkled Lagrangian and Legendrian embeddings to be wrinkled embeddings f into a symplectic or contact manifold M whose Gauss map $G(df)$ lands in the Lagrangian Grassmannian. The precise definition, together with all related terminology, is given in Sect. 2.

The point of working with wrinkled Lagrangian and Legendrian embeddings instead of regular Lagrangian and Legendrian embeddings is the following theorem, the proof of which takes up Sects. 3, 4 and 5 (a more precise breakdown of its proof is given below) (Fig. 7).

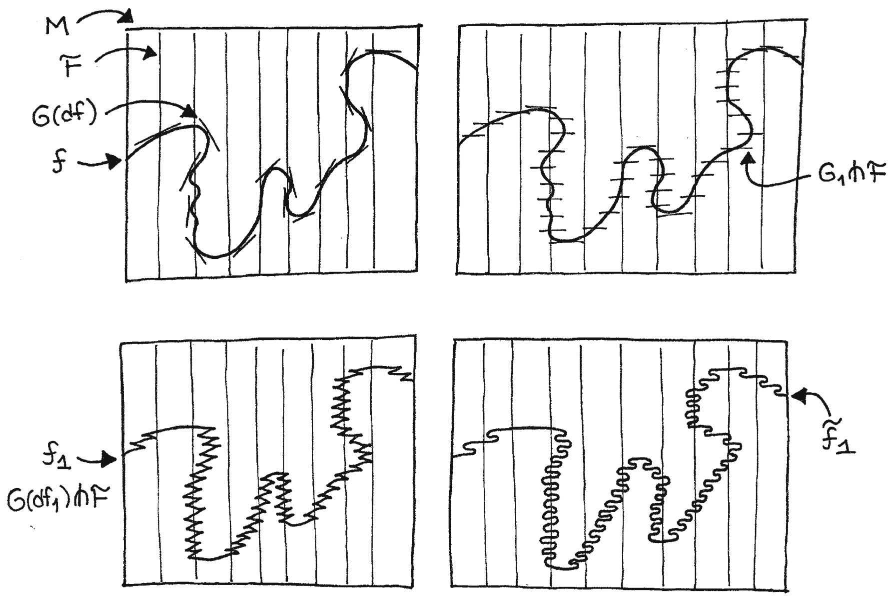

Fig. 7 The strategy of the proof

Theorem 1.15 *Any tangential rotation G_t of a regular Lagrangian or Legendrian embedding f can be C^0-approximated by the Gauss maps $G(df_t)$ of a homotopy f_t of wrinkled Lagrangian or Legendrian embeddings.*

Such a statement is of course false if we demand that the homotopy f_t consists only of regular Lagrangian or Legendrian embeddings. The additional flexibility provided by Theorem 1.15 trivially implies the following result: if there exists a tangential rotation G_t of a regular Lagrangian or Legendrian embedding f such that G_1 is transverse to an ambient foliation $\mathcal{F}$, then there exists a homotopy of wrinkled Lagrangian or Legendrian embeddings f_t such that f_1 is transverse to $\mathcal{F}$. By a regularization process (which is C^0 small but C^1 large) we can smooth out the wrinkles of f_t and obtain a homotopy of regular Lagrangian or Legendrian embeddings $\widetilde{f}_t$. The embedding $\widetilde{f}_1$ is no longer transverse to $\mathcal{F}$, we must of course pay a price when we pass from f_1 to $\widetilde{f}_1$. The price is the following: the regularization process causes Σ^{10} folds to appear where the embedding used to be wrinkled, with Σ^{110} pleats on the equator of each wrinkle. But these Σ^{110} pleats are not necessary, we can use a surgery of singularities technique to get rid of them. The result of the surgery is a union of double folds, as in the conclusion of our h-principle. See Fig. 7 for an illustration of the strategy. We formalize the process described in this paragraph in Sect. 6, where we also present applications of the h-principle.

The heart of the matter is therefore to prove the C^0-approximation result for wrinkled Lagrangian and Legendrian embeddings stated in Theorem 1.15. The steps in the proof of this result are roughly as follows.

Step 1 (Section 3) We first restrict the class of tangential rotations under consideration. A suitable class is that of simple rotations, which are those that fix a hyperplane field in the tangent space to the embedding, leaving only one degree of freedom to rotate. The key result we prove is that any tangential rotation can be C^0 approximated by a piecewise simple tangential rotation. By piecewise simple we mean that there is some subdivision of the time interval [0, 1] such that on each subinterval the Lagrangian plane field G_t always contains a fixed field of isotropic $(n-1)$-planes (the field of course depends on the subinterval). Using this result we reduce to proving Theorem 1.15 for simple tangential rotations, but now allowing for the possibility that f is wrinkled to start with.

Step 2 (Section 4) We use a refinement of the holonomic approximation lemma of Eliashberg and Mishachev [13] to construct a homotopy of our embedding f (which now may have wrinkles!) such that the Gauss map of the homotopy approximates our simple rotation G_t near the wrinkles. This is achieved by wiggling f via an ambient Hamiltonian isotopy, so no new wrinkles are needed at this stage. The refinement in question is a version of the holonomic approximation lemma in which cutoffs are carefully controlled. We established this refinement in [1]. The control in the cutoffs allows us to perform the wiggling of f in such a way that the simplicity condition is preserved up to an error which can be made arbitrarily small.

Step 3 (Section 5) We construct by hand a local relative wrinkling model which allows us to add wrinkles to the homotopy produced in the previous step so that the resulting Gauss map globally approximates G_t on the whole submanifold, completing the proof of Theorem 1.15. The simplicity of G_t is essential in order for us to reduce the general problem to an explicit local model. Our model is analogous to the model used by Eliashberg and Mishachev in [16] but some care is needed in order to adapt their construction to the Lagrangian and Legendrian settings without the inevitable cutoffs introducing uncontrolled error terms.

This completes the outline of the proof. We conclude this introduction with some brief comments on the techniques of surgery of singularities and wrinkling.

1.9 Surgery of singularities

In his thesis [9], Eliashberg developed a technique to modify the singularity locus of a Σ^2-nonsingular map between smooth manifolds by means of a surgery construction, see Fig. 8 for an example. This technique yields an

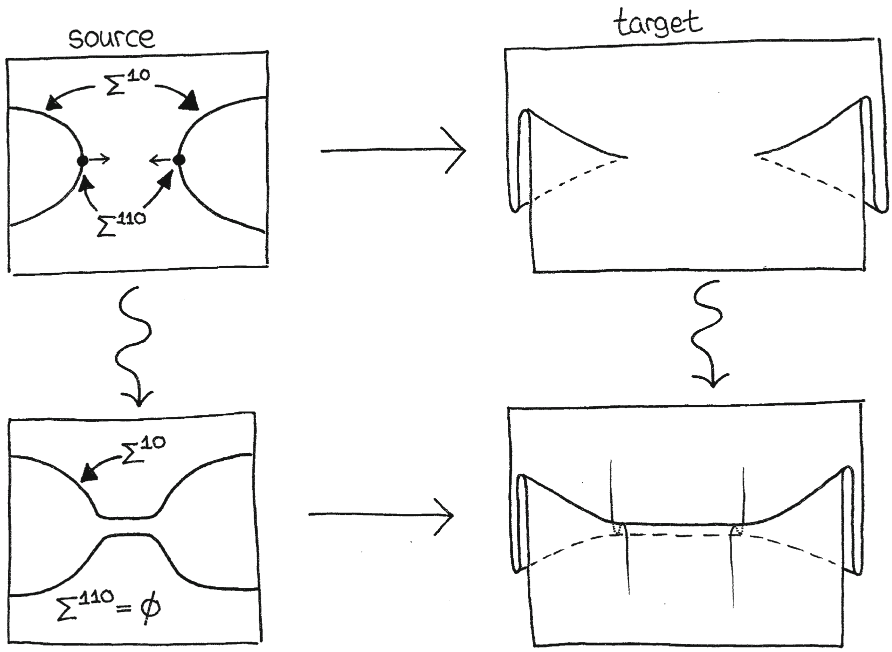

Fig. 8 An example of the surgery of singularities

h-principle for the simplification of singularities of Σ^2-nonsingular smooth maps. Almost thirty years later, Entov adapted this surgery technique to the setting of Lagrangian and Legendrian fronts, also in his thesis [20]. The main point in Entov's construction is to write down the generating functions that produce Eliashberg's surgeries, but some additional subtleties arise such as the Maslov co-orientation. As a consequence of Entov's results, one obtains an h-principle for the simplification of singularities of Σ^2-nonsingular Lagrangian or Legendrian fronts, which we now briefly discuss.

Suppose that $f : L \to M$ is a Σ^2-nonsingular Lagrangian or Legendrian embedding into a symplectic or contact manifold M equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. We recall that Σ^2-nonsingularity means that $\dim(df(T_qL) \cap T_{f(q)}\mathcal{F}) < 2$ for all $q \in L$, hence $\Sigma^2 = \varnothing$. The Thom–Boardman stratification of the singularity locus $\Sigma = \Sigma^1$ therefore consists of a flag of submanifolds $\Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$, where $\dim(\Sigma^{1^k}) = n-k$. This flag, together with certain co-orientation data which we won't be precise about right now, is called the chain of singularities associated to the embedding f and the foliation $\mathcal{F}$. More generally, given any Lagrangian distribution D defined along $f(L)$ (not necessarily tangent to an ambient foliation), we say that D is Σ^2-nonsingular if $\dim(df(T_qL) \cap D_{f(q)}) < 2$ for all $q \in L$. For such Lagrangian distributions D we can similarly define an associated chain of singularities consisting of a flag $\Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$ together with certain co-orientation data.

We say that two chains of singularities are equivalent if the flags of submanifolds are isotopic in L, with the corresponding co-orientation data also matching up under the isotopy. Entov's main result can be phrased as follows.

Theorem 1.16 (Entov) *Let $f : L \to M$ be a Σ^2-nonsingular Lagrangian or Legendrian embedding into a symplectic or contact manifold M equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. Let D_t be a homotopy of Σ^2-nonsingular Lagrangian distributions defined along $f(L)$, fixed outside of a compact subset and such that $D_0 = T\mathcal{F}|_{f(L)}$. We moreover assume that $f \pitchfork \mathcal{F}$ outside of that compact subset. Then there exists a C^0-small compactly supported Hamiltonian isotopy $\varphi_t : M \to M$ such that the chain of singularities of $\varphi_1 \circ f$ with respect to $\mathcal{F}$ is equivalent to the chain of singularities of f with respect to D_1, together with a union of nested double folds.*

Suppose that $G(df) \pitchfork D_1$. Then the chain of singularities associated to f and D_1 is empty and the conclusion of Entov's theorem is the same as the one in our Theorem 1.11. It is no coincidence that both Entov's result and Theorem 1.11 only work up to a union of double folds. Although homotopically trivial, one cannot hope to get rid of these double folds in general. The rigidity of Lagrangian and Legendrian folds was first explored by Entov in [21] and by Ferrand and Pushkar in [22] and [23]. We note that for singularities of smooth maps as considered by Eliashberg in [8] and [9] the situation is slightly better: one can always absorb these double folds into an already existing fold locus with the only condition that this locus is nonempty.

The main limitation of the surgery technique is that it requires Σ^2-nonsingularity of the initial embedding to even get started. A generic Lagrangian or Legendrian embedding is Σ^2-nonsingular only when the Lagrangian or Legendrian has dimension ≤ 2. This restricts significantly the possible applications of the surgery h-principle beyond the case of Lagrangian or Legendrian surfaces. Even in the 2-dimensional case, Σ^2-type singularities will generically arise in 1-parametric families, preventing a satisfactory parametric result from being formulated.

This limitation is not serious in the smooth version of the problem because one can easily get rid of Σ^2-type singularities by using a different technique, for example one can use Gromov's convex integration (the partial differential relation in question is ample, see Section 2.4 of [26]). Unfortunately, these techniques seem to be inadequate to get rid of the Σ^2-type singularities of Lagrangian and Legendrian fronts. We bypass this issue in the present article by using a different strategy, namely the wrinkling philosophy. Indeed, we will prove in Sect. 6.3 the following version of Entov's Theorem 1.16 in which the condition of Σ^2-nonsingularity is dropped.

Theorem 1.17 *Let $f : L \to M$ be a Lagrangian or Legendrian embedding into a symplectic or contact manifold M equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. Let D_t be a homotopy of Lagrangian distributions defined along $f(L)$, fixed outside of a compact subset, such that $D_0 = T\mathcal{F}|_{f(L)}$ and such that D_1 is Σ^2-nonsingular. We moreover assume that $f \pitchfork \mathcal{F}$ outside of that compact subset. Then there exists a C^0-small compactly supported Hamiltonian isotopy $\varphi_t : M \to M$ such that $\varphi_1 \circ f$ is Σ^2-nonsingular with respect to $\mathcal{F}$ and moreover such that the chain of singularities of $\varphi_1 \circ f$ with respect to $\mathcal{F}$ is equivalent to the chain of singularities of f with respect to D_1, together with a union of nested double folds.*

Remark 1.18 Theorem 1.13, the h-principle for Σ^2-nonsingular embeddings, is an immediate consequence of Theorem 1.17.

1.10 The wrinkling philosophy

Many h-principles can be proved by interpolating between local Taylor approximations. To achieve this interpolation near a subset of positive codimension, one can use the extra dimension to wiggle the subset in and out, creating extra room. This room ensures that no big derivatives arise when interpolating from one Taylor polynomial to another. This idea has been present throughout the history of the h-principle starting with the immersion theory of Smale-Hirsch-Phillips [27,37,40] and Gromov's method of flexible sheaves [24,26]. The wiggling strategy was reformulated into a simple but general statement by Eliashberg and Mishachev in [13,14] with their holonomic approximation lemma.

In many cases, however, one wishes to prove a global h-principle on the whole manifold (which might be closed) and there is no extra dimension available for wiggling. In the absence of additional hypotheses (such as ampleness), the wrinkling philosophy provides a strategy for proving global h-principles. The idea is to wrinkle the manifold back and forth upon itself. One can then interpolate between local Taylor approximations along the wrinkles. The wrinkling process creates the extra room needed so that this interpolation does not create big derivatives. One pays an unavoidable price, namely the singularities caused by the wrinkles. However, these are very simple singularities which can be explicitly understood (Fig. 9).

In their papers [10–12,15–17], Eliashberg and Mishachev exploit this wrinkling strategy to prove a number of results in flexible geometric topology. Together with Galatius, they give a further application in [18]. The theorem on wrinkled embeddings from [15], which is particularly relevant for our purposes, has gained greater significance after it was used by Murphy in [34] to establish the existence of loose Legendrians in high-dimensional contact man-

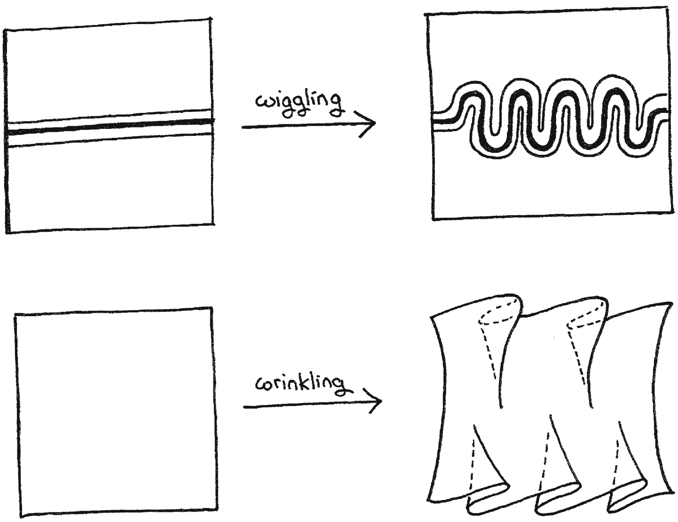

Fig. 9 The difference between wiggling and wrinkling

ifolds. Our paper provides a different application of the wrinkled embeddings theorem to flexible symplectic and contact topology.

Warning 1.19 At this point we should alert the reader that Murphy's wrinkled Legendrians are not the same as our wrinkled Lagrangian and Legendrian embeddings. The two notions should not be confused, despite the terribly similar terminology for which the author can only apologize and excuse himself in the desire to be consistent with the existing literature [15].

To be clear: in Murphy's wrinkled Legendrians, the wrinkles occur in the Legendrian front. In the wrinkled Lagrangian and Legendrian embeddings under consideration in this paper, the wrinkles occur in the Lagrangian or Legendrian submanifold itself (see Remark 2.5 below).

2 Lagrangian and Legendrian wrinkles

2.1 Wrinkled embeddings

We start by recalling the definition of wrinkled embeddings, from [15]. Throughout we denote a point $q \in \mathbb{R}^n$ by $q = (\hat{q}, q_n)$, where $\hat{q} = (q_1, \dots, q_{n-1})$.

Definition 2.1 A wrinkled embedding is a topological embedding $f : L^n \to X^{n+r}$ which is a smooth embedding away from a disjoint union of finitely many topologically trivial embedded $(n-1)$-spheres $S \subset L$, with f equivalent (up

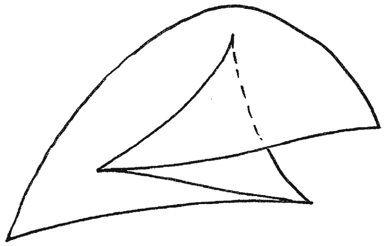

Fig. 10 One half of a standard wrinkle

to diffeomorphism) on $Op(S)$ to the local model $\mathcal{W}_{n,r} : Op(S^{n-1}) \subset \mathbb{R}^n \to \mathbb{R}^{n+r}$ given by (Fig. 10)

$$(q_1, \ldots, q_n) \mapsto (q_1, \ldots, q_{n-1}, \eta, 0, \ldots, 0, h),$$
$$\text{where } \eta(q) = q_n^3 + 3\left(||\hat{q}||^2 - 1\right) q_n \text{ and}$$
$$h(q) = \int_0^{q_n} \left(||\hat{q}||^2 + u^2 - 1\right)^2 du.$$

We recall that by topologically trivial we mean that each sphere is the boundary $S = \partial B$ of an embedded $n-$ball $B \subset L$. We say that f has a wrinkle along each S. The wrinkle itself is the germ of f in a neighborhood of S. By definition all the wrinkles are equivalent and the above formula gives an explicit model.

The mapping $\mathcal{W}_{n,r}$ has singularities along S^{n-1}. On the upper and lower hemispheres $S^{n-1} \cap \{q_n > 0\}$ and $S^{n-1} \cap \{q_n < 0\}$, the singularities are semi-cubical cusps. More precisely, near each point of $S^{n-1} \backslash S^{n-2}$, the model $\mathcal{W}_{n,r}$ is locally equivalent to the following map near the origin, see Fig. 11.

$$(q_1, \ldots, q_n) \mapsto \left(q_1, \ldots, q_{n-1}, q_n^2, 0, \ldots, 0, q_n^3\right).$$

On the equator $S^{n-2} = S^{n-1} \cap \{q_n = 0\}$, the singularities are the birth/death of semi-cubical zig-zags. More precisely, near each point of S^{n-2}, the model $\mathcal{W}_{n,r}$ is locally equivalent to the following map near the origin, see Fig. 12.

$$(q_1, \ldots, q_n) \mapsto \left(q_1, \ldots, q_{n-1}, q_n^3 - 3q_1 q_n, 0, \ldots, 0, \int_0^{q_n} \left(u^2 - q_1\right)^2 du\right).$$

Warning 2.2 Observe that a wrinkled embedding has singularities along the wrinkles, but these are not singularities of tangency with respect to any foliation. These are (non-generic) singularities of the smooth map, in other words, points in the source where the rank of the differential is strictly less than the possible maximum. Throughout the paper we will be talking about both types

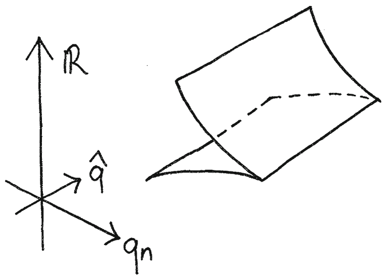

Fig. 11 A wrinkled embedding has cusps on the complement of the equator of each wrinkle

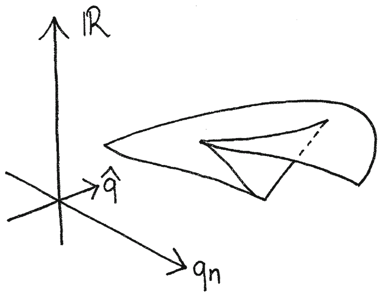

Fig. 12 A wrinkled embedding has birth/deaths of zig-zags on the equator of each wrinkle

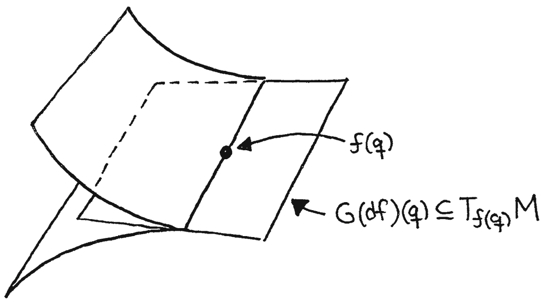

Fig. 13 A wrinkled embedding has a well-defined Gauss map everywhere, including points in the wrinkling locus

of singularities but it should always be clear from the context which type we are referring to in each case.

A wrinkled embedding has a well defined Gauss map $G(df) : L \to Gr_n(X)$, where $Gr_n(X)$ is the Grasmannian of n-planes in TX. For each $q \in L$ there is a unique n-dimensional subspace $G(df)(q) \subset T_{f(q)}X$ tangent to $f(L)$ at $f(q)$. At regular points $q \in L$ we have of course $G(df)(q) = df(T_qL)$, but $G(df)(q)$ is defined even at singular points, see Fig. 13.

2.2 Wrinkled Lagrangian and Legendrian embeddings

Let (M, ω) be a symplectic manifold.

Definition 2.3 A wrinkled Lagrangian embedding is a topological embedding $f : L^n \to (M^{2n}, \omega)$ which is a smooth Lagrangian embedding away from a disjoint union of finitely many topologically trivial embedded $(n-1)$-spheres $S \subset L$, with f equivalent (up to symplectomorphism) on $Op(S) \subset L$ to the local model $\mathcal{L}_n : Op(S^{n-1}) \subset \mathbb{R}^n \to (T^*\mathbb{R}^n, dp \wedge dq)$ given by

$$(q_1, \dots, q_n) \mapsto \left(q_1\,,\, \dots\,,\, q_{n-1}\,,\, \eta\,,\, \frac{\partial H}{\partial q_1} - h\frac{\partial \eta}{\partial q_1}\,,\, \dots\,,\, \frac{\partial H}{\partial q_{n-1}} - h\frac{\partial \eta}{\partial q_{n-1}}\,,\, h\right)$$

$$\text{where}\quad \eta(q) = q_n^3 + 3(||\hat{q}||^2 - 1)q_n, \quad h(q) = \int_0^{q_n} (||\hat{q}||^2 + u^2 - 1)^2 du$$

$$\text{and}\quad H(q) = \int_0^{q_n} h(\hat{q}, u)\frac{\partial \eta}{\partial q_n}(\hat{q}, u)du.$$

The wrinkled Lagrangian embedding $\mathcal{L}_n$ is obtained from the wrinkled embedding $\mathcal{W}_{n,n}$ in the following way. Let (q, p) be the standard coordinates on $T^*\mathbb{R}^n = \mathbb{R}^n(q_1, \dots, q_n) \times \mathbb{R}^n(p_1, \dots, p_n)$. Keeping $p_n \circ \mathcal{W}_{n,n} = h$ fixed, for $j < n$ we replace the zero functions $p_j \circ \mathcal{W}_{n,n} = 0$ with the only possible functions (up to initial conditions) which will make the embedding Lagrangian. Informally, integrate h in the direction $\partial/\partial q_n$ and differentiate the resulting function in the directions $\partial/\partial q_j$, $j < n$. Note that this construction produces a Lagrangian object out of a smooth object, independently of which functions η and h one applies the construction to. Taking η and h to be the functions defining the local model for a wrinkled embedding we obtain the local model for a wrinkled Lagrangian embedding. The corresponding definition for Legendrians is entirely analogous. Let (M, ξ) be a contact manifold.

Definition 2.4 A wrinkled Legendrian embedding is a topological embedding $f : L^n \to (M^{2n+1}, \xi)$ which is a smooth Legendrian embedding away from a disjoint union of finitely many topologically trivial embedded $(n-1)$-spheres $S \subset L$, with f equivalent (up to contactomorphism) on $Op(S) \subset L$ to the local model $\widehat{\mathcal{L}}_n : Op(S^{n-1}) \subset \mathbb{R}^n \to \big(J^1(\mathbb{R}^n, \mathbb{R}), \xi_{std}\big)$ given by

$$(q_1, \dots, q_n) \mapsto \left(q_1\,,\, \dots\,,\, q_{n-1}\,,\, \eta\,,\, \frac{\partial H}{\partial q_1} - h\frac{\partial \eta}{\partial q_1}\,,\, \dots\,,\, \frac{\partial H}{\partial q_{n-1}} - h\frac{\partial \eta}{\partial q_{n-1}}\,,\, h\,,\, H\right)$$

$$\text{where}\quad \eta(q) = q_n^3 + 3(||\hat{q}||^2 - 1)q_n, \quad h(q) = \int_0^{q_n} (||\hat{q}||^2 + u^2 - 1)^2 du,$$

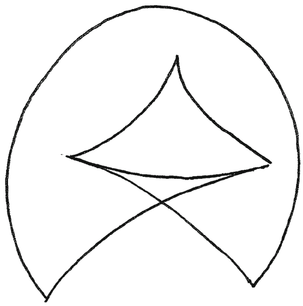

Fig. 14 The Legendrian front which generates one half of a Legendrian wrinkle. The cusps and swallowtail have a higher order of tangency than the standard cusps or swallowtails which one finds in the front projection of a regular Legendrian. To be more precise, the cusps which appear in the front projection of a Legendrian wrinkle are locally equivalent to $y^2 = x^5$, whereas the standard cusps are locally equivalent to $y^2 = x^3$

$$\text{and} \quad H(q) = \int_0^{q_n} h(\hat{q}, u) \frac{\partial \eta}{\partial q_n}(\hat{q}, u) du.$$

We recall that $J^1(\mathbb{R}^n, \mathbb{R}) = T^*\mathbb{R}^n(q, p) \times \mathbb{R}(z)$ with the standard contact structure $\xi_{std} = \ker(dz - pdq)$. The Legendrian model $\widehat{\mathcal{L}_n}$ is the Legendrian lift of the Lagrangian model $\mathcal{L}_n$ under the Lagrangian projection $J^1(\mathbb{R}^n, \mathbb{R}) \to T^*\mathbb{R}^n$, $(q, p, z) \mapsto (q, p)$. Consider also the front projection $J^1(\mathbb{R}^n, \mathbb{R}) \to J^0(\mathbb{R}^n, \mathbb{R}) = \mathbb{R}^n \times \mathbb{R}$, $(q, p, z) \mapsto (q, z)$. It is conceptually useful to understand the Legendrian front of the model $\widehat{\mathcal{L}_n}$, which is the map $Op(S^{n-1}) \subset \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}$ given by $q \mapsto \big((\hat{q}, \eta), H\big)$. On each of the hemispheres in $S^{n-1} \setminus S^{n-2}$, the front has semi-quintic cusps. On the equator $S^{n-2} \subset S^{n-1}$, the front has semi-quintic swallowtail singularities. See Fig. 14 for an illustration.

When we need to specify that a Lagrangian or Legendrian embedding $f : L \to M$ is not wrinkled, we will call f regular. Observe that the Gauss map $G(df)$ of a wrinkled Lagrangian or Legendrian embedding $f : L \to M$ lands in the Lagrangian Grassmannian $\Lambda(M)$, just like a regular Lagrangian or Legendrian embedding.

Warning 2.5 The zig-zags of a wrinkled Legendrian embedding are different from the zig-zags which appear in the loose Legendrians and wrinkled Legendrians of Murphy [34]. Indeed, the zig-zags of Murphy's wrinkled Legendrians occur in the front projection, whereas the zig-zags of our wrinkled Legendrian embeddings occur in the Legendrian submanifold itself. Moreover, the Lagrangian projection of Murphy's wrinkled Legendrians is not embedded (there is a Reeb chord in the zig-zag), whereas the Lagrangian projection of our wrinkled Legendrian is a wrinkled Lagrangian embedding, which is in particular a topological embedding. Finally, the cusps of Murphy's wrinkled

Legendrians are semi-cubic in the front projection whereas the cusps of our wrinkled Legendrian embeddings are semi-quintic in the front projection. So the two notions of wrinkled Legendrian are quite different, although of course they share the feature of exploiting the wrinkling philosophy in the context of symplectic and contact geometry.

We should also mention that a wrinkled embedding [15] is not a wrinkled map in the sense of [10], though of course the two are closely related. We make sure to always include the word 'embedding' throughout the text when referring to our wrinkled Lagrangian and Legendrian embeddings in the hope of minimizing confusion, but in any case this is the only flavor of wrinkling that will appear.

2.3 Parametric families of wrinkles

We will also consider families f^z parametrized by a smooth compact manifold Z, possibly with boundary. A family of regular Lagrangian or Legendrian embeddings $f^z : L \to M$ parametrized by Z is simply a smooth map $Z \times L \to M$, $(z, q) \mapsto f^z(q)$, such that for each $z \in Z$ the map f^z is a regular Lagrangian or Legendrian embedding. If we allow the embeddings f^z to be wrinkled, then we must allow the wrinkles to appear and disappear as the parameter z varies. Indeed, in the smooth case considered in [15], Eliashberg and Mishachev allow wrinkled embeddings to have the following local model $\mathcal{E}_{n,r} : Op(0) \subset \mathbb{R}^n \to \mathbb{R}^{n+r}$ near finitely many points. These are embryos of wrinkles, instances of birth/death.

$$(q_1, \dots, q_n) \mapsto \Big(q_1, \dots, q_{n-1}, \mu, 0, \dots, 0, e\Big),$$
$$\text{where} \quad \mu(q) = q_n^3 + 3||\hat{q}||^2 q_n \quad \text{and} \quad e = \int_0^{q_n} (||\hat{q}||^2 + u^2)^2 du.$$

In the symplectic or contact case, we can deduce corresponding local forms for Lagrangian or Legendrian embryos by integrating the function e in the direction $\partial/\partial q_n$ and then differentiating in the directions $\partial/\partial q_j$, $j < n$, just like we did in the definition of Lagrangian and Legendrian wrinkles. However, we wish to be slightly more precise in the way in which we allow wrinkles to be born or die and so we give the following definition of a family of wrinkled Lagrangian or Legendrian embeddings. We use the fibered terminology, which is a convenient language and is largely self-explanatory (the reader who wishes to see further details may consult for example [15]).

Definition 2.6 A fibered wrinkled Lagrangian embedding $f^z : L^n \to (M^{2n}, \omega)$ parametrized by an m-dimensional manifold Z is a topological embedding $f : Z \times L \to Z \times M$, $(z, q) \mapsto (z, f^z(q))$ such that f is a fibered smooth Lagrangian embedding away from a disjoint union of finitely many topologically trivial embedded $(m+n-1)$-spheres $S \subset Z \times L$, with f equivalent (up to fibered symplectomorphism) on $Op(S) \subset Z \times L$ to the local fibered model $\mathcal{L}^{\mathcal{F}}_{n,m} : Op(S^{m+n-1}) \subset \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}^m \times (T^*\mathbb{R}^n, dp \wedge dq)$ given by

$$(z_1, \dots, z_m, q_1, \dots, q_n) \mapsto \Bigg(z_1, \dots, z_m,\ q_1, \dots, q_{n-1},\ \eta\ ,\ \frac{\partial H}{\partial q_1}$$
$$-h\frac{\partial \eta}{\partial q_1}\ ,\ \dots\ ,\ \frac{\partial H}{\partial q_{n-1}} - h\frac{\partial \eta}{\partial q_{n-1}}\ ,\ h\Bigg),$$
$$\text{where} \quad \eta(z,q) = q_n^2 + 3(||z||^2 + ||\hat{q}||^2 - 1)q_n,$$
$$h(z,q) = \int_0^{q_n} (||z||^2 + ||\hat{q}||^2 + u^2 - 1)^2 du$$
$$\text{and} \qquad H(z,q) = \int_0^{q_n} h(z,\hat{q},u)\frac{\partial \eta}{\partial q_n}(z,\hat{q},u)du.$$

If we restrict $\mathcal{L}^{\mathcal{F}}_{n,m}$ to the half space $\{z_1 \geq 0\}$ we get the local model for the fibered half-wrinkles near the boundary ∂Z of the parameter space. We can define fibered wrinkled Legendrian embeddings in the exact same way, with the local model $\widehat{\mathcal{L}}^{\mathcal{F}}_{n,m} = (\mathcal{L}^{\mathcal{F}}_{n,m}, H) : Op(S^{m+n-1}) \subset \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}^m \times \big(J^1(\mathbb{R}^n, \mathbb{R}), \xi_{std}\big)$. When we talk about families of wrinkled Lagrangian or Legendrian embeddings parametrized by a compact manifold, we will always assume that the family is fibered in the sense just described.

2.4 Exact homotopies

Taking $Z = [0, 1]$ in the definition of fibered wrinkled Lagrangian or Legendrian embeddings, we obtain the notion of a homotopy of wrinkled Lagrangian or Legendrian embeddings $f_t : L \to M$, $t \in [0, 1]$, in which wrinkles are allowed to be born and to die as time goes by. The notion of exactness for homotopies of regular Lagrangian embeddings can be extended to the wrinkled case in a straightforward way.

Definition 2.7 Let $f_t : L \to M$ be a homotopy of (possibly wrinkled) Lagrangian embeddings. We say that f_t is exact if the following condition holds. For the mapping $F : L \times [0, 1] \to M$ defined by $(q, t) \mapsto f_t(q)$, consider the closed form $i_{\partial/\partial t}F^*\omega$ on $L \times [0, 1]$. We demand that this form is exact when pulled back to L by each of the inclusions $L \hookrightarrow L \times [0, 1]$, $q \mapsto (q, t)$ (Fig. 15).

Remark 2.8 Recall that if $f_t : L \to M$ is a homotopy of regular Lagrangian embeddings, then for small time $t > 0$ one can interpret f_t as a closed 1-

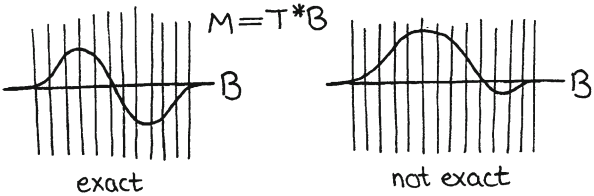

Fig. 15 The difference between an exact and a non-exact deformation of the zero section $B \hookrightarrow T^*B$. On the left, the areas cancel out, whereas on the right they do not. Exactness can be thought of as an area condition

form α_t on L by identifying a neighborhood of the zero section in T^*L with a Weinstein neighborhood of $f_0(L)$ in M. In this case exactness of f_t amounts to asking that α_t is exact for every $t \in [0, 1]$.

The importance of this definition stems from the following fact. If $f_t : L \to M$ is a compactly supported exact homotopy of regular Lagrangian embeddings, then there exists a (compactly supported) ambient Hamiltonian isotopy $\varphi_t : M \to M$ such that $f_t = \varphi_t \circ f_0$. We will always want to ensure that all homotopies of Lagrangian embeddings, regular or wrinkled, are exact. In the contact case, exactness is automatic. For convenience, we shall therefore refer to all homotopies of Legendrian embeddings, regular or wrinkled, as exact.

When a homotopy f_t is fixed on a closed subset $A \subset L$ (usually $A = L \backslash U$ is the complement of an open set U where we are performing some geometric manipulation), the notions of exactness will be understood relative to $Op(A)$. In this way, the ambient Hamiltonian isotopy inducing f_t can be taken to be the identity on $Op\big(f(A)\big) \subset M$.

2.5 Regularization of wrinkles

Wrinkles can be regularized as follows. Consider the local model $\mathcal{W}_{n,n}(q) = (\hat{q}, \eta, 0, \ldots, 0, h)$ introduced in Sect. 2.1. Let $\phi : \mathbb{R}^n \to \mathbb{R}$ be a C^∞-small function such that $\partial\phi/\partial q_n > 0$ on $S^{n-1} \subset \mathbb{R}^n$ and such that $\text{supp}(\phi) \subset Op(S^{n-1})$. Let $\widetilde{h} = h + \phi$ and observe that $\widetilde{\mathcal{W}}_{n,n}(q) = (\hat{q}, \eta, 0, \ldots, 0, \widetilde{h})$ is a smooth regular embedding such that $\widetilde{\mathcal{W}}_{n,n} = \mathcal{W}_{n,n}$ outside of $Op(S^{n-1})$, see Fig. 16.

Next, require further that

$$\int_0^{q_n} \phi(\hat{q}, u) \frac{\partial \eta}{\partial q_n}(\hat{q}, u) = 0$$

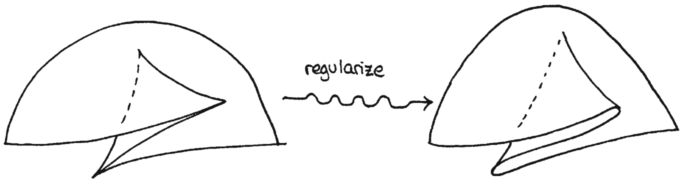

Fig. 16 Regularization of the standard wrinkle

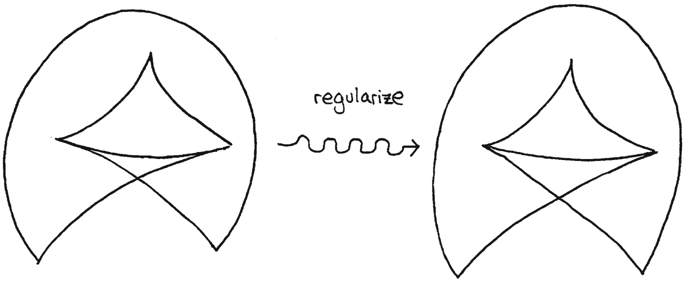

Fig. 17 The regularization can be also understood in terms of the front projection. The effect is to replace the semi-quintic cusps and swallowtails with semi-cubic cusps and swallowtails

whenever $q = (\hat{q}, q_n) \notin \text{supp}(\phi)$, and consider the modified integral

$$\widetilde{H}(q) = \int_0^{q_n} \widetilde{h}(\hat{q}, u) \frac{\partial \eta}{\partial q_n}(\hat{q}, u) du.$$

We obtain a regular Lagrangian embedding $\widetilde{\mathcal{L}}_n : Op(S^{n-1}) \to (T^*\mathbb{R}^n, dp \wedge dq)$ such that $\widetilde{\mathcal{L}}_n = \mathcal{L}_n$ outside of $Op(S^{n-1})$ by the formula

$$(q_1, \dots, q_n) \mapsto \Big(q_1, \dots, q_{n-1}, \eta, \frac{\partial \widetilde{H}}{\partial q_1} - \widetilde{h} \frac{\partial \eta}{\partial q_1}, \dots, \frac{\partial \widetilde{H}}{\partial q_{n-1}} - \widetilde{h} \frac{\partial \eta}{\partial q_{n-1}}, \widetilde{h} \Big).$$

The Legendrian counterpart of the regularization is the local model $(\widetilde{\mathcal{L}}_n, \widetilde{H})$. See Fig. 17 for an illustration of the regularization process in the front projection. Given a wrinkled Lagrangian or Legendrian embedding $f : L \to M$, we can apply this local procedure to every wrinkle and obtain a regular Lagrangian or Legendrian embedding $\widetilde{f}$. Similarly, a fibered wrinkled Lagrangian or Legendrian embedding f^z can be regularized to a fibered regular Lagrangian or Legendrian embedding $\widetilde{f^z}$. If $f_t : L \to M$ is an exact homotopy of wrinkled Lagrangian embeddings, then $\widetilde{f_t} : L \to M$ is an exact homotopy of regular Lagrangian embeddings.

The change in the order of tangency as well as the geometric meaning of the condition $\int_0^{q_n} \phi(\hat{q}, u) \frac{\partial \eta}{\partial q_n}(\hat{q}, u) du = 0$ can be better appreciated if we focus on the complement of the equator. See Fig. 18 for an illustration of the regularization process near a cusp point.

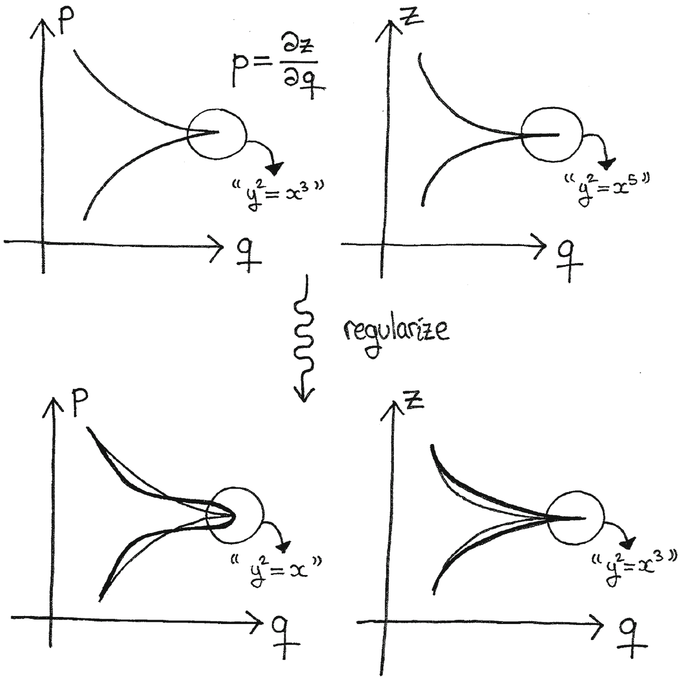

Fig. 18 Effect of the regularization process away from the equator in both the Lagrangian and front projections. The equation $\int_0^{q_n} \phi(\hat{q}, u)\frac{\partial \eta}{\partial q_n}(\hat{q}, u)du = 0$ manifests itself as an area condition in the bottom left

Remark 2.9 Observe that the regularization process $f \mapsto \widetilde{f}$ depends on the choice of ϕ. However, the space of possible ϕ is convex and therefore $\widetilde{f}$ is well defined up to a contractible choice. Different choices alter $\widetilde{f}$ by an ambient Hamiltonian isotopy supported on a neighborhood of the image of the wrinkling locus.

Remark 2.10 In the Lagrangian case, let $T^*\mathbb{R}^n$ be foliated by the fibres of the standard projection $\pi : T^*\mathbb{R}^n \to \mathbb{R}^n$ and in the contact case, let $J^1(\mathbb{R}^n, \mathbb{R}) = T^*\mathbb{R}^n \times \mathbb{R}$ be foliated by the fibres of the front projection $\pi \times id : T^*\mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n \times \mathbb{R}$. Observe that the standard Lagrangian and Legendrian wrinkles are transverse to these foliations. Moreover, when we regularize the Lagrangian or Legendrian wrinkle we obtain a regular Lagrangian or Legendrian embedding whose singularities of tangency with respect to the corresponding foliation consist of Σ^{10} folds away from the equator and of Σ^{110} pleats on the equator.

2.6 Sharpening the wrinkles

Let $D^{\pm} = \{q \in S^{n-1} | \pm q_n \geq 0\}$ be the north and south hemispheres of the unit sphere $S^{n-1} \subset \mathbb{R}^n$ and let D^{n-1} be the closed unit disk in $\mathbb{R}^{n-1}$, which we think of as sitting in $\mathbb{R}^n$ via the inclusion $\mathbb{R}^{n-1} = \mathbb{R}^{n-1} \times 0 \subset \mathbb{R}^n$. The standard Lagrangian wrinkle $\mathcal{L}_n : Op(S^{n-1}) \subset \mathbb{R}^n \to T^*\mathbb{R}^n$ is equivalent on $Op(D^{\pm}) \setminus Op(\partial D^{\pm})$ to the following local model $\mathcal{C}_n : \mathbb{R}^n \to T^*\mathbb{R}^n$ on $Op(D^{n-1}) \setminus Op(\partial D^{n-1})$.

$$\mathcal{C}_n(q_1, \dots, q_n) = \left(q_1, \dots, q_{n-1}, q_n^2, 0, \dots, q_n^3\right).$$

Note that $\mathcal{C}_n$ is the product of $\mathcal{C}_1 : \mathbb{R} \to T^*\mathbb{R}$ and the zero section $\mathbb{R}^{n-1} \hookrightarrow T^*\mathbb{R}^{n-1}$. Scaling the model $\mathcal{C}_n$ by any small number $\varepsilon > 0$ in the direction of the cotangent fibres yields a sharpened Lagrangian cusp $\varepsilon\, \mathcal{C}_n : \mathbb{R}^n \to T^*\mathbb{R}^n$. Explicitly, we set

$$\varepsilon\, \mathcal{C}_n(q_1, \dots, q_n) = \left(q_1, \dots, q_{n-1}, q_n^2, 0, \dots, \varepsilon q_n^3\right).$$

Later on it will be useful for us to be able to sharpen the cusps of a Lagrangian wrinkle at will. This sharpening can be achieved by interpolating between the two models $\mathcal{C}_n$ and $\varepsilon\, \mathcal{C}_n$. The key property of the sharpening construction is that the interpolation can be achieved by a C^1-small perturbation. The precise result that we will need is the following, where we recall the notation $q = (\hat{q}, q_n)$, $\hat{q} = (q_1, \dots, q_{n-1})$ (Fig. 19).

Lemma 2.11 *For $\delta, \varepsilon > 0$ there exists an exact homotopy $\mathcal{C}_{n,t} : \mathbb{R}^n \to T^*\mathbb{R}^n$ such that the following properties hold.*

- *$\mathcal{C}_{n,0} = \mathcal{C}_n$.*
- *$\mathcal{C}_{n,t} = \mathcal{C}_n$ when $|q_n| > 2\delta$ or $||\hat{q}|| > 1 - \delta$.*
- *$\mathcal{C}_{n,1} = \varepsilon\, \mathcal{C}_n$ when $|q_n| < \delta$ and $||\hat{q}|| < 1 - 2\delta$.*
- *$dist_{C^1}(\mathcal{C}_n, \mathcal{C}_{n,t}) \leq K\delta$, where K is a constant independent of δ and ε.*

The same Lemma also holds for the Legendrian cusp $\widehat{\mathcal{C}_n} = (\mathcal{C}_n, C) : \mathbb{R}^n \to T^*\mathbb{R}^n \times \mathbb{R} = J^1(\mathbb{R}^n, \mathbb{R})$, where $C(q) = \frac{2}{5} q_n^5$. We prove the Lagrangian and Legendrian versions simultaneously.

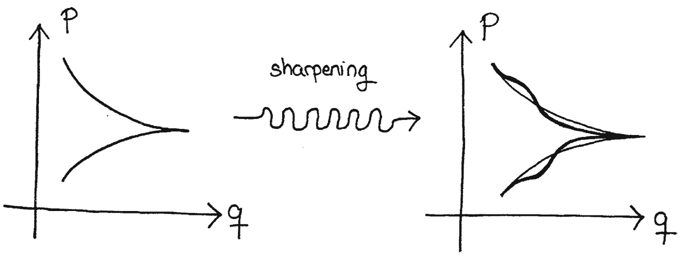

Fig. 19 Sharpening the Lagrangian cusp. Since we define the sharpening by means of a generating function, the area condition which is necessary for exactness is automatically satisfied, as shown on the picture

Proof Fix $A > 1$. Given $\delta, \varepsilon > 0$ arbitrarily small, there exists a function $\psi : \mathbb{R} \times \mathbb{R} \to [0, 1]$ satisfying the following properties:

- $\psi(x, y) = \varepsilon$ for $(x, y) \in [-\delta, \delta] \times [-1 + 2\delta, 1 - 2\delta]$,
- $\varepsilon \leq \psi(x, y) \leq 1$ for $(x, y) \in [-2\delta, 2\delta] \times [-1 + \delta, 1 - \delta] \setminus [-\delta, \delta] \times [-1 + 2\delta, 1 - 2\delta]$,
- $\psi(x, y) = 1$ for $(x, y) \notin [-2\delta, 2\delta] \times [-1 + \delta, 1 - \delta]$.
- $|\partial\psi/\partial x|,\ |\partial\psi/\partial y| \leq A/\delta$.
- $|\partial^2\psi/\partial x^2|,\ |\partial^2\psi/\partial x\partial y|,\ |\partial^2\psi/\partial y^2| \leq A/\delta^2$.
- $\partial\psi/\partial y = 0$ when $|y| < 1 - 2\delta$.

Set $\psi_t = (1 - t) + t\psi$ and $C_t(q) = \frac{2}{5}\psi_t(q_n, ||\hat{q}||)q_n^5$. The front $q \mapsto \big((\hat{q}, q_n^2), C_t\big) \in \mathbb{R}^n \times \mathbb{R}$ generates the Lagrangian and Legendrian cusps $\mathcal{C}_{n,t}$ and $\widehat{\mathcal{C}}_{n,t} = (\mathcal{C}_{n,t}, C_t)$ respectively. To be explicit, we have

$$\mathcal{C}_{n,t}(q) = \Bigg(\hat{q}, q_n^2, \frac{2}{5}\frac{\partial\psi_t}{\partial y}(q_n, ||\hat{q}||)\frac{q_1 q_n^5}{||\hat{q}||}, \ldots,$$
$$\frac{2}{5}\frac{\partial\psi_t}{\partial y}(q_n, ||\hat{q}||)\frac{q_{n-1} q_n^5}{||\hat{q}||}, \frac{1}{5}\frac{\partial\psi_t}{\partial x}(q_n, ||\hat{q}||)q_n^4 + \psi_t(q_n, ||\hat{q}||)q_n^3\Bigg).$$

The first three properties stated in the Lemma are clearly satisfied. The fourth property follows from the uniform bounds on the first and second partial derivatives of ψ. □

Next we explain how to sharpen the birth/deaths of zig-zags on the equator of each wrinkle. The standard Lagrangian wrinkle $\mathcal{L}_n : Op(S^{n-1}) \subset \mathbb{R}^n \to T^*\mathbb{R}^n$ is equivalent on $Op(S^{n-2}) \subset \mathbb{R}^n$ to the following local model $\mathcal{G}_n : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$ on $Op(S^{n-2} \times 0) \subset S^{n-2} \times \mathbb{R}^2$.

$$\mathcal{G}_n(\widetilde{q}, q_{n-1}, q_n) = \left(\widetilde{q}, q_{n-1}, \tau, 0, \frac{\partial G}{\partial q_{n-1}} - g\frac{\partial \tau}{\partial q_{n-1}}, g\right),$$
$$q = (\widetilde{q}, q_{n-1}, q_n) \in S^{n-2} \times \mathbb{R} \times \mathbb{R},$$
$$\text{where} \quad \tau(q_{n-1}, q_n) = q_n^3 - 3q_{n-1}q_n,$$
$$g(q_{n-1}, q_n) = \int_0^{q_n} (u^2 - q_{n-1})^2 du$$
$$\text{and} \qquad G(q_{n-1}, q_n) = \int_0^{q_n} g(q_{n-1}, u)\frac{\partial \tau}{\partial q_n}(q_{n-1}, u)\, du.$$

We remark that $\mathcal{G}_n$ is the product of $\mathcal{G}_2 : \mathbb{R}^2 \to T^*\mathbb{R}^2$ with the zero section $S^{n-2} \hookrightarrow T^*S^{n-2}$. For any $\varepsilon > 0$, the sharpened model $\varepsilon\,\mathcal{G}_n : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$ is given by

Fig. 20 Sharpening the Lagrangian birth/death of zig-zags

$$\varepsilon\,\mathcal{G}_n(\widetilde{q}, q_{n-1}, q_n) = \left(\widetilde{q}, q_{n-1}, \tau, 0, \varepsilon\left(\frac{\partial G}{\partial q_{n-1}} - g\frac{\partial \tau}{\partial q_{n-1}}\right), \varepsilon g\right).$$

The following result allows us to interpolate between $\mathcal{G}_n$ and $\varepsilon\,\mathcal{G}_n$ while maintaining C^1-control throughout the perturbation (Fig. 20).

Lemma 2.12 *For any $\delta, \varepsilon > 0$ there exists an exact homotopy $\mathcal{G}_{n,t} : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$ such that the following properties hold.*

- $\mathcal{G}_{n,0} = \mathcal{G}_n$.
- $\mathcal{G}_{n,t} = \mathcal{G}_n$ *when* $|q_{n-1}| > 2\delta$ *or* $|q_n| > 2\delta$.
- $\mathcal{G}_{n,1} = \varepsilon\,\mathcal{G}_n$ *when* $|q_{n-1}| < \delta$ *and* $|q_n| < \delta$.
- $dist_{C^1}(\mathcal{G}_n, \mathcal{G}_{n,t}) \leq K\delta$, *where K is a constant independent of δ and ε.*

As before, the same Lemma also holds for the Legendrian counterpart $\widehat{\mathcal{G}}_n = (\mathcal{G}_n, G) : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2) \times \mathbb{R} = J^1(S^{n-2} \times \mathbb{R}^2, \mathbb{R})$ and we prove both versions simultaneously.

Proof Fix $A > 1$. Given $\delta, \varepsilon > 0$ arbitrarily small there exists a function $\phi : \mathbb{R}^2 \to [0, 1]$ satisfying the following properties.

- $\phi(x, y) = \varepsilon$ for $(x, y) \in [-\delta, \delta]^2$,
- $\varepsilon \leq \phi(x, y) \leq 1$ for $(x, y) \in [-2\delta, 2\delta]^2 \setminus [-\delta, \delta]^2$,
- $\phi(x, y) = 1$ for $(x, y) \notin [-2\delta, 2\delta]^2$.
- $|\partial\phi/\partial x|, |\partial\phi/\partial y| \leq A/\delta$.
- $|\partial^2\phi/\partial x^2|, |\partial^2\phi/\partial x\partial y|, |\partial^2\phi/\partial y^2| \leq A/\delta^2$.
- $\partial\phi/\partial y = 0$ when $|y| < \delta$.

Set $\phi_t = (1-t) + t\phi_t$ and $G_t(q) = \phi_t(q_{n-1}, q_n)G(q)$. The front $q \mapsto \big((\widetilde{q}, q_{n-1}, \tau), G_t\big)$ generates the Lagrangian and Legendrian birth/deaths of zig-zags $\mathcal{G}_{n,t}$ and $\widehat{\mathcal{G}}_{n,t} = (\mathcal{G}_{n,t}, G_t)$ respectively. To be explicit, we have

$$\mathcal{G}_{n,t}(\widetilde{q}, q_{n-1}, q_n) = \Bigg(\widetilde{q}, q_{n-1},\ \tau\ , 0\ ,\ \frac{\partial G_t}{\partial q_{n-1}} - \left(\frac{\partial \phi_t}{\partial y}\frac{G}{(\frac{\partial \tau}{\partial q_n})} + \phi_t\, g\right)\frac{\partial \tau}{\partial q_{n-1}}\ ,\ \frac{\partial \phi_t}{\partial y}\frac{G}{(\frac{\partial \tau}{\partial q_n})} + \phi_t\, g\Bigg).$$

The first three properties stated in the Lemma are clearly satisfied. The fourth property follows from the uniform bounds on the first and second partial derivatives of ϕ. □

Remark 2.13 The sharpening construction can also be applied to a family of wrinkled Lagrangian or Legendrian embeddings. To do this, one needs to work instead with the local model for the fibered wrinkle and repeat the above construction in the fibered setting. The proofs only differ in notation.

3 Lagrangian and Legendrian rotations

3.1 Tangential rotations

In Sect. 1.5 we introduced the notion of a tangential rotation, which decouples a Gauss map $G(df) : L \to \Lambda(M)$ from its underlying Lagrangian or Legendrian embedding $f : L \to M$. We repeat the definition below for convenience. Recall that $\Pi : \Lambda(M) \to M$ denotes the Lagrangian Grassmannian of a symplectic or contact manifold M.

Definition 3.1 A tangential rotation of a regular Lagrangian or Legendrian embedding $f : L \to M$ is a compactly supported deformation $G_t : L \to \Lambda(M)$, $t \in [0, 1]$, of $G_0 = G(df)$ such that $\Pi \circ G_t = f$.

We will also need to consider tangential rotations of wrinkled Lagrangian and Legendrian embeddings. As in the unwrinkled case, a tangential rotation of a wrinkled Lagrangian or Legendrian embedding $f : L \to M$ is a compactly supported deformation $G_t : L \to \Lambda(M)$, $t \in [0, 1]$, of $G_0 = G(df)$ such that $\Pi \circ G_t = f$.

3.2 Simple tangential rotations

Let $G_t : L \to \Lambda(M)$ be a tangential rotation of a possibly wrinkled Lagrangian or Legendrian embedding $f : L \to M$. A priori, the one-parameter family of Lagrangian planes $G_t(q)$ could rotate around wildly inside $T_{f(q)}M$. It will be useful for us to restrict these rotations to be of a particularly simple type. See Fig. 21 for an illustration of the desired simplicity.

Definition 3.2 A tangential rotation $G_t : L \to \Lambda(M)$ of a possibly wrinkled Lagrangian or Legendrian embedding $f : L \to M$ is simple if there exists a field of $(n-1)$-dimensional isotropic planes $H^{n-1} \subset TM$ defined along some open subset $\mathcal{O} \subset M$ such that

- on $f^{-1}(\mathcal{O})$ we have $H \subset \mathrm{im}(G_t)$ for all $t \in [0, 1]$.
- on $L \setminus f^{-1}(\mathcal{O})$ the rotation G_t is constant.

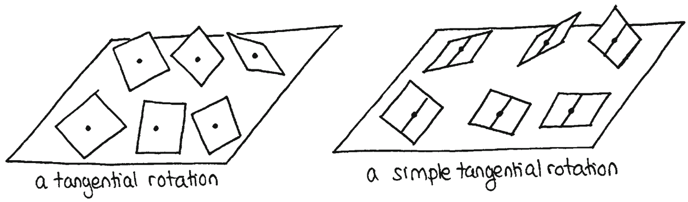

Fig. 21 The difference between a non-simple tangential rotation and a simple tangential rotation. Observe that in the simple case, the rotating planes G_t are constrained so that the $(n-1)$ directions contained in H are kept fixed, leaving only one degree of freedom

We say that G_t is simple with respect to H.

If f is regular, then we can think of $H \subset df(TL)$ as a hyperplane field in TL. When f is wrinkled we need to be a little bit careful near the wrinkling locus so it will be best to think of H as an ambient $(n-1)$-plane field in TM.

Remark 3.3 Our definition of simple tangential rotations is slightly more restrictive than what one might expect by comparing with the definition given by Eliashberg and Mishachev in [15] for the smooth analogue of this notion. This is the case because the Lagrangian or Legendrian wrinkling model that we are able to construct below is somewhat more restrictive than the model used in their proof.

We will also need the notion of piecewise simplicity. A tangential rotation G_t of a regular Lagrangian or Legendrian embedding f is piecewise simple if we can subdivide the time interval $0 = t_0 < \cdots < t_k = 1$ so that the following property holds. We demand that there exist $(n-1)$-dimensional isotropic plane fields $H^j \subset \text{im}(G_{t_j})$, which extend over open subsets $\mathcal{O}_j \subset M$, such that for all $t \in [t_j, t_{j+1}]$ we have $G_t = G_{t_j}$ outside of $f^{-1}(\mathcal{O}_j)$ and $H^j \subset \text{im}(G_t)$ on $f^{-1}(\mathcal{O}_j)$. We will prove below that any tangential rotation of a regular Lagrangian or Legendrian embedding can be C^0-approximated as accurately as desired by a piecewise simple tangential rotation. In order to do this we first translate the notion of a tangential rotation into the language of jet spaces.

3.3 Rotations of 2-jets

Let $f : L \to M$ be a regular Lagrangian embedding. Fix once and for all a Riemannian metric on L. For $\delta > 0$ small enough, the Weinstein theorem guarantees the existence of a symplectomorphism Φ between a neighborhood $\mathcal{N}$ of $f(L)$ in (M, ω) and $(T^*_\delta L, dp \wedge dq)$, where $T^*_\delta L = \{(q,p) \in T^*_q L : ||p|| < \delta\}$. We call Φ the Weinstein parametrization (Fig. 22). The zero section $L \hookrightarrow T^*_\delta L$ corresponds under Φ to the embedding $f : L \to M$. More generally, for any open subset $U \subset L$ and any function $h : U \to \mathbb{R}$ such

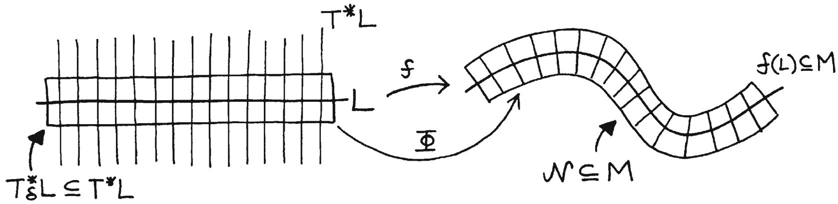

Fig. 22 A Weinstein neighborhood $\mathcal{N}$ of $f(L)$ in M

that $||dh|| < \delta$, the section $dh : U \to T^*_\delta L$ corresponds under Φ to a regular Lagrangian embedding $f_h : U \to M$ which is graphical over $f|_U$.

Similarly, if $f : L \to M$ is a regular Legendrian embedding, then for some $\delta > 0$ small enough there exists a contactomorphism Φ between a neighborhood $\mathcal{N}$ of $f(L)$ in (M, ξ) and $J^1_\delta(L, \mathbb{R}) = T^*_\delta L \times (-\delta, \delta)$, which is equipped with the standard contact structure. We still call Φ the Weinstein parametrization. For any open subset $U \subset L$ and any function $h : U \to \mathbb{R}$ such that $|h| < \delta$ and $||dh|| < \delta$, we obtain a regular Legendrian embedding $f_h : U \to M$ which is graphical over $f|_U$. The embedding f_h corresponds under Φ to the section $j^1(h) : U \to J^1_\delta(L, \mathbb{R})$.

In order to capture the tangential information contained in 1-jets we must consider 2-jets. The Riemannian metric fixed on L induces the following trivialization of the 2-jet space $J^2(L, \mathbb{R})$.

$$J^2(L, \mathbb{R}) = \{(q, z, p, Q), \quad q \in L, \quad z \in \mathbb{R}, \quad p : T_q L \to \mathbb{R}, \quad Q : T_q L \to \mathbb{R}\},$$

where p is a linear form and Q is a quadratic form. Explicitly, given a germ of a function $h : Op(q) \subset L \to \mathbb{R}$, we set $j^2(h)(q) = \big(q, h(q), dh_q, \text{Hess}(h)_q\big) \in J^2(L, \mathbb{R})$. We obtain a vector bundle $J^2(L, \mathbb{R}) \to L$, where the linear structure is induced by the above trivialization.

Example 3.4 When $L = \mathbb{R}^n$ with the standard Euclidean metric and standard coordinates $q = (q_1, \dots, q_n)$, we have a canonical identification $T_q\mathbb{R}^n \simeq \mathbb{R}^n$ for each $q \in \mathbb{R}^n$. Under this identification, $dh(v) = \sum_{i=1}^n (\partial h/\partial q_i) v_i$ and $\text{Hess}(h)(v) = \sum_{i,j=1}^n (\partial^2 h/\partial q_i \partial q_j) v_i v_j$ for all $v = (v_1, \dots, v_n) \in \mathbb{R}^n$.

Definition 3.5 A 2-jet rotation of L is a compactly supported deformation $s_t : L \to J^2(L, \mathbb{R})$, $t \in [0, 1]$, of the zero section $s_0 = 0$ which is of the form $s_t(q) = \big(q, 0, 0, Q_t(q)\big)$ for some family of quadratic forms $Q_t : TL \to \mathbb{R}$.

In other words, a 2-jet rotation is a deformation of the zero section whose 1-jet component is zero at all times. The corresponding notion of simplicity for 2-jet rotations is the following.

Definition 3.6 A 2-jet rotation $s_t : L \to J^2(L, \mathbb{R})$ is simple if there exists a hyperplane field $H \subset TL$ defined along an open subset $U \subset L$ containing

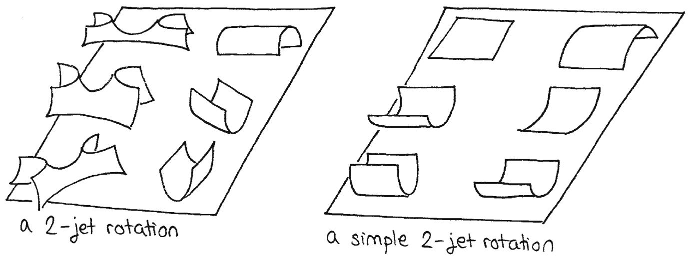

Fig. 23 The difference between a non-simple 2-jet rotation and a simple 2-jet rotation

$\mathrm{supp}(s_t)$ such that $H \subset \ker(Q_t)$ for all $t \in [0, 1]$. We say that s_t is simple with respect to H.

Remark 3.7 Observe in particular that Q_t has rank ≤ 1. However, the condition of simplicity is stronger, we demand that the kernel always contains a fixed $(n-1)$-dimensional distribution. See Fig. 23 for an illustration of 2-jet simplicity.

In the same vein, we say that a 2-jet rotation $s_t : L \to J^2(L, \mathbb{R})$ is piecewise simple if there exists a subdivision $0 = t_0 < \cdots < t_k = 1$ of the time interval $[0, 1]$ such that on each subinterval $[t_j, t_{j+1}]$ we have $s_t = s_{t_j} + r_t^j$ for some simple 2-jet rotation $r_t^j : L \to J^2(L, \mathbb{R})$.

Remark 3.8 The proper language for this discussion would naturally extend our definitions to include the concepts of l- and $\perp$-holonomic sections of the r-jet bundle associated to any fibre bundle. These ideas were introduced by Gromov in [26] in the context of convex integration. We explore these notions further in the context of holonomic approximation in our paper [1], the results of which will be crucially used below.

Given a regular Lagrangian or Legendrian embedding $f : L \to M$, a Weinstein parametrization Φ of a neighborhood $\mathcal{N}$ of $f(L)$ in M and a 2-jet rotation $s_t : L \to J^2(L, \mathbb{R})$, we can define a tangential rotation $G(\Phi, s_t) : L \to \Lambda(M)$ of f associated to Φ and s_t. Explicitly, we set $G(\Phi, s_t)(q) = G(df_{h_t})(q)$ at each point $q \in L$, where $h_t : Op(q) \subset L \to \mathbb{R}$ is any function germ such that $j^2(h_t)(q) = s_t(q)$ and $f_{h_t} : Op(q) \subset L \to M$ is the Lagrangian or Legendrian embedding corresponding to h_t under Φ. Observe that if s_t is simple, then $G(\Phi, s_t)$ is also simple.

Conversely, given a regular Lagrangian or Legendrian embedding $f : L \to M$, a Weinstein parametrization Φ and a tangential rotation $G_t : L \to \Lambda(M)$ of f, there exists a unique 2-jet rotation $s_t : L \to J^2(L, \mathbb{R})$ such that $G(\Phi, s_t) = G_t$. To be more precise, s_t might only be defined in a small time interval $[0, \varepsilon] \subset [0, 1]$, since the Lagrangian planes $G_t(q)$ could at some point stop being graphical over $df(T_qL)$ with respect to Φ, see Fig. 24.

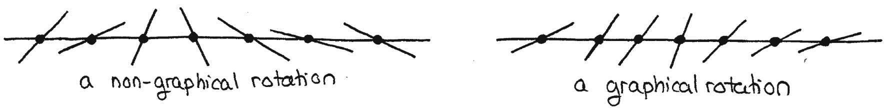

Fig. 24 The difference between a graphical and a non-graphical tangential rotation

Definition 3.9 When s_t is defined for all $t \in [0, 1]$, we say that G_t is graphical.

The Weinstein parametrization Φ is implicit in the definition. Observe again that if G_t is simple, then s_t is also simple. The notions of piecewise simplicity also coincide under this correspondence.

3.4 Approximation by simple rotations

Let $I^n = [-1, 1]^n$ denote the unit n-dimensional cube. The following lemma will allow us to replace any tangential rotation of a Lagrangian or Legendrian embedding by a piecewise simple tangential rotation.

Lemma 3.10 *Let $s_t : I^n \to J^2(\mathbb{R}^n, \mathbb{R})$ be a 2-jet rotation such that $s_t = 0$ on $Op(\partial I^n)$. Then there exists a piecewise simple 2-jet rotation $r_t : I^n \to J^2(\mathbb{R}^n, \mathbb{R})$ which is C^0-close to s_t and such that $r_t = 0$ on $Op(\partial I^n)$.*

Lemma 3.10 is an immediate consequence of a more general approximation result which we prove in [1]. For completeness we present below the outline of the argument in our concrete setting. The idea goes back to Gromov's iterated convex hull extensions in [26], which used similar decompositions into so-called principal subspaces. Indeed, in convex integration one is also forced to work one pure partial derivative at a time. These decompositions are studied carefully in Spring's book [41].

For our purposes, we only need to remark that any homogeneous degree 2 polynomial can be written as a sum of squares of linear polynomials. Explicitly, we have the polynomial identity $X_i X_j = \frac{1}{2}\big((X_i + X_j)^2 - X_i^2 - X_j^2\big)$. We can think of a 2-jet rotation as a parametric family of Taylor polynomials which are homogeneous of degree 2. By applying the above identity we obtain a decomposition $s_t = \sum r_t^{i,j}$, where the 2-jet rotation $r_t^{i,j}$ is simple with respect to the integrable hyperplane field $\tau_{i,j} = \ker(dq_i + dq_j)$ and the sum is taken over all $1 \le i \le j \le n$. Moreover, it follows that if $s_t = 0$ on $Op(\partial I^n)$, then $r_t^{i,j} = 0$ on $Op(\partial I^n)$ for all i, j (Fig. 25).

Once we have this decomposition, we can subdivide the interval $[0, 1]$ very finely and add a fraction of each $r_t^{i,j}$ at a time to obtain the desired piecewise simple approximation of s_t. The parametric version is proved in the exact same way. The statement reads as follows.

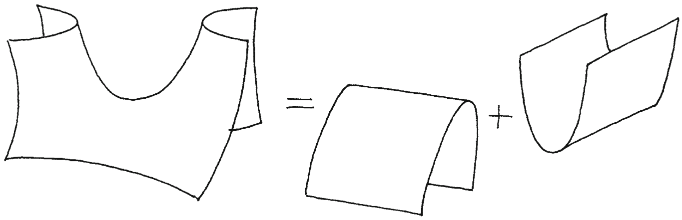

Fig. 25 Decomposing a homogeneous degree 2 polynomial into a sum of squares of linear polynomials

Lemma 3.11 *Let $s^z_t : I^n \to J^2(\mathbb{R}^n, \mathbb{R})$ be a family of 2-jet rotations parametrized by a compact manifold Z such that $s^z_t = 0$ on $Op(\partial I^n)$ and such that $s^z_t = 0$ for $z \in Op(\partial Z)$. Then there exists a family of piecewise simple 2-jet rotations $r^z_t : I^n \to J^2(\mathbb{R}^n, \mathbb{R})$ which is C^0-close to s^z_t, such that $r^z_t = 0$ on $Op(\partial I^n)$ and such that $r^z_t = 0$ for $z \in Op(\partial Z)$.*

To be more precise, for the piecewise simple family we demand that there exists a single subdivision $0 = t_0 < \cdots < t_k = 1$ of the time interval $[0, 1]$ such that every r^z_t is simple on each piece $[t_j, t_{j+1}]$. We can translate Lemmas 3.10 and 3.11 from the world of jet spaces back into the world of symplectic and contact topology. The precise consequence that we wish to extract is the following.

Proposition 3.12 *Let $G_t : L \to \Lambda(M)$ be a tangential rotation of a regular Lagrangian or Legendrian embedding $f : L \to M$. Then we can C^0-approximate G_t as much as desired by a piecewise simple tangential rotation $R_t : L \to \Lambda(M)$.*

Proof By using a partition of unity and a fine enough subdivision $0 = t_0 < \cdots < t_k = 1$ of the interval $[0, 1]$, we can localize in space and time to obtain a tangential rotation $\widetilde{G}_t : L \to \Lambda(M)$ which is C^0-close to G_t and such that on each subinterval $[t_j, t_{j+1}]$ the rotation $\widetilde{G}_t$ is constant outside of some ball $B_j \subset L$. In the Lagrangian case, let Φ_j be a symplectic isomorphism of the symplectic vector bundle $(TM|_{f(B_j)}, \omega) \to B_j$ such that $\Phi_j \cdot G(df) = \widetilde{G}_{t_j}$. In the Legendrian case, we ask that Φ_j satisfies the same property but is instead a symplectic isomorphism of the symplectic vector bundle $(\xi|_{f(B_j)}, d\alpha) \to B_j$, where $\xi = \ker(\alpha)$ on the ball B_j.

Consider the tangential rotation $S^j_t = (\Phi_j)^{-1} \cdot \widetilde{G}_t$, $t \in [t_j, t_{j+1}]$. Observe that $S^j_t = G(df)$ on $Op(\partial B_j)$. By further subdividing the time interval if necessary and picking new isomorphisms Φ_j, we may assume that S^j_t is graphical. In other words, S^j_t corresponds to a 2-jet rotation $s^j_t : B_j \to J^2(B_j, \mathbb{R})$ such that $s^j_t = 0$ on $Op(\partial B_j)$. Lemma 3.10 asserts the existence of a piecewise simple 2-jet rotation $r^j_t : B_j \to J^2(B_j, \mathbb{R})$ which is C^0-close to s^j_t and such that

$r_t^j = 0$ on $Op(\partial B_j)$. We obtain a corresponding piecewise simple tangential rotation $R_t^j : B_j \to \Lambda(M)$ which is C^0-close to S_t^j and such that $R_t^j = G(df)$ on $Op(\partial B_j)$. Set $R_t = \Phi_j \cdot R_t^j$, $t \in [t_j, t_{j+1}]$ on B_j. Outside of B_j we extend by setting $R_t = \widetilde{G}_t$, which is constant for $t \in [t_j, t_{j+1}]$. This piecewise definition yields a tangential rotation $R_t : L \to \Lambda(M)$, $t \in [0, 1]$, where each piece $R_t|_{[t_j, t_{j+1}]}$ is itself a piecewise simple tangential rotation. Hence R_t is also a piecewise simple tangential rotation. Moreover, R_t is everywhere C^0-close to $\widetilde{G}_t$, hence also to G_t. □

Remark 3.13 From the proof we can also deduce the relative version of Proposition 3.12. If $G_t = G(df)$ on $Op(A)$ for some closed subset $A \subset L$, then we can arrange it so that $R_t = G(df)$ on $Op(A)$.

The parametric version is proved in the same way. The corresponding relative version also holds. As in the case of 2-jet rotations, by a family of piecewise simple tangential rotations we mean a family of tangential rotations such that for some subdivision $0 = t_0 < \cdots < t_k = 1$ of the time interval $[0, 1]$, every tangential rotation of the family is simple on each subinterval $[t_j, t_{j+1}]$. The precise statement that we will need reads as follows.

Proposition 3.14 *Let $G_t^z : L \to \Lambda(M)$ be a family of tangential rotations of regular Lagrangian or Legendrian embeddings $f^z : L \to M$ parametrized by a compact manifold Z such that $G_t^z = G(df^z)$ for $z \in Op(\partial Z)$. Then we can C^0-approximate the family G_t^z as much as desired by a family of piecewise simple tangential rotations $R_t^z : L \to \Lambda(M)$ such that $R_t^z = G(df^z)$ for $z \in Op(\partial Z)$.*

Remark 3.15 Although we won't need this fact, we note that the piecewise simple rotations produced by our approximation process are all piecewise simple with respect to integrable hyperplane fields.

4 Wiggling embeddings

4.1 Regular approximation near the $(n-1)$-skeleton

Let $G_t : L \to \Lambda(M)$ be a tangential rotation of a regular Lagrangian or Legendrian embedding $f : L \to M$. It is in general impossible to globally C^0-approximate G_t by the Gauss maps $G(df_t)$ of an exact homotopy of regular Lagrangian or Legendrian embeddings $f_t : L \to M$, $f_0 = f$. However, it is always possible to achieve this approximation in a wiggled neighborhood of any reasonable subset of L which has positive codimension, see Fig. 26. For simplicity, we will restrict ourselves to the following class of stratified subsets.

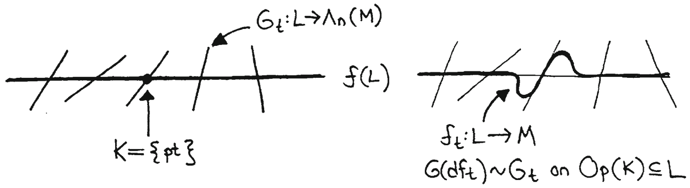

Fig. 26 We can always approximate G_t by Gauss maps $G(df_t)$ in a neighborhood of any reasonable subset $K \subset L$ of positive codimension

Definition 4.1 A closed subset $K \subset L$ is called a polyhedron if it is a subcomplex of some smooth triangulation of L.

In [1] we prove several refinements of the holonomic approximation lemma. The following result is a straightforward application of our holonomic approximation lemma for l-holonomic sections.

Theorem 4.2 *Let $K \subset L$ be a polyhedron of positive codimension and let $G_t : L \to \Lambda(M)$ be a tangential rotation of a regular Lagrangian or Legendrian embedding $f : L \to M$. Then there exists an exact homotopy of regular Lagrangian or Legendrian embeddings $f_t : L \to M$, $f_0 = f$, such that $G(df_t)$ is C^0-close to G_t on $Op(K) \subset L$.*

Remark 4.3 We can arrange it so that f_t is C^0-close to f on all of L and so that $f_t = f$ outside of a slightly bigger neighborhood of K in L. Moreover, the result also holds in relative and parametric forms.

Remark 4.4 As far as the author can tell, Theorem 4.2 is not an immediate consequence of Eliashberg and Mishachev's holonomic approximation lemma [14] or of any of the other standard h-principle techniques. However it does follow immediately from the holonomic approximation lemma for l-holonomic sections which we established in [1]. The subtlety stems from the pervasive danger of cutoffs in symplectic topology. .

In Sect. 5 we will prove that *any tangential rotation G_t can be globally C^0-approximated by the Gauss maps $G(df_t)$* of *an exact homotopy of wrinkled Lagrangian or Legendrian embeddings f_t*. This is the main technical ingredient in the proof of the h-principle for the simplification of singularities in Sect. 6 below. In the course of the proof of this global C^0-approximation theorem we will need to use a result of the same flavour as Theorem 4.2, taking K to be the $(n-1)$-skeleton of a triangulation of L. The idea is to construct the homotopy f_t by first wiggling f near the $(n-1)$-skeleton. Then one can apply a wrinkling construction in the interior of each of the top dimensional simplices to complete the approximation.

However, on the nose Theorem 4.2 is not quite sufficient for our purposes. The issue is that the local wrinkling model which we construct in Sect. 5 can only be applied if the tangential rotation is simple. Initially this is not a problem because we can use Proposition 3.12 to first approximate any given rotation by a piecewise simple rotation. We can then attempt to deal with each simple piece in the decomposition separately, working step by step. Unfortunately the following additional difficulty arises. Suppose that at a given step we apply Theorem 4.2 near the $(n-1)$-skeleton of L. We might find that our fixed decomposition is no longer piecewise simple from the viewpoint of the freshly wiggled embedding. If this is the case, then we cannot continue on to the next step, because the local wrinkling model can only be applied to simple rotations. To fix this issue we need a stronger version of Theorem 4.2 which allows us to control the wiggles with respect to any fixed simple tangential rotation. We state and prove this stronger version in the next section.

4.2 Keeping things simple

The precise result that we need is the following application of our holonomic approximation lemma for $\perp$-holonomic sections from [1]. The choice of a Riemannian metric on L and a Weinstein parametrization of a neighborhood of $f(L)$ in M is implicit throughout. We use the language of 2-jet rotations introduced in Sect. 3.3.

Theorem 4.5 *Let $K \subset L$ be a polyhedron of positive codimension and let $G_t : L \to \Lambda(M)$ be a graphical simple tangential rotation of a regular Lagrangian or Legendrian embedding $f : L \to M$. Then there exists a graphical simple tangential rotation $R_t : L \to \Lambda(M)$ of f and an exact homotopy of regular Lagrangian or Legendrian embeddings $f_t : L \to M$, $f_0 = f$, such that the following properties hold.*

- *$G(df_t)$ is C^0-close to G_t on $Op(K) \subset L$*
- *$G(df_t)$ is C^0-close to R_t on all of L.*
- *R_t is simple with respect to the same hyperplane field as G_t.*
- *$f_t = f$ and $R_t = G(df)$ outside of a slightly bigger neighborhood of K in L.*

Remark 4.6 Tthe second property implies that f_t is everywhere C^0-close to f.

Remark 4.7 The relative form of Theorem 4.5 also holds. If $G_t = G(df)$ on $Op(A)$ for some closed subset $A \subset L$, then we can arrange it so that $f_t = f$ and $R_t = G(df)$ on $Op(A) \subset L$.

Let us explain the difference between Theorems 4.2 and 4.5 and how this difference deals with the difficulty discussed at the end of Sect. 4.1. Denote by

$H \subset TL$ the hyperplane field with respect to which G_t is simple. First note that the Gauss map $G(df_t)$ of the exact homotopy f_t produced by Theorem 4.2 is an arbitrarily good approximation of G_t near K, but we have no control on $G(df_t)$ away from K. Compare with the Gauss map $G(df_t)$ of the exact homotopy f_t produced by Theorem 4.5, which is not only an arbitrarily good approximation of G_t near K, but everywhere on L only differs from $G(df)$ by a rotation which is simple with respect to H (up to an error which can be made arbitrarily small). Hence the lack of global control on $G(df_t)$ is restricted to the one degree of freedom completementary to H in TL. We don't know what $G(df_t)$ does within this one degree of freedom, but we record it and give it a name: R_t. The upshot is that from the viewpoint of f_t the rotation G_t is still simple with respect to the same hyperplane field H (up to an error which can be made arbitrarily small). This will allow us to complete the approximation of G_t by the introduction of wrinkles on f_t which are parallel to H and of magnitude $G_t - R_t$. This discussion will be made precise when it is time for us to wrinkle. In the meantime, we proceed to wiggle.

Proof of Theorem 4.5 Fix a Riemannian metric on L. By definition of graphicality, we can think of G_t as a 2-jet rotation $s_t : L \to J^2(L, \mathbb{R})$ which is simple with respect to some hyperplane field $H \subset TL$. We can therefore apply the (1-parametric) holonomic approximation lemma for $\perp$-holonomic sections from [1] to s_t. The output is a family of functions $h_t : L \to \mathbb{R}$, $h_0 = 0$ and an isotopy $F_t : L \to L$ such that the following properties hold.

- $j^2(h_t)$ is C^0-close to s_t on $Op\big(F_t(K)\big) \subset L$.
- $j^1(h_t)$ is C^0-small on all of L.
- $\mathrm{Hess}(h_t)|_H$ is C^0-small on all of L.
- F_t is C^0-small.
- $F_t^* H$ is C^0-close to H.
- $h_t = 0$ and $F_t = id_L$ outside of a slightly bigger neighborhood of K in L.

The C^1-smallness of h_t allows us to think of $dh_t \circ F_t : L \to T^*L$ (in the Lagrangian case) or of $j^1(h_t) \circ F_t : L \to J^1(L, \mathbb{R})$ (in the Legendrian case) as an exact homotopy of regular Lagrangian or Legendrian embeddings $f_t : L \to M$. We define the simple tangential rotation $R_t : L \to \Lambda(M)$ by specifying its corresponding simple 2-jet rotation $r_t : L \to J^2(L, \mathbb{R})$ as follows. Write $r_t(q) = \big(q, 0, 0, Q_t(q)\big) \in J^2(L, \mathbb{R})$ for $Q_t : TL \to \mathbb{R}$ a family of quadratic forms and set $Q_t(q) = \mathrm{Hess}(h_t)|_{F_t(q)} \circ p : TL \to \mathbb{R}$ to obtain the desired R_t, where $p : TL \to TL$ is the orthogonal projection with kernel H. Note that $\mathrm{Hess}(h_t)|_{F_t(q)}$ does not define a quadratic form on T_qL, but rather a quadratic form on $T_{F_t(q)}L$, so let us explain more carefully what we mean by $Q_t(q)$. With respect to our fixed Riemannian metric, a quadratic form on T_qL which vanishes on H_q is determined by a co-orientation of H_q and a non-negative number, namely its norm. Since F_t is C^0-small a co-orientation of

H_q induces co-orientation of $H_{F_t(q)}$. Therefore $Q_t(q)$ is uniquely determined by demanding that it vanishes on H_q and that its norm is equal to the norm of $\mathrm{Hess}(h_t)|_{F_t(q)} \circ p$ (which is a quadratic form on $T_{F_t(q)}L$ with kernel $H_{F_t(q)}$).

All the properties listed in Theorem 4.5 can now be easily checked to hold. The only property which may need clarification is the third one. To verfity it observe that the C^0-smallness of $\mathrm{Hess}(h_t)$ on H implies that $\mathrm{Hess}(h_t)|_{F_t(q)}$ is C^0-close to $\mathrm{Hess}(h_t)|_{F_t(q)} \circ p$. But $\mathrm{Hess}(h_t)|_{F_t(q)} \circ p$ is C^0-close to $Q_t(q)$, since they both have the same norm and their kernels, $H_{F_t(q)}$ and H_q respectively, are C^0-close (the degree of accuracy determined by how C^0-close F_t is to the identity).

Finally, we observe that the C^0-approximation bounds satisfied by the resulting $G(df_t)$ and R_t can be improved as much as desired by demanding the corresponding degree of C^0-approximation in the invokation or our holonomic approximation lemma for $\perp$-holonomic sections. Note that the R_t produced by the proof will depend on the desired degree of C^0-approximation. □

Remark 4.8 The condition that $F_t^* H$ is C^0-close to H was not used in this proof but we include it for the sake of intuition given that in our construction of the refined holonomic approximation [1] it is crucial to have the wiggles be almost parallel to H.

The above argument also works for families. In the parametric case, we note that the polyhedron K may also vary with the parameter. To be more precise, we have the following definition.

Definition 4.9 A closed subset $K \subset Z \times L$ is called a fibered polyhedron if it is a subcomplex of a smooth triangulation of $Z \times L$ which is in general position with respect to the fibres $z \times L$, $z \in Z$.

More precisely, the requirement is that the n-plane field $V \subset T(Z \times L)$ tangent to the fibres of the projection $Z \times L \to Z$ is transverse to each k-simplex α^k in the triangulation when $k \geq n$ and that $V + T\alpha^k \subset T(Z \times L)|_{\alpha^k}$ has dimension $n + k$ when $k \leq n$. The crucial consequence of this definition is that for every $z \in Z$ the subset $K^z \subset L$ given by $K \cap (z \times L) = z \times K^z$ is a polyhedron in L. If K has positive codimension in $Z \times L$, then K^z has positive codimension in L for all $z \in Z$. The more restrictive notion of general position considered by Thurston in [46] is not necessary for our purposes but we can also ask for it if we want to since we will rely on his existence result for triangulations in general position to foliations, which he proves in this stronger sense.

The parametric version of Theorem 4.5 is proved in the same way, by adding a parameter in the notation everywhere and invoking our parametric holonomic approximation lemma for $\perp$-holonomic sections from [1]. The statement reads as follows. We note that the relative version also holds, as in Remark 4.7.

Theorem 4.10 *Let $K \subset Z \times L$ be a fibered polyhedron of positive codimension and let $G_t^z : L \to \Lambda(M)$ be a family of graphical simple tangential rotations of regular Lagrangian or Legendrian embeddings $f^z : L \to M$ parametrized by a compact manifold Z such that $G_t^z = G(df^z)$ for $z \in Op(\partial Z)$. Then there exists a family of graphical simple tangential rotations $R_t^z : L \to \Lambda(M)$ of f^z and a family of exact homotopies of regular Lagrangian or Legendrian embeddings $f_t^z : L \to M$, $f_0^z = f^z$, such that the following properties hold.*

- *$G(df_t^z)$ is C^0-close to G_t^z on $Op(K^z) \subset L$.*
- *$G(df_t^z)$ is C^0-close to R_t^z on all of L.*
- *R_t^z is simple with respect to the same hyperplane field as G_t^z.*
- *$f_t^z = f^z$ and $R_t^z = G(df^z)$ outside of a slightly bigger neighborhood of K^z in L.*
- *$f_t^z = f^z$ and $R_t^z = G(df^z)$ for $z \in Op(\partial Z)$.*

4.3 Wiggling the wrinkles

In this section we extend Theorems 4.5 and 4.10, which were stated for regular Lagrangian or Legendrian embeddings, to the case of wrinkled Lagrangian or Legendrian embeddings. In the wrinkled case, we cannot invoke our holonomic approximation lemma for $\perp$-holonomic sections from [1] directly because a wrinkled Lagrangian or Legendrian embedding is not regular near the wrinkles. The sharpening construction described in Sect. 2.6 will allow us to resolve this issue, since the sharper the wrinkles, the better they can be approximated locally by a regular Lagrangian or Legendrian submanifold.

Given a wrinkled Lagrangian or Legendrian embedding $f : L \to M$, recall that the subset on which f is wrinkled consists of a disjoint union $W = \bigcup_j S_j$ of finitely many $(n-1)$-dimensional embedded spheres $S_j \subset L$. Each sphere S_j has an $(n-2)$-dimensional equator $E_j \subset S_j$ on which f has birth/deaths of zig-zags. The complement $S_j \setminus E_j$ consists of two hemispheres on which f has cusps.

We say that a polyhedron $K \subset L$ is compatible with the wrinkles of f if the following condition holds. We demand that the wrinkling locus $W = \bigcup_j S_j$ is contained in the $(n-1)$-skeleton of K and that the union of the equators $\bigcup_j E_j$ is contained in the $(n-2)$-skeleton of K. In the same way we can define what it means for a fibered polyhedron $K \subset Z \times L$ to be compatible with the fibered wrinkles of a family $f^z : L \to M$ of wrinkled Lagrangian or Legendrian embeddings parametrized by a compact manifold Z.

We now prove the analogue of Theorem 4.5 for wrinkled Lagrangian and Legendrian embeddings. The precise statement is the following.

Theorem 4.11 *Let $K \subset L$ be a polyhedron of positive codimension which is compatible with the wrinkles of a wrinkled Lagrangian or Legendrian* *embedding $f : L \to M$. Let $G_t : L \to \Lambda(M)$ be a graphical simple tangential rotation of f. Then there exists a graphical simple tangential rotation $R_t : L \to \Lambda(M)$ of f and a family of exact homotopies $f_t : L \to M$, $f_0 = f$, of wrinkled Lagrangian or Legendrian embeddings such that all of the properties listed in Theorem 4.5 hold.*

Proof Consider first a single wrinkle S in the wrinkling locus $W \subset L$ of f. The wiggling on S is performed in two steps. First we will wiggle f near the equator $E \subset S$ and then we will wiggle f near the remaining part of S. In both cases this wiggling is achieved by replacing the singular Lagrangian or Legendrian submanifold $f(L)$ with a regular approximation to which holonomic approximation can be applied. We then use the resulting ambient Hamiltonian isotopy to induce a wiggling of f. We will restrict our attention to the Lagrangian case for the sake of concreteness, but the Legendrian case is no different.

In Sect. 2.6 we introduced the Lagrangian local model $\mathcal{G}_n$ for the birth/death of zig-zags. Recall that $\mathcal{G}_n : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$ is given by

$$\mathcal{G}_n(\widetilde{q}, q_{n-1}, q_n) = \left(\widetilde{q}, q_{n-1}, \tau, 0, \frac{\partial G}{\partial q_{n-1}} - g\frac{\partial \tau}{\partial q_{n-1}}, g\right),$$

$$q = (\widetilde{q}, q_{n-1}, q_n) \in S^{n-2} \times \mathbb{R} \times \mathbb{R}.$$

$$\text{where} \quad \tau(q_{n-1}, q_n) = q_n^3 - 3q_{n-1}q_n,$$

$$g(q_{n-1}, q_n) = \int_0^{q_n} (u^2 - q_{n-1})^2 du$$

$$\text{and} \qquad G(q_{n-1}, q_n) = \int_0^{q_n} g(q_{n-1}, u)\frac{\partial \tau}{\partial q_n}(q_{n-1}, u)\, du.$$

Near the equator $E \subset S$, our wrinkled Lagrangian embedding $f : L \to M$ is locally equivalent to $\mathcal{G}_n$ near $S^{n-2} \times 0 \subset S^{n-2} \times \mathbb{R}^2$. Working in this local model, we can think of G_t as a tangential rotation of $\mathcal{G}_n$ which is simple with respect to an $(n-1)$-plane field $H \subset T\big(T^*(S^{n-2} \times \mathbb{R}^2)\big)$. Consider the zero section $\mathcal{Z} : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$, which is a Lagrangian cylinder. Observe that $\mathcal{Z}|_{S^{n-2}\times 0} = \mathcal{G}_n|_{S^{n-2}\times 0}$, and moreover that $G(d\mathcal{Z})|_{S^{n-2}\times 0} = G(d\mathcal{G}_n)|_{S^{n-2}\times 0}$. Extend $G_t|_{S^{n-2}\times 0}$ to $S_t : S^{n-2} \times \mathbb{R}^2 \to \Lambda\big(T^*(S^{n-2} \times \mathbb{R}^2)\big)$, a tangential rotation of $\mathcal{Z}$ which is simple with respect to H.

Let $\delta > 0$ and set $\mathcal{N} = S^{n-2} \times (-\delta, \delta)^2 \subset S^{n-2} \times \mathbb{R}^2$. Apply Theorem 4.5 to the regular Legendrian embedding $\mathcal{Z}$, the simple tangential rotation S_t and the stratified subset $S^{n-2} \times 0 \subset S^{n-2} \times \mathbb{R}^2$. We obtain an exact homotopy of regular Legendrian embeddings $\mathcal{Z}_t : S^{n-2} \times \mathbb{R}^2 \to T^*(S^{n-2} \times \mathbb{R}^2)$ which we may assume is constant outside of $\mathcal{N}$. Recall that $G(d\mathcal{Z}_t)$ is C^0-close to G_t near $S^{n-2} \times 0$. Recall also that $G(d\mathcal{Z}_t)$ is everywhere C^0-close to a tangential

Fig. 27 The sharpening construction applied to the equator

rotation R_t which is also simple with respect to H and which is supported on $\mathcal{N}$.

Write $\mathcal{Z}_t = \varphi_t \circ \mathcal{Z}$ for an ambient Hamiltonian isotopy φ_t which we may assume constant outside of $Op(\mathcal{N}) \subset T^*(S^{n-2} \times \mathbb{R}^2)$. Let $\varepsilon > 0$ and consider the sharpening $\mathcal{G}_{n,t}$ of $\mathcal{G}_n$ described in Sect. 2.6 with respect to the parameters δ and ε (Fig. 27). Recall that $\text{dist}_{C^1}(\mathcal{G}_n, \mathcal{G}_{n,t}) \leq A\delta$ for some constant $A > 0$ independent of δ and ε, so by taking $\delta > 0$ small enough we may replace $\mathcal{G}_n$ with $\mathcal{G}_{n,1}$ from the onset up to an error which is proportional to δ. Recall also that sharpening is supported on $S^{n-2} \times (-2\delta, 2\delta)^2$ and is ε-sharp on $\mathcal{N} = S^{n-2} \times (\delta, \delta)$. For details see Sect. 2.6.

Consider now $\varphi_t \circ \mathcal{G}_{n,1}$. Note that on $Op(S^{n-2} \times 0)$ the Gauss map of this composition is C^0-close to G_t. Indeed, $\mathcal{G}_{n,1}$ and $\mathcal{Z}$ are tangent along $S^{n-2} \times 0$ and when we invoke Theorem 4.5 to construct $\mathcal{Z}_t$ we can demand as much accuracy in the approximation as we want. Next, observe that $\mathcal{Z}_t$ is supported on $\mathcal{N}$ and on that neighborhood $G(d\mathcal{Z}_t)$ is C^0-close to a tangential rotation which is simple with respect to H. Let $\pi : T^*(S^{n-2} \times \mathbb{R}^2) \to S^{n-2} \times \mathbb{R}^2$ denote the standard projection. As $\varepsilon \to 0$ in the sharpening $\mathcal{G}_{n,1}$, for each $q \in \mathcal{N}$ the tangent plane $G(d\mathcal{G}_{n,1})(q)$ converges to the horizontal plane tangent to the zero section at the point $\pi(q)$, and hence $G\big(d(\varphi_t \circ \mathcal{G}_{n,1})\big)(q)$ converges to $G(d\mathcal{Z}_t)\big(\pi(q)\big)$. It follows that by taking $\varepsilon > 0$ as small as is necessary, we can use R_t to exhibit a tangential rotation of $\mathcal{G}_{n,1}$ which is simple with respect to H and which is arbitrarily C^0-close to $G\big(d(\varphi_t \circ \mathcal{G}_{n,1})\big)$ on all of L. We have therefore achieved the required global approximation up to an error which is proportional to δ. Since we can take $\delta > 0$ to be arbitrarily small, this completes the wiggling near the equator.

Once we have wiggled f near the equator E we proceed to wiggle f on the two hemispheres $D^{\pm}$ of the complement $S \backslash E$. Near the interior of each of the two disks D^+ and D^- the map f is equivalent to the local model $\mathcal{C}_n : \mathbb{R}^n \to T^*\mathbb{R}^n$ on $Op(D^{n-1}) \backslash Op(\partial D^{n-1})$, where we recall from Sect. 2.6 that

$$\mathcal{C}_n(q_1, \dots, q_n) = \big(q_1, \dots, q_{n-1}, q_n^2, 0, \dots, q_n^3\big) .$$

Our input this time is a simple tangential rotation G_t of the local model $\mathcal{C}_n|_{D^{n-1}}$ which we assume to be constant on $Op(\partial D^{n-1})$. The strategy is the same as before. Consider the zero section $\mathcal{Z} : \mathbb{R}^n \to T^*\mathbb{R}^n$ and extend $G_t|_{D^{n-1}}$ to a simple tangential rotation S_t of $\mathcal{Z}$. Then apply (the relative version of) Theorem 4.5 to $\mathcal{Z}$, S_t and D^{n-1} to obtain an exact homotopy of regular Lagrangian embeddings $\mathcal{Z}_t$ which is induced by an ambient Hamiltonian isotopy φ_t fixing the boundary. For a suitable choice of parameters δ and ε, the concatenation of the sharpening homotopy $\mathcal{C}_{n,t}$ described in Sect. 2.6 followed by the isotopy $\varphi_t \circ \mathcal{C}_{n,1}$ gives the required wiggling on $S \backslash E$.

This process can now be repeated on all wrinkles S until we have achieved the desired wiggling on the locus W where f fails to be a regular Lagrangian embedding. The proof of Theorem 4.11 is completed by applying (the relative version) of Theorem 4.5 on the regular locus. □

The analogue of the parametric Theorem 4.10 for families of wrinkled Lagrangian or Legendrian embeddings also holds, where we demand that the fibered polyhedron $K \subset Z \times L$ is compatible with the wrinkles. The proof only differs in notation and the precise statement reads as follows.

Theorem 4.12 *Let $K \subset Z \times L$ be a fibered polyhedron of positive codimension which is compatible with the wrinkles of a family of wrinkled Lagrangian or Legendrian embeddings $f^z : L \to M$ parametrized by a compact manifold Z. Let $G_t^z : L \to \Lambda(M)$ be a family of graphical simple tangential rotations of f^z such that $G_t^z = G(df^z)$ for $z \in Op(\partial Z)$. Then there exists a family of graphical simple tangential rotations $R_t^z : L \to \Lambda(M)$ of f^z and an exact homotopy of wrinkled Lagrangian or Legendrian embeddings $f_t^z : L \to M$, $f_0^z = f^z$, such that all of the properties listed in Theorem 4.10 hold.*

Remark 4.13 Observe that no wrinkles appear or disappear in the homotopies of wrinkled Lagrangian or Legendrian embeddings produced by Theorems 4.11 and 4.12. In [15], Eliashberg and Mishachev refer to the analogous smooth homotopy as an isotopy of wrinkled embeddings. We call the process 'wiggling embeddings'.

5 Wrinkling embeddings

5.1 Wrinkled approximation on the whole manifold

As we already mentioned, we cannot in general hope to globally C^0-approximate a tangential rotation $G_t : L \to \Lambda(M)$ of a regular Lagrangian or Legendrian embedding $f : L \to M$ by the Gauss maps $G(df_t)$ of a homotopy f_t of regular Lagrangian or Legendrian embeddings. In the previous section, we showed that the approximation can nevertheless be achieved by such a regular homotopy in a small neighborhood of any polyhedron $K \subset L$ of positive

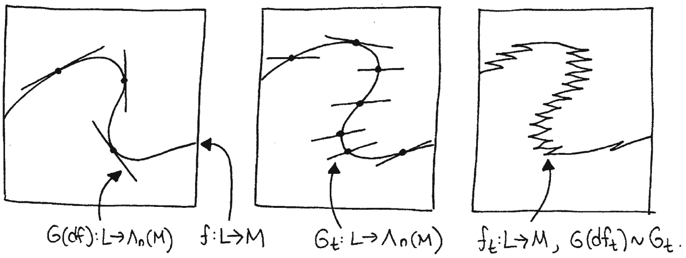

Fig. 28 The wrinkling theorem in action

codimension. In this section we show that the approximation can be globally achieved on the whole manifold L *if we allow the homotopy* f_t *to be wrinkled*. See Fig. 28 for an illustration. More, precisely we have the following theorem, which is the main result of this section.

Theorem 5.1 *Let* $G_t : L \to \Lambda(M)$ *be a tangential rotation of a regular Lagrangian or Legendrian embedding* $f : L \to M$. *Then there exists a compactly supported exact homotopy of wrinkled Lagrangian or Legendrian embeddings* $f_t : L \to M$, $f_0 = f$ *such that* $G(df_t)$ *is* C^0*-close to* G_t.

By Proposition 3.12 we can reduce Theorem 5.1 to the following statement.

Theorem 5.2 *Let* $G_t : L \to \Lambda(M)$ *be a graphical simple rotation of a wrinkled Lagrangian or Legendrian embedding* $f : L \to M$. *Then there exists a compactly supported exact homotopy of wrinkled Lagrangian or Legendrian embeddings* $f_t : L \to M$, $f_0 = f$ *such that* $G(df_t)$ *is* C^0*-close to* G_t.

The parametric version of Theorem 5.1 reads as follows.

Theorem 5.3 *Let* $G_t^z : L \to \Lambda(M)$ *be a family of tangential rotations of regular Lagrangian or Legendrian embeddings* $f^z : L \to M$ *parametrized by a compact manifold* Z *such that* $G_t^z = G(df^z)$ *for* $z \in Op(\partial Z)$. *Then there exists a family of compactly supported exact homotopies of wrinkled Lagrangian or Legendrian embeddings* $f_t^z : L \to M$, $f_0^z = f^z$ *such that* $G(df_t^z)$ *is* C^0*-close to* G_t^z *and such that* $f_t^z = f^z$ *for* $z \in Op(\partial Z)$.

As in the non-parametric case, by Proposition 3.14 we can reduce Theorem 5.3 to the following statement.

Theorem 5.4 *Let* $G_t^z : L \to \Lambda(M)$ *be a family of graphical simple rotations of wrinkled Lagrangian or Legendrian embeddings* $f^z : L \to M$ *parametrized by a compact manifold* Z *such that* $G_t^z = G(df^z)$ *for* $z \in Op(\partial Z)$. *Then there exists a family of compactly supported exact homotopies of wrinkled Lagrangian or Legendrian embeddings* $f_t^z : L \to M$, $f_0^z = f^z$ *such that* $G(df_t^z)$ *is* C^0*-close to* G_t^z *and such that* $f_t^z = f^z$ *for* $z \in Op(\partial Z)$.

The proof of Theorems 5.2 and 5.4 consists of two steps. The first step is the construction of a local wrinkling model, which we carry out in Sect. 5.2. The second step is to combine this local wrinkling model with the wiggling results established in Sect. 4 to obtain the desired global approximation. We carry out this second step in Sect. 5.3.

5.2 Local wrinkling model

We begin by describing the local model for the oscillating function that will generate the wrinkles. This is essentially the same local model used by Eliashberg and Mishachev in [15]. In fact, our local wrinkling model for Lagrangians and Legendrians is obtained from theirs by simply integrating and differentiating the formulae, just like we did in Sect. 2 with the definition of wrinkled Lagrangian and Legendrian embeddings.

The basic geometric idea behind the construction is quite straightforward. One wishes to wrinkle the Lagrangian or Legendrian submanifold back and forth so that the wrinkles are parallel to the rotating planes $G_t(q)$. Since we model the wrinkles on a highly oscillating function, the Gauss map of the resulting wrinkled embedding gives an arbitrarily good approximation of G_t. There is a delicate part of the construction regarding the embryos of the zigzags because the oscillating function is forced to have a derivative with the 'wrong sign' in some neighboring region. However, we will impose bounds on the size of this bad derivative to ensure that its effect is not significant.

Construction 5.5 (The oscillating function) First, we fix some notation. We will localize our problem from a general n-dimensional manifold L to the unit cube $I^n = [-1, 1]^n \subset \mathbb{R}^n$. A point $q = (q_1, \dots, q_n) \in I^n$ will be written as $q = (\hat{q}, q_n)$, where $\hat{q} = (q_1, \dots, q_{n-1})$. We will consider rotations which are simple with respect to the (constant) hyperplane field $H^{n-1} \subset TI^n$ spanned by the vectors $\partial/\partial q_1, \dots, \partial/\partial q_{n-1}$. Hence the last coordinate q_n will play a special role in our discussion. We will also need a time parameter, which will be denoted by t. Sometimes it will be convenient to consider time as another spatial parameter, in which case we will think of the domain of our local model as $[0, 1] \times I^n$.

Consider the family of curves $Z_s \subset \mathbb{R}^2$, $s \in \mathbb{R}$, given by parametric equations

$$x_s(u) = \frac{15}{8} \int_0^u (w^2 - s)^2 dw, \qquad y_s(u) = \frac{1}{2}(u^3 - 3su).$$

The curve Z_s is a graph of a continuous function $z_s : \mathbb{R} \to \mathbb{R}$ which is smooth for $s < 0$ and smooth on $\mathbb{R} \setminus \{-s^{5/2}, s^{5/2}\}$ for $s \geq 0$, where we

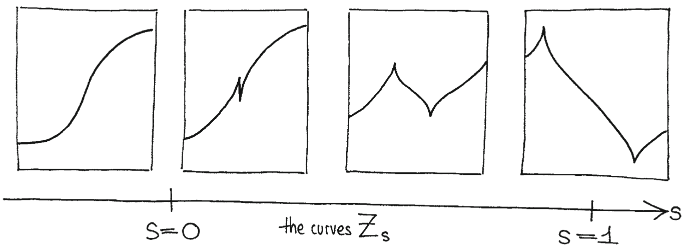

Fig. 29 The family of curves Z_s gives the local model for the birth/death of semi-cubical zig-zags

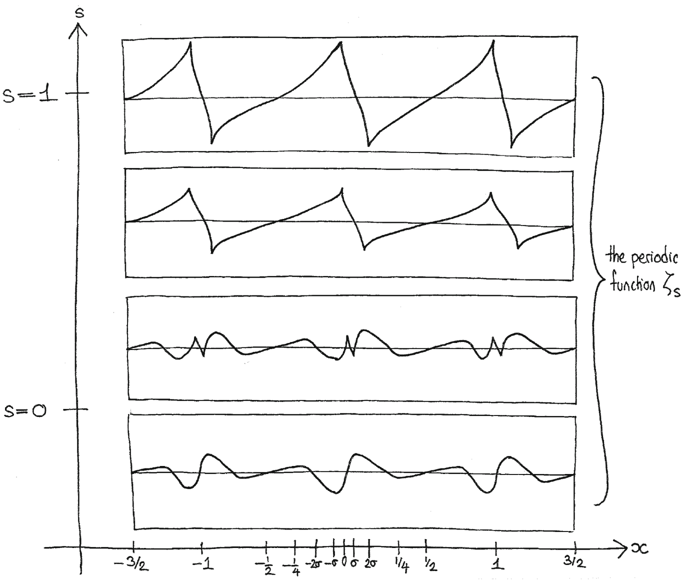

Fig. 30 The family ζ_s. Observe that for $s = 0$ the derivative $d\zeta_0/dx$ blows up near zero but is everywhere bounded below by $-\alpha$, where the parameter α can be taken to be arbitrarily small. This lower bound also holds everywhere for $s < 0$ and outside of $[-\sigma s^{5/2}, \sigma s^{5/2}]$ for $s > 0$

note that $x_s(\pm\sqrt{s}) = \pm s^{5/2}$. See Fig. 29 for an illustration. We note that the constants $15/8$ and $1/2$ are chosen for convenience in the calculation but are otherwise immaterial.

Remark 5.6 Observe that the composition $y_s(u) = z_s\big(x_s(u)\big)$ is smooth for all $s \in \mathbb{R}$.

Let $\sigma, \alpha > 0$ be small and choose an odd 1-periodic family of functions $\zeta_s : \mathbb{R} \to \mathbb{R}$, $s \in [-1, 1]$, illustrated in Fig. 30, which satisfies the following properties.

$$\zeta_s(x) \begin{cases} = z_s(\frac{x}{\sigma}) \text{ for } x \in Op\left(\left[-\sigma s^{5/2}, \sigma s^{5/2}\right]\right), s \in [0,1], \\ = z_s(\frac{x}{\sigma}) \text{ for } x \in Op(0), s \in [-1,0], \\ \geq 0 \quad\quad \text{for x } \in \left[-\frac{1}{2}, -\frac{1}{4}\right], s \in [-1,1], \\ \leq 0 \quad\quad \text{for x } \in \left[\frac{1}{4}, \frac{1}{2}\right], s \in [-1,1]. \end{cases}$$

$$\frac{d\zeta_1}{dx}(x) \begin{cases} \leq -\frac{4}{\sigma} \quad \text{for } x \in (-\sigma, \sigma), \\ \geq 1 \quad\quad \text{for x } \in [-2\sigma, -\sigma) \cup (\sigma, 2\sigma], \\ \in [1,2] \text{ for x } \in \left[-\frac{1}{2}, -2\sigma\right] \cup [2\sigma, \frac{1}{2}]. \end{cases}$$

$$\frac{d\zeta_s}{dx}(x) \begin{cases} \leq -\frac{4}{\sigma} \quad\quad \text{for } x \in (-\sigma s^{5/2}, \sigma s^{5/2}), s \in (0,1], \\ \geq -\alpha \quad\quad \text{for } x \in [-2\sigma, -\sigma s^{5/2}) \cup (\sigma s^{5/2}, 2\sigma], s \in (0,1], \\ \geq -\alpha \quad\quad \text{for } x \in [-2\sigma, 2\sigma], s \in [-1,0], \\ \in [-\alpha, 2] \text{ for } x \in \left[-\frac{1}{2}, -2\sigma\right] \cup \left[2\sigma, \frac{1}{2}\right], s \in [-1,1]. \end{cases}$$

Let $D^n = \{x \in \mathbb{R}^n : ||x|| \leq 1\}$ denote the closed unit $n-$dimensional disk. We now use the family ζ_s to define a model $\xi = \xi_{\sigma,\alpha,\gamma,\delta,N} : D^n(t,\hat{q}) \times [-1,1](q_n) \to \mathbb{R}$ which like ζ_s depends on $\sigma, \alpha > 0$ but also depends on three more parameters $\gamma, \delta > 0$ and $N \in \mathbb{N}$. The parameters $\sigma, \alpha, \gamma, \delta, 1/N$ are all taken to be small (in particular we demand that they are all < 1), but it is the relative smallness between the parameters that will play a crucial role in what follows.

Fix a non-increasing function $\eta : [0,1] \to \mathbb{R}$ such that

- $\eta(x) = 1$ for $x \in [0, 1-2\delta]$,
- $\eta(x) = -\delta$ for $x \in [1-\delta, 1]$.

Fix a non-increasing cutoff function $\rho : [0,1] \to \mathbb{R}$ such that

- $\rho(x) = 1$ for $x \in [0, 1-\delta]$
- $\rho(x) = 0$ for x near 1.

Fix also another non-increasing cutoff function $\psi : [0,1] \to [0,1]$ such that

- $\psi(x) = 1$ for $x \in \left[0, 1 - \frac{1}{4N+2}\right]$,
- $\psi(x) = 0$ for x near 1.

We define our oscillating model ξ by the following formula, see Fig. 31 for an illustration.

$$\xi(t,q) = \gamma \; \rho\big(||(t,\hat{q})||\big) \; \psi\big(|q_n|\big) \; \zeta_{\eta(||(t,\hat{q})||)}\left(\frac{2N+1}{2} q_n\right),$$
$$(t,\hat{q}) \in D^n, \quad q_n \in [-1,1].$$

Given $t \in [0,1]$, $q \in I^n$ and $b, c > 0$, let $C = C(t,q,b,c)$ denote the box

$$C = (t,q) + \big(bD^n\big) \times [-c,c] \subset \mathbb{R}^{n+1}(t,\hat{q},q_n)$$

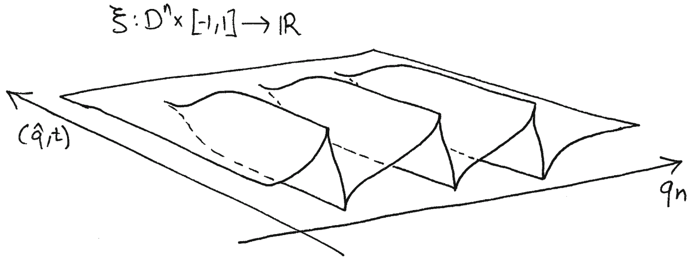

Fig. 31 One-half of the oscillating function ξ

which is a copy of $D^n \times [-1, 1]$ centered at (t, q) and scaled by b and c in the $(t, \hat{q})$ and q_n directions respectively. Let $\psi : C \to D^n \times [-1, 1]$ be the obvious diffeomorphism obtained by translating and rescaling. Define $\xi_C = \xi \circ \psi : C \to \mathbb{R}$. The oscillating function ξ_C also depends on the parameters $\sigma, \alpha, \gamma, \delta$ and N. We will call $\psi^{-1}\big(D^n \times 1\big)$ and $\psi^{-1}\big(D^n \times -1\big)$ the top and bottom of the box C respectively. We will also need to consider the slightly smaller boxes

$$\widehat{C} = C\left(t, q, (1-\delta)b, \left(1 - \frac{1}{4N+2}\right)c\right) \subset C$$
$$\text{and} \quad \widetilde{C} = C\left(t, q, (1-2\delta)b, \left(1 - \frac{1}{4N+2}\right)c\right) \subset \widehat{C}.$$

Observe that ξ_C has wild oscillations on $\widetilde{C}$ which die out on $\widehat{C} \backslash \widetilde{C}$, so that ξ_C is smooth on $C \backslash \widehat{C}$ and $\xi_C = 0$ on $Op(\partial C)$.

Finally, we modify our local model ξ to make it Lagrangian. We do this by integrating and differentiating as in the definition of wrinkled Lagrangian embeddings. Define $\ell : D^n(t, \hat{q}) \times [-1, 1](q_n) \to T^*\mathbb{R}^n(q, p)$ by the formula

$$\ell(t, q) = \left(q_1, \ldots, q_n, \frac{\partial K}{\partial q_1}, \ldots, \frac{\partial K}{\partial q_{n-1}}, \xi\right),$$
$$\text{where} \quad K(t, q) = \int_{-1}^{q_n} \xi(t, \hat{q}, u) du \tag{1}$$

Observe that ℓ is defined in terms of ξ, hence also depends on the parameters $\sigma, \alpha, \gamma, \delta$ and N. Observe also that ξ is odd in the q_n variable, hence $K = 0$ on $Op\big(\partial(D^n \times [-1, 1])\big)$. It follows that ℓ has a Legendrian lift (ℓ, K) which agrees with the zero section on $Op\big(\partial(D^n \times [-1, 1])\big)$.

Given any box C we can similarly define a translated and scaled version ℓ_C of ℓ which has support in C. This completes the construction of our local wrinkling model.

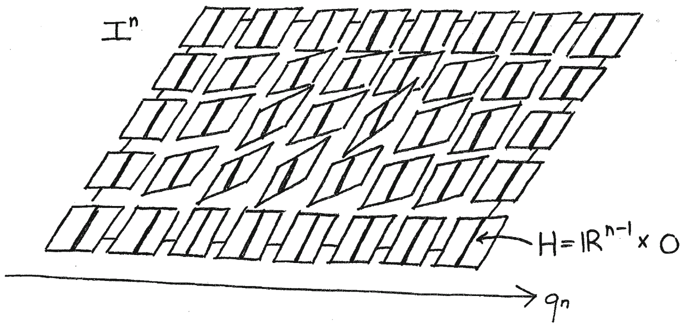

Fig. 32 A tangential rotation which is quasi-graphical and simple with respect to H

Remark 5.7 The function ξ is not smooth and hence ℓ is also not smooth. However, ξ can be smoothly reparametrized and therefore so can ℓ. We will revisit this nuance later on but it will not cause us any trouble.

We are now ready to state and prove the local wrinkling lemma. Note that a tangential rotation $G_t : I^n \to \Lambda(T^*I^n)$ of the inclusion of the zero section $i : I^n \hookrightarrow T^*I^n$ is simple with respect to the hyperplane field $H = \operatorname{span}(\partial/\partial q_1, \ldots, \partial/\partial q_{n-1}) \subset TI^n$ if it can be written as

$$G_t = \operatorname{span}\left(\frac{\partial}{\partial q_1}, \ldots, \frac{\partial}{\partial q_{n-1}}, \cos(\lambda_t)\frac{\partial}{\partial q_n} + \sin(\lambda_t)\frac{\partial}{\partial p_n}\right)$$

for some angle function $\lambda_t : I^n \to \mathbb{R}$. According to our previous definition we say that G_t is graphical when $\operatorname{im}(\lambda_t) \subset (-\pi/2, \pi/2)$. We will say that G_t is quasi-graphical when $\operatorname{im}(\lambda_t) \subset (-\pi, \pi)$ (Fig. 32).

Lemma 5.8 (Local wrinkling for Lagrangians) *Let $G_t : I^n \to \Lambda(T^*I^n)$ be a tangential rotation of the zero section $i : I^n \hookrightarrow T^*I^n$ which is quasi-graphical and simple with respect to H and such that $G_t = G(di)$ on $Op(\partial I^n)$. Then there exists an exact homotopy of wrinkled Lagrangian embeddings $f_t : I^n \to T^*I^n$, $f_0 = i$, such that the following properties hold.*

- *$G(df_t)$ is C^0-close to G_t.*
- *$f_t = i$ on $Op(\partial I^n)$.*

Proof Let $\tau > 0$ be small. We will be precise about exactly how small we need τ to be later on. Wrinkling is dangerous and unnecessary where λ_t is close to zero, so we will first use our oscillating model ℓ to define a similar model which does not oscillate on the subset of $[0, 1] \times I^n$ in which $|\lambda_t| < \tau$.

Remark 5.9 Although we want to think of time as a spatial parameter, observe that $\lambda_t \neq 0$ on the boundary face $1 \times I^n \subset \partial([0, 1] \times I^n)$, so we are not quite in the relative setting. To remedy this, we extend the time interval from [0, 1] to [0, 2] by setting $\lambda_t = \lambda_{2-t}$ for $t \in [1, 2]$. We can then work with the box $[0, 2] \times I^n$ as our local model, which has the advantage that $\lambda_t = 0$ on $Op\big(\partial([0, 2] \times I^n)\big)$. We can later restrict back to only considering times $t \in [0, 1]$ and forget about the rest.

Let $\Omega_\tau = \{(t, q) \in [0, 2] \times I^n : |\lambda_t(q)| > \tau\}$. We call a box $C = C(t, q, b, c) \subset [0, 2] \times I^n$ special if $|\lambda_t(q)| < 2\tau$ for (t, q) near the top and bottom of C. Choose special boxes $C_1, \ldots, C_m \subset [0, 2] \times I^n$ which are contained in Ω_τ and such that the smaller boxes $\widetilde{C}_1, \ldots, \widetilde{C}_m$ are still special and cover $\Omega_{2\tau}$. This can be achieved if δ is sufficiently small and N is sufficiently big. Write ψ_j for the parametrizing diffeomorphisms $\psi_j : C_j \to D^n \times [-1, 1]$ as above. We can assume that the sets $\psi_j^{-1}\big(D^n \times ([-1, 1] \cap \mathbb{Q})\big) \subset [0, 2] \times I^n$ are disjoint. Therefore for each integer N there exists a number $\sigma(N) > 0$ such that for all $\sigma < \sigma(N)$ the subsets

$$\psi_j^{-1}\left(D^n \times \left[\frac{2k}{2N+1} - \widetilde{\sigma}, \frac{2k}{2N+1} + \widetilde{\sigma}\right]\right), \quad \widetilde{\sigma} = \frac{4\sigma}{2N+1}, \quad -N \leq k \leq N,$$

are also disjoint. When we let $N \to \infty$ below, we will let $\sigma \to 0$ accordingly so that we always have $\sigma < \sigma(N)$.

For each box $C_j \subset \Omega_\tau$ we have an oscillating Lagrangian model ℓ_{C_j}. Let $\text{sign}(j) = \text{sign}(\lambda_t|_{C_j}) \in \{\pm 1\}$. Define the Lagrangian oscillating model w_t adapted to G_t by setting $w_t(q) = \sum_j \text{sign}(j)\ell_{C_j}(t, q)$. More precisely, we set

$$w_t(q) = \left(q_1, \ldots, q_n, \frac{\partial H_t}{\partial q_1}, \ldots, \frac{\partial H_t}{\partial q_{n-1}}, \sum_j \text{sign}(j)\xi_{C_j}\right),$$
$$\text{where } H_t = \sum_j \int_{-1}^{q_n} \text{sign}(j)\xi_{C_j}(t, \hat{q}, u)du$$

Observe that $w_t = 0$ and $H_t = 0$ outside of Ω_τ. At this point we can restrict back to the time interval $[0, 1] \subset [0, 2]$, which is all that we really cared about.

Consider the function $F_t(q, p) = \frac{1}{2}\cot\big(\lambda_t(q)\big)p_n^2$. For each $t \in [0, 1]$ we consider F_t as an autonomous Hamiltonian function. Therefore F_t yields a Hamiltonian isotopy $\varphi_t^s : T^*\Omega_\tau \to T^*\mathbb{R}^n$ such that the vector field $X_t = \partial_s\varphi_t^s(q)|_{s=0}$ is the symplectic dual of $dF_t(q) = \cot\big(\lambda_t(q)\big)p_n dp_n - \frac{1}{2}\text{cosec}^2\big(\lambda_t(q)\big)p_n^2 d\lambda_t(q)$. Hence we have

$$X_t(q, p) = \cot\big(\lambda_t(q)\big)p_n\frac{\partial}{\partial q_n} + \frac{1}{2}\text{cosec}^2\big(\lambda_t(q)\big)p_n^2\sum_{j=1}^n \frac{\partial\lambda_t}{\partial q_j}\frac{\partial}{\partial p_j}.$$

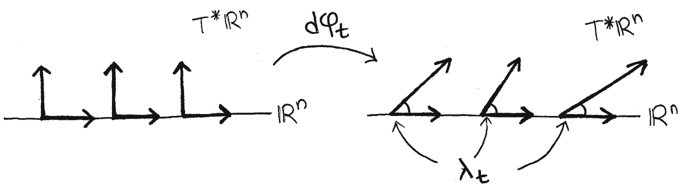

Fig. 33 Along the zero section $\mathbb{R}^n \subset T^*\mathbb{R}^n$ we have $d\varphi_t(\partial/\partial q_n) = \partial/\partial q_n$ and $d\varphi_t(\partial/\partial p_n) = \cot(\lambda_t)\partial/\partial q_n + \partial/\partial p_n$

It follows by explicit computation that

$$\varphi_t^s(q, p) = \big(\hat{q},\ q_n + \cot\big(\lambda_t(q)\big)p_n s\,,\ \ p_1 + \frac{1}{2}\mathrm{cosec}^2\big(\lambda_t(q)\big)p_n^2\frac{\partial \lambda_t}{\partial q_1}s\ ,\ \ldots\ ,\ p_n + \frac{1}{2}\mathrm{cosec}^2\big(\lambda_t(q)\big)p_n^2\frac{\partial \lambda_t}{\partial q_n}s\bigg).$$

We set $\varphi_t = \varphi_t^1$, which is well defined for all $t \in [0, 1]$ on $Op(\Omega_\tau) \subset T^*\mathbb{R}^n$. Note that φ_t is itself a Hamiltonian isotopy. Note also that $\varphi_t = id$ on $\Omega_\tau \subset T^*\Omega_\tau$ since $p_1, \ldots, p_n = 0$ on the zero section. Note moreover that on Ω_τ we have (Fig. 33)

$$\frac{\partial \varphi_t}{\partial q_j} = \frac{\partial}{\partial q_j} \quad \text{for } j = 1, \ldots, n, \qquad \frac{\partial \varphi_t}{\partial p_j} = \frac{\partial}{\partial p_j} \quad \text{for } j < n$$
$$\text{and} \quad \frac{\partial \varphi_t}{\partial p_n} = \cot(\lambda_t)\frac{\partial}{\partial q_n} + \frac{\partial}{\partial p_n}.$$

Hence in particular on Ω_τ we have

$$d\varphi_t\left(\mathrm{span}\left(\frac{\partial}{\partial q_1}, \ldots, \frac{\partial}{\partial q_{n-1}}, \frac{\partial}{\partial p_n}\right)\right) = G_t.$$

Set $f_t = \varphi_t \circ w_t$. We recall from Remark 5.7 that each ℓ_{C_j} is not smooth, hence w_t is not smooth, hence the same is true for f_t. However, we can pre-compose w_t with a reparametrization of the domain so that w_t and hence also f_t is smooth. Note moreover that this reparametrization can be taken to be C^0-smal and supported in an arbitrarily neighborhood of the wrinkles. Note finally that reparametrizing f_t doesn't change the image of f_t and therefore it also doesn't change the image of the Gauss map $G(df_t)$, which is what we actually care about. By abusing notation we will also use f_t to denote the reparametrized smooth map whenever this is convenient. See Figs. 34 and 35 for an illustration of f_t.

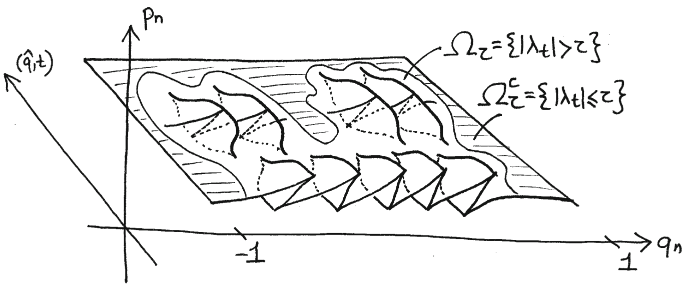

Fig. 34 The p_n-coordinate of the map f_t. The cusps are semi-cubic

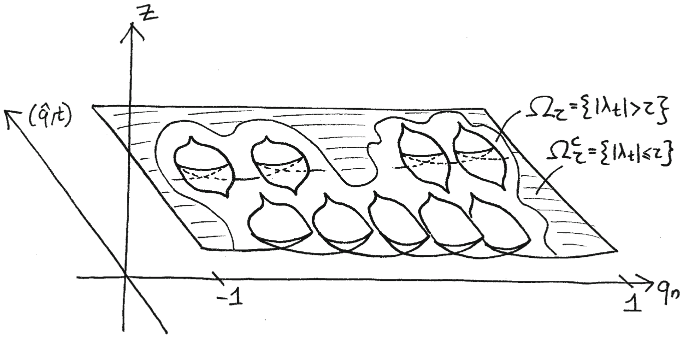

Fig. 35 The z-coordinate of the Legendrian lift of f_t. In other words, this is the Legendrian front of f_t. The cusps are semi-quintic

Claim 5.10 *For any $\varepsilon > 0$ we can choose parameters $\tau, \delta, \sigma, \alpha, \gamma$ and N so that $dist_{C^0}(G(df_t), G_t) < \varepsilon$.*

We recall the parameters at play. The game is all about controlling the different rates at which the parameters tend to zero or infinity, so it will be important to be precise in the interdependence of the parameters and in the order of quantifiers.

- τ is the cutoff angle of G_t under which we will perform no wrinkling.
- δ is proportional to the width of the shell between a box C and the smaller box $\widetilde{C}$.
- $1/\sigma$ is the order of magnitude of ζ_s' on the regions where it is large and negative (inside the wrinkles).
- α controls the magnitude of the 'bad' negative derivative ζ_s' when the wrinkles die out.
- γ is the height of the oscillating model ξ.
- N is proportional to the number of wrinkles in ξ.

We begin by fixing $\varepsilon > 0$ arbitrarily small. To choose τ, observe that

$$d\varphi_t\left(\frac{\partial}{\partial q_n}+\beta\frac{\partial}{\partial p_n}\right)=\frac{\partial}{\partial q_n}+\beta\left(\cot(\lambda_t)\frac{\partial}{\partial q_n}+\frac{\partial}{\partial p_n}\right), \qquad \beta\in\mathbb{R}$$

and hence if $\operatorname{sign}(\beta) = \operatorname{sign}(\lambda_t)$, then the scalar product of $\partial/\partial q_n$ and $d\varphi_t(\partial/\partial q_n + \beta\partial/\partial p_n)$ is positive and moreover we have

$$\measuredangle\left(\frac{\partial}{\partial q_n},\ d\varphi_t\left(\frac{\partial}{\partial q_n}+\beta\frac{\partial}{\partial p_n}\right)\right) < |\lambda_t|.$$

Recall that on the subset $\Omega_\tau\backslash\Omega_{2\tau}$ we have $\tau < |\lambda_t| \leq 2\tau$. Suppose that $\tau < \varepsilon/4$. It follows that if $\operatorname{sign}(\beta) = \operatorname{sign}(\lambda_t)$, then

$$\measuredangle\left(\frac{\partial}{\partial q_n},\ d\varphi_t\left(\frac{\partial}{\partial q_n}+\beta\frac{\partial}{\partial p_n}\right)\right) < 2\tau < \frac{\varepsilon}{2} \quad \text{on } \Omega_\tau\backslash\Omega_{2\tau}.$$

Once $\tau < \varepsilon/4$ is fixed, we choose δ small enough so that the construction of $w_t = \sum_j \operatorname{sign}(j)\ell_{C_j}$ (which depends implicitly on τ) is possible. The other parameters must be chosen somewhat more judiciously. Our first task is to understand the geometry of the initial local model $(s,u) \mapsto \big(x_s(u), y_s(u)\big)$ in order to control the error produced when we modify the model to make it Lagrangian.

Consider the Lagrangian version in $T^*\mathbb{R}^2 = \mathbb{R}^4(q_1, q_2, p_1, p_2)$ given by the formula

$$m(s,u) = \Big(s, x_s(u), r_s(u), y_s(u)\Big) \in T^*\mathbb{R}^2,$$
$$r_s(u) = \int_0^u \partial_s\big(y_s(u)\big)\partial_u\big(x_s(u)\big) - \partial_u\big(y_s(u)\big)\partial_s\big(x_s(u)\big)du,$$

where we recall that

$$x_s(u) = \frac{15}{8}\int_0^u (w^2-s)^2dw, \qquad y_s(u) = \frac{1}{2}(u^3-3su).$$

We also have the corresponding scaled version

$$m_{\gamma,N}(s,u) = \Big(s, \frac{1}{N}x_s(u), \frac{\gamma}{N}r_s(u), \gamma y_s(u)\Big) \in T^*\mathbb{R}^2.$$

If $\gamma \to 0$ and $N \to \infty$ in such a way that $N\gamma \to \infty$, then the Gauss map $G(dm_{\gamma,N})$ converges (on compact subsets of the (s,u) plane) to the distribution spanned by the vectors $\partial/\partial q_1 = (1,0,0,0)$ and $\partial/\partial p_2 = (0,0,0,1)$. The proof is the following explicit computation.

It will we convenient to carry out our calculations in terms of the function $F(s,u) = \frac{1}{3}(u^3 - 3su)$ and its derivative $F_u(s,u) = u^2 - s$. Note that the zero set $\{F_u = 0\}$ is precisely the wrinkling locus of m. We compute:

$$\begin{aligned}
\partial_u\big(x_s(u)\big) &= \frac{15}{8}F_u^2, \qquad \partial_s\big(x_s(u)\big) = -\frac{15}{4}F\\
\partial_u\big(y_s(u)\big) &= \frac{3}{2}F_u, \qquad \partial_s\big(y_s(u)\big) = -\frac{3}{2}u\\
\partial_u\big(r_s(u)\big) &= \frac{15}{8}F_u\Big(-\frac{3}{2}uF_u + 3F\Big) = -\frac{15}{16}F_u(u^3+3su).\\
\frac{\partial m_{\gamma,N}}{\partial s} &= \Big(1,\ \frac{1}{N}\partial_s\big(x_s(u)\big),\ \frac{\gamma}{N}\partial_s\big(r_s(u)\big),\ \gamma\partial_s\big(y_s(u)\big)\Big)\\
&\to \big(1,0,0,0\big) \text{ as } \gamma\to 0,\ N\to\infty,\\
\frac{\partial m_{\gamma,N}}{\partial u} &= \Big(0,\ \frac{1}{N}\frac{15}{8}F_u^2,\ -\frac{\gamma}{N}\frac{15}{16}(u^3+3su)F_u,\ \gamma\frac{3}{2}F_u\Big)\\
&= -\gamma F_u\Big(0,\ \frac{1}{N\gamma}F_u,\ -\frac{1}{N}\frac{15}{16}(u^3+3su),\ \frac{3}{2}\Big)
\end{aligned}$$

and hence provided that $N\gamma\to\infty$ we have

$$\operatorname{span}\Big(\frac{\partial m_{\gamma,N}}{\partial s},\ \frac{\partial m_{\gamma,N}}{\partial u}\Big) \longrightarrow \operatorname{span}\Big(\frac{\partial}{\partial q_1},\ \frac{\partial}{\partial p_2}\Big).$$

With some minor modifications we can extend our computations to the scaled $n-$dimensional model for the Lagrangian wrinkle as it appears in ℓ (see equation 1).

$$\begin{aligned}
(t,q) \mapsto \Big(&q_1,\ldots,q_{n-1},\frac{\sigma x_\eta(q_n)}{2N+1},\ \frac{\sigma\gamma r_\eta(q_n)}{2N+1}\frac{\partial||(t,\hat{q})||}{\partial q_1}\eta',\ \ldots,\\
&\frac{\sigma\gamma r_\eta(q_n)}{2N+1}\ \frac{\partial||(t,\hat{q})||}{\partial q_{n-1}}\eta',\ \gamma y_\eta(q_n)\Big)
\end{aligned}$$

where $\eta = \eta\big(||(t,\hat{q})||\big)$. Indeed, the only difference comes from the terms $\partial_j\eta = \eta'\partial_j||(t,\hat{q})||$ for $j<n$ and their partial derivatives, which give an error that tends to zero as $\gamma\to 0$ and $N\to\infty$. The conclusion is that provided we have $N\gamma\to\infty$, the Gauss map converges to the distribution

$$V = \operatorname{span}\Big(\frac{\partial}{\partial q_1},\ldots,\frac{\partial}{\partial q_{n-1}},\frac{\partial}{\partial p_n}\Big).$$

Recall that we must ensure $\sigma < \sigma(N)$ so that the singularity loci $\Sigma(\ell_{C_j})\subset [0,1]\times I^n$ are disjoint. Hence if we let $N\to\infty$, then we must also allow for $\sigma\to 0$. But this only helps us in the above computation so there is no issue.

Consider next the oscillating model ℓ defined above. Let $\Sigma \subset [-1, 1] \times I^n$ be the locus on which ℓ is not smooth. The set Σ consists of a disjoint union of spheres with cuspidal equators. Let E be the compact region bounded by Σ. If $\gamma \to 0$ and $N \to \infty$ so that $N\gamma \to \infty$, then the above computations show that on $Op(E)$ the Gauss map of ℓ converges to the distribution V. In the complement of $Op(E)$, the model ℓ is smooth and for $j < n$ we have $\partial\ell/\partial q_j \to \partial/\partial q_j$ as $\gamma \to 0$. On the subset $B = [-1+2\delta, 1-2\delta]^n \times [-1 + \frac{1}{4N+2}, 1 - \frac{1}{4N+2}]$ the Gauss map of ℓ converges to V, indeed on the remaining part $B \backslash Op(E)$ the derivative $dp_n(\partial\ell/\partial q_n) = \partial\xi/\partial q_n$ is strictly positive and scales by $N\gamma$ while $dp_j(\partial\ell/\partial q_n)$ scales by γ for $j < n$. On $I^n \backslash B$ we cannot control $\partial\ell/\partial q_n$ so precisely but we assert that outside of $Op(E)$ there still holds the following lower bound:

$$dp_n(\partial\ell/\partial q_n) = \partial\xi/\partial q_n \geq -(N+1)\gamma\alpha.$$

To confirm this assertion, we compute

$$\frac{\partial\xi}{\partial q_n} = \gamma\rho(t, \hat{q})\Bigg(\operatorname{sign}(q_n)\psi'(|q_n|)\zeta_{\eta(t,\hat{q})}\left(\frac{2N+1}{2}q_n\right) + \frac{2N+1}{2}\psi(|q_n|)\zeta'_{\eta(t,\hat{q})}\left(\frac{2N+1}{2}q_n\right)\Bigg).$$

Since $\psi' \leq 0$ and $\operatorname{sign}(\zeta_{\eta(t,\hat{q})}(\frac{2N+1}{2}q_n)) = -\operatorname{sign}(q_n)$ in the region where $\psi' \neq 0$, the first term is always non-negative. For the second term we use our assumption that $\zeta'_s \geq -\alpha$ and the desired inequality follows.

We deduce from this inequality that if we let $\gamma, \alpha \to 0$ and $N \to \infty$ so that $N\gamma \to \infty$ and $N\gamma\alpha \to 0$, then on the complement of $Op(E)$ we have $\liminf dp_n(\partial\ell/\partial q_n) \geq 0$. Of course we also still have $dq_j(\partial\ell/\partial q_n) = 0$ for $j < n$, $dq_n(\partial\ell/\partial q_n) = 1$ and $dp_j(\partial\ell/\partial q_n) \to 0$ as $\gamma \to 0$.

Next we proceed to study the model $w_t = \sum_j \operatorname{sign}(j)\ell_{C_j}$ which is adapted to our rotation G_t. Assume first for simplicity that $\lambda_t \geq 0$, so that $\operatorname{sign}(j) = 1$ for all j. Let $\widetilde{\Sigma} \subset [0, 1] \times I^n$ be the non-smooth locus of w_t. The set $\widetilde{\Sigma}$ is again a disjoint union of spheres which have cuspidal equators. Let $\widetilde{E}$ be the compact subset bounded by $\widetilde{\Sigma}$. Note that $\widetilde{E} \subset \Omega_\tau$. On $\Omega_{2\tau} \backslash \widetilde{E}$ all the derivatives $\partial\xi_{C_j}/\partial q_n$ are bounded below by a positive constant times $-N\gamma\alpha$ and at each point there is at least one of them which is bounded below by a constant times $N\gamma$. This last assertion holds because the boxes $\widetilde{C}_j \subset C_j$ cover $\Omega_{2\tau}$. Inside $\widetilde{E}$ all the derivatives $\partial\xi_{C_j}/\partial q_n$ are bounded above by a positive constant times $N\gamma$ and at each point there is exactly one derivative $\partial\xi_{C_j}/\partial q_n$ for which is bounded above by a constant times $-N\gamma/\sigma$. This last derivative corresponds to the ℓ_{C_j} whose non-smooth locus bounds the component of $\widetilde{E}$ containing the point we're looking at. We recall that we are letting $\sigma \to 0$ with the only requirement that $\sigma < \sigma(N)$. Hence if $N \to \infty$ and $\gamma, \alpha, \sigma \to 0$ in such a way that this condition holds and if additionally we have $N\gamma \to \infty$ and $N\gamma\alpha \to 0$, then on the region $\Omega_{2\tau} \cup Op(\widetilde{E})$ the Gauss map of w_t converges to the distribution V and on and on $\Omega_\tau \setminus \big(\Omega_{2\tau} \cup Op(\widetilde{E})\big)$ we know that $\partial w_t/\partial q_j \to \partial/\partial q_j$ for $j < n$ and that $\partial w_t/\partial q_n$ gets arbitrarily close to the sector

$$\mathcal{C} = \operatorname{span}\left\{\frac{\partial}{\partial q_n} + \beta\frac{\partial}{\partial p_n} : \ \beta \geq 0\right\} \subset T(T^*\mathbb{R}^n)|_{\mathbb{R}^n}.$$

Consider next the general case where we don't assume that $\operatorname{sign}(j) = 1$ for all j. Since $\Omega_\tau = \{\lambda_t > \tau\} \cup \{\lambda_t < -\tau\}$ is a disjoint union, we can repeat the above reasoning on each component and reach the same conclusion, provided that we modify that definition of the subset $\mathcal{C}$ as follows

$$\mathcal{C} = \operatorname{span}\left\{\frac{\partial}{\partial q_n} + \beta\frac{\partial}{\partial p_n} : \ \operatorname{sign}(\beta) = \operatorname{sign}(\lambda_t)\right\} \subset T(T^*\mathbb{R}^n)|_{\Omega_\tau}.$$

We now return to the wrinkled Lagrangian embedding $f_t = \varphi_t \circ w_t$. Recall that along the zero section the linear symplectic isomorphism $d\varphi_t$ is the map which sends

$$\frac{\partial}{\partial q_j} \mapsto \frac{\partial}{\partial q_j}, \ \ j = 1, \dots, n, \qquad \frac{\partial}{\partial p_j} \mapsto \frac{\partial}{\partial p_j}, \ \ j < n$$
$$\text{and} \quad \frac{\partial}{\partial p_n} \mapsto \cot(\lambda_t)\frac{\partial}{\partial q_n} + \frac{\partial}{\partial p_n},$$

so that $d\varphi_t(V) = G_t$ along the zero section. Recall also that we chose $\tau = \tau(\varepsilon)$ so that on $\Omega_\tau \setminus \Omega_{2\tau}$ we have $\measuredangle\big(d\varphi_t(v), \partial/\partial q_n\big) < \varepsilon/2$ for all $v \in \mathcal{C}$. Under the above convergence assumptions it follows that we have $\limsup \measuredangle(\partial f_t/\partial q_n, \partial/\partial q_n) \leq \varepsilon/2$ on $\Omega_\tau \setminus \big(\Omega_{2\tau} \cup Op(\widetilde{E})\big)$ and hence also $\limsup \operatorname{dist}(G(df_t), T\mathbb{R}^n) \leq \varepsilon/2$. Therefore $\limsup \operatorname{dist}\big(G(df_t), G_t\big) \leq \operatorname{dist}\big(G(df_t), T\mathbb{R}^n\big) + \operatorname{dist}\big(T\mathbb{R}^n, G_t\big) < \varepsilon/2 + 2\tau < \varepsilon$ on $\Omega_\tau / \big(\Omega_{2\tau} \cup Op(\widetilde{E})\big)$. Outside of Ω_τ we have $\operatorname{dist}\big(G(df_t), G_t) = \operatorname{dist}\big(T\mathbb{R}^n, G_t) < \tau < \varepsilon$. If we assume that on $\Omega_{2\tau} \cup Op(\widetilde{E})$ the Gauss map of w_t converges to the distribution V, then for f_t we have

$$G(df_t) \to d\varphi_t(V) = G_t \qquad \text{on } \Omega_{2\tau} \cup Op(\widetilde{E}).$$

Therefore to conclude the proof of Claim 5.10, and hence also of Lemma 5.8, it suffices to show that we can arrange that $\gamma, \alpha \to 0$ and $N \to \infty$ in such a way that $N\gamma \to \infty$ and $N\gamma\alpha \to 0$. This is clearly possible, for instance we can set $\gamma = N^{-1/2}$ and $\alpha = N^{-2/3}$. $\square$

The analogous result for Legendrians is stated and proved in the same way. Observe as in the Lagrangian case that a tangential rotation $G_t : I^n \to \Lambda\big(J^1(I^n, \mathbb{R})\big)$ of the inclusion of the zero section $i : I^n \hookrightarrow J^1(I^n, \mathbb{R})$ is simple with respect to the hyperplane field $H = \text{span}(\partial/\partial q_1, \ldots, \partial/\partial q_{n-1}) \subset TI^n$ if it can be written as

$$G_t = \text{span}\big(\partial/\partial q_1, \ldots, \partial/\partial q_{n-1}, \cos(\lambda_t)\partial/\partial q_n + \sin(\lambda_t)\partial/\partial p_n\big)$$

for some function $\lambda_t : I^n \to \mathbb{R}$. According to our previous definition we say that G_t is graphical when $\text{im}(\lambda_t) \subset (-\pi/2, \pi/2)$. We will say that G_t is quasi-graphical when $\text{im}(\lambda_t) \subset (-\pi, \pi)$.

Lemma 5.11 (Local wrinkling for Legendrians) *Let $G_t : I^n \to \Lambda\big(J^1(I^n, \mathbb{R})\big)$ be a tangential rotation of the zero section $i : I^n \hookrightarrow J^1(I^n, \mathbb{R})$ which is quasi-graphical and simple with respect to H and such that $G_t = G(di)$ on $Op(\partial I^n)$. Then there exists an exact homotopy of wrinkled Legendrian embeddings $f_t : I^n \to J^1(I^n, \mathbb{R})$, $f_0 = i$, such that the following properties hold.*

- *$G(df_t)$ is C^0-close to G_t.*
- *$f_t = i$ on $Op(\partial I^n)$.*

Proof We proceed exactly like we did in the proof of Lemma 5.8. The Legendrian model is simply given by the Legendrian lift $\widehat{\ell} = (\ell, K)$ of the Lagrangian model ℓ which exists because of the exactness condition $K = 0$ on $[-1, 1] \times Op(\partial I^n)$. □

The parametric versions read as follows. Note that we also localize the problem from a general m-dimensional parameter space Z to the unit cube $I^m = [-1, 1]^m$.

Lemma 5.12 (Parametric local wrinkling for Lagrangians) *Let $G_t^z : I^n \to \Lambda(T^*I^n)$ be a family of tangential rotations of the zero section $i : I^n \hookrightarrow T^*I^n$ parametrized by the unit cube I^m which are all quasi-graphical and simple with respect to H, such that $G_t^z = G(di)$ on $Op(\partial I^n)$ and such that $G_t^z = G(di)$ for $z \in Op(\partial I^m)$. Then there exists a family of exact homotopies of wrinkled Lagrangian embeddings $f_t^z : I^n \to T^*I^n$, $f_0^z = i$, such that the following properties hold.*

- *$G(df_t^z)$ is C^0-close to G_t^z.*
- *$f_t^z = i$ on $Op(\partial I^n)$.*
- *$f_t^z = i$ for $z \in Op(\partial I^m)$.*

Lemma 5.13 (Parametric local wrinkling for Legendrians) *Let $G_t^z : I^n \to \Lambda\big(J^1(I^n, \mathbb{R})\big)$ be a family of tangential rotations of the zero section $i : I^n \hookrightarrow$*

$J^1(I^n, \mathbb{R})$ parametrized by the unit cube I^m which are all quasi-graphical and simple with respect to H, such that $G^z_t = G(di)$ on $Op(\partial I^n)$ and such that $G^z_t = G(di)$ for $z \in Op(\partial I^m)$. Then there exists a family of exact homotopies of wrinkled Legendrian embeddings $f^z_t : I^n \to J^1(I^n, \mathbb{R})$, $f^z_0 = i$, such that the following properties hold.

- *$G(df^z_t)$ is C^0-close to G^z_t.*
- *$f^z_t = i$ on $Op(\partial I^n)$.*
- *$f^z_t = i$ for $z \in Op(\partial I^m)$.*

Lemmas 5.12 and 5.13 are proved in the same way as Lemmas 5.8 and 5.11, adapting our construction to the fibered case as in [15]. To be more precise, in the local model for the oscillating function ξ we replace the box $D^n(t, \hat{q}) \times [-1, 1](q_n)$ by the box $D^n(t, \hat{q}) \times [-1, 1](q_n) \times D^m(z)$ and set

$$\xi(t, q, z) = \gamma\, \rho(||z||)\, \rho\big(||(t, \hat{q})||\big)\, \psi\big(|q_n|\big)\, \zeta_{\eta(||(t,\hat{q})||)}\left(\frac{2N+1}{2} q_n\right),$$
$$(t, \hat{q}) \in D^n, \;\; q_n \in [-1, 1], \;\; z \in D^m.$$

The rest of the proof can then be repeated carrying the parameter $z \in D^m$ along for the ride.

5.3 Wrinkling the wiggles

We are now ready to prove that tangential rotations can be globally approximated by Gauss maps of wrinkled embeddings (Fig. 36).

Proof of Theorem 5.2 For simplicity we spell out the details only for the Lagrangian case, but the Legendrian case is entirely analogous. Let $G_t : L \to \Lambda(M)$ be a graphical simple rotation of a wrinkled Lagrangian embedding $f : L \to M$. Fix a Riemannian metric on L. Let Δ be a triangulation of L which is compatible with the wrinkles of f as in Sect. 4.3.

Set $K = \Delta^{n-1}$, the $(n-1)$-skeleton of Δ. By Theorem 4.11, there exists an exact homotopy of wrinkled Lagrangian embeddings $\widetilde{f}_t : L \to M$, $\widetilde{f}_0 = f$,

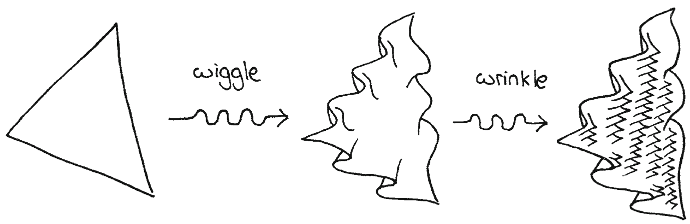

Fig. 36 The two-step process applied to a given simplex D. First we wiggle, then we wrinkle

(which is in fact an isotopy in the sense of Remark 4.13) and a tangential rotation $R_t : L \to \Lambda(M)$ of f such that the following properties hold.

- $G(d\widetilde{f_t})$ is C^0-close to G_t on $Op(K)$.
- $G(d\widetilde{f_t})$ is C^0-close to R_t on all of L.
- R_t is graphical and simple with respect to the same hyperplane field H as G_t.
- $\widetilde{f_t} = f$ and $R_t = G(df)$ outside of a slightly bigger neighborhood of K in L.

Take an open n-simplex D in Δ^n, so that $\widetilde{f_t}|_D : D \to M$ is an exact homotopy of regular Lagrangian embeddings. Suppose first that ∂D is disjoint from the wrinkling locus of f and take a slightly larger disk $\widetilde{D} \supset D$ in L. With respect to a Weinstein parametrization of a tubular neighborhood of $f(\widetilde{D})$ in M, the graphical homotopy $\widetilde{f_t}|_D$ corresponds to a homotopy $dh_t \circ F_t : D \to T^*\widetilde{D}$. Here $F_t : D \to \widetilde{D}$ is an isotopy (which may wiggle D outside of itself but by C^0-smallness can be assumed to satisfiy $F_t(D) \subset \widetilde{D}$) and $h_t : F_t(D) \to \mathbb{R}$ is a homotopy of real valued functions. The reader can review the proof of Theorem 4.5 (to which Theorem 4.11 reduces away from the wrinkling locus) to see where F_t and h_t come from.

We can assume that the hyperplane field H is almost constant along $\widetilde{D}$, in the following sense. For any $\varepsilon > 0$, there exists a $\delta > 0$ for which we can cover the compact subset where G_t is not identically $G(df)$ with a finite union of radius δ metric balls $B_j = B_\delta(q_j)$, $q_j \in L$, such that in exponential coordinates from q_j the hyperplane field $H|_{B_j}$ is ε-close to being constant (in other words, the angle between H at different points varies by less than ε). For balls which intersect the wrinkling locus of f (along which the hyperplane field H may jump discontinuously) we demand that the restriction of each hyperplane field is ε-close to being constant. For that fixed $\varepsilon > 0$ we can from the onset subdivide the triangulation Δ fine enough so that every n-simplex is contained in one of the balls B_j. We can then use the exponential coordinates on B_j to replace the hyperplane field H with a constant hyperplane field, modifying G_t and R_t accordingly, at the cost of a C^0-error uniformly proportional to ε. Since we were free to choose $\varepsilon > 0$ arbitrarily small, we can ensure that the error resulting from straightening out H is arbitrarily small and in particular smaller than whatever C^0-accuracy is desired in the conclusion of Theorem 5.2. Henceforth we shall use these coordinates for $\widetilde{D}$ and H on B_j, composed with an inclusion into the unit cube $I^n = [-1, 1]^n$ by a linear Euclidean isometry so that we have $\widetilde{D} \subset I^n$ and $H = \mathbb{R}^{n-1} \times 0$ as in our local wrinkling model. Since our local model is a relative construction, we won't really care about the precise way that $\widetilde{D}$ sits in I^n.

With respect to our Weinstein parametrization, $G_t|_D$ and $R_t|_D$ correspond to simple 2-jet rotations g_t and r_t, which we recall are maps $D \to J^2(D, \mathbb{R})$. Consider the difference $g_t - r_t$. Observe from the conclusion of Theorem 4.5

that near ∂D this difference is C^0-small. Let s_t be the composition $(g_t - r_t) \circ F_t^{-1}$ on $F_t(D) \subset \widetilde{D}$ but cut off so that $s_t = 0$ near $\partial F_t(D)$. We precompose by F_t^{-1} to account for the reparametrization F_t used to construct $\widetilde{f}_t$. Notice that even after the cutoff we have that s_t is C^0-close to $(g_t - r_t) \circ F_t^{-1}$ on $F_t(D)$, with the degree of C^0-closeness determined by how much C^0-closeness we demanded in the invokation of Theorem 4.5 which produced $\widetilde{f}_t$. Equivalently, $s_t\big(F_t(q)\big)$ is C^0-close to $g_t(q) - r_t(q)$.

We are now ready to wrinkle. We will produce a wrinkling of the zero section to approximate s_t and then add this wrinkling to the graph of dh_t to approximate $s_t + r_t \sim g_t$. Apply the local wrinkling Lemma 5.8 to the simple rotation of $I^n \subset T^*I^n$ determined by the 2-jet rotation $s_t : I^n \to J^2(\mathbb{R}^n, \mathbb{R})$. The result is an exact homotopy of wrinkled Lagrangian embeddings $\widehat{f}_t : I^n \to T^*I^n$ as in the statement of the Lemma. Note in particular that $\widehat{f}_t$ is the inclusion of the zero section in the complement of $F_t(D) \subset I^n$. Now consider the following addition of the two Lagrangian embeddings $f_t = \widehat{f}_t + \widetilde{f}_t$. We define an exact homotopy of wrinkled Lagrangian embeddings $D \to T^*\widetilde{D}$ by means of the formula

$$q \mapsto \widehat{f}\big(F_t(q)\big) + dh_t\big(\pi \circ \widehat{f}_t\big(F_t(q)\big)\big), \qquad q \in D,$$

where $\pi : T^*\widetilde{D} \to \widetilde{D}$ is the cotangent bundle projection and the addition sign corresponds to the fibrewise addition of cotangent vectors based at the same point.

Note that in the invokation of our local wrinkling lemma we can demand that $\pi \circ \widehat{f}$ be arbitrarily C^0-close to the identity, so $\pi \circ \widehat{f}_t\big(F_t(q)\big)$ is arbitrarily C^0-close to $F_t(q)$ and hence the Lagrangian plane tangent to the graph of dh_t over the point $\pi \circ \widehat{f}\big(F_t(q)\big)$ is C^0-close to $r_t(q)$. Note also that each Lagrangian plane tangent to $\widehat{f}_t$ over that same point is C^0-close to $s_t\big(F_t(q)\big)$. Hence the Lagrangian plane tangent to f_t at the point $f_t(q)$ is C^0-close to the Lagrangian plane corresponding to $r_t(q) + s_t\big(F_t(q)\big)$, which is itself C^0-close to $G_t(q)$, as required.

In remains to explain how to adapt the proof when the n-simplex D has boundary intersecting the wrinkling locus of f. The issue is that any larger disk $\widetilde{D} \supset D$ would necessarily intersect the wrinkling locus. However, this is straightforward to fix: take a smooth Lagrangian disk $\Lambda \subset M$ which contains $f(D)$ in its interior and such that $\widetilde{f}_t(D)$ is graphical over Λ with respect to a Weinstein parametrization of a tubular neighborhood of Λ. A quick look at the proof of Theorem 4.11 is sufficient to convince oneself of the existence of this disk. Then with respect to this Weinstein parametrization we can still write $\widetilde{f}_t$ as a composition $dh_t \circ F_t$ for $F_t : f(D) \to \Lambda$ an isotopy and $h_t : F_t(D) \to \mathbb{R}$ a homotopy of real valued functions. The rest of the proof now proceeds as before. $\square$

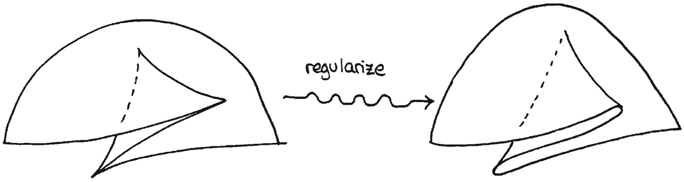

Fig. 37 One-half of a regularized wrinkle. In this picture, the ambient foliation should be thought of as being vertical

The proof of the parametric Theorem 5.4 follows the same outline, using the parametric Theorem 4.10 instead of Theorem 4.5 and using the parametric Lemmas 5.12 and 5.13 instead of Lemmas 5.8 and 5.11. The only essential difference is that in order to localize the parameter space from an arbitrary m-dimensional manifold Z to the unit cube $Z = I^m$ we need to choose a triangulation Δ of $Z \times L$ with sufficiently small simplices, which is compatible with the wrinkles of f and which is in general position with respect to the fibres $z \times L \subset Z \times L$, $z \in Z$. The existence of such a triangulation was proved by Thurston in [46]. Once we know that such a triangulation exists, we can take the fibered polyhedron $K = \Delta^{n+m-1} \subset Z \times L$ and work simplex by simplex.

6 The simplification of singularities

6.1 Wrinkles, swallowtails and double folds

We now return to the setting described in Sect. 1. Let M be a symplectic or contact manifold and let $\mathcal{F}$ be a foliation of M by Lagrangian or Legendrian leaves. Suppose that $f : L \to M$ is a wrinkled Lagrangian or Legendrian embedding which is transverse to $\mathcal{F}$. We can apply the regularization procedure described in Sect. 2.5 to f and obtain a regular Lagrangian or Legendrian embedding $\widetilde{f} : L \to M$. We already observed in Remark 2.10 that $\widetilde{f}$ only has Σ^1-type singularities with respect to $\mathcal{F}$, see Fig. 37 for an illustration. More precisely, $\Sigma(\widetilde{f}, \mathcal{F})$ consists of a disjoint union of regularized wrinkles, which are defined as follows.

Definition 6.1 A regularized wrinkle of a regular Lagrangian or Legendrian embedding $g : L \to M$ with respect to a foliation $\mathcal{F}$ is a connected component of the singularity locus $\Sigma(g, \mathcal{F})$ which consists of a topologically trivial codimension 1 sphere $S \subset L$ such that we can decompose $S = D_1 \cup E \cup D_2$ into two hemispheres D_1 and D_2 and an equator E satisfying the following.

- the equator E consists of Σ^{110} pleats.
- the disks D_1 and D_2 consist of Σ^{10} folds.

For a concrete local model, one can take the standard Lagrangian or Legendrian wrinkle defined in Sect. 2.2, after regularizing as described in Sect. 2.5. In the Lagrangian case, the foliation $\mathcal{F}$ of the cotangent bundle is given by the fibres of the standard projection $\pi : T^*\mathbb{R}^n \to \mathbb{R}^n$. In the Legendrian case, the foliation $\mathcal{F}$ of the 1-jet space $J^1(\mathbb{R}^n, \mathbb{R}) = T^*\mathbb{R}^n \times \mathbb{R}$ is given by the fibres of the front projection $\pi \times id : T^*\mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n \times \mathbb{R}$.

Remark 6.2 If the foliation $\mathcal{F}$ is induced by a Lagrangian fibration $\pi : M^{2n} \to B^n$, then for any regular Lagrangian embedding $f : L^n \to M^{2n}$ the following two conditions are equivalent.

- the singularities of tangency of g with respect to $\mathcal{F}$ consist of a union of regularized wrinkles.
- the front $\pi \circ g : L^n \to B^n$ is a generalized wrinkled mapping in the sense of [15].

In the contact case where $\pi : M^{2n+1} \to B^{n+1}$ is a Legendrian fibration (which we think of as the front projection), we can think of regularized wrinkles in the following way. The singularities of tangency of a regular Legendrian embedding consist of a union $W = \bigcup_j S_j$ of regularized wrinkles if and only if the front of the embedding has cusps on each sphere S_j together with swallowtails on the equator E_j of each S_j.

Regularized wrinkles are also close relatives of the double folds introduced in Sect. 1.3. We recall the definition for convenience.

Definition 6.3 A double fold is a pair of topologically trivial $(n-1)$-spheres S_1 and S_2 in the fold locus $\Sigma^{10} \subset L$ which have opposite Maslov co-orientations and such that $S_1 \cup S_2$ is the boundary of an embedded annulus $A \subset L$.

Indeed, the Entov surgery of [20] can be used to open up a regularized wrinkle along its equator, producing a double fold. This is achieved by taking one of the two hemispheres of a regularized wrinkle $S \subset L$ and pushing it slightly away from S while keeping it fixed on the equator E. We obtain an embedded disk $D \subset L$ contained in an arbitrarily small neighborhood of S in L such that $\partial D = E$ and $\mathrm{int}(D) \cap S = \varnothing$. In fact, we require that $\mathrm{int}(D)$ is outside of the n-ball $B \subset L$ bounded by S. The surgery construction removes the Σ^{110} pleats from E and trades them for Σ^{10} folds on two parallel copies of D. One of these two parallel copies of D is surgered onto one of the hemispheres of S and the other parallel copy is surgered onto the other hemisphere, so that the end result consists of a disjoint union of two parallel spheres on which the embedding has Σ^{10} folds. The Maslov co-orientations on the two resulting spheres are opposite of each other. Hence we end up with the desired double fold. See [20] for the details of the surgery construction and see Fig. 38 for an illustration.

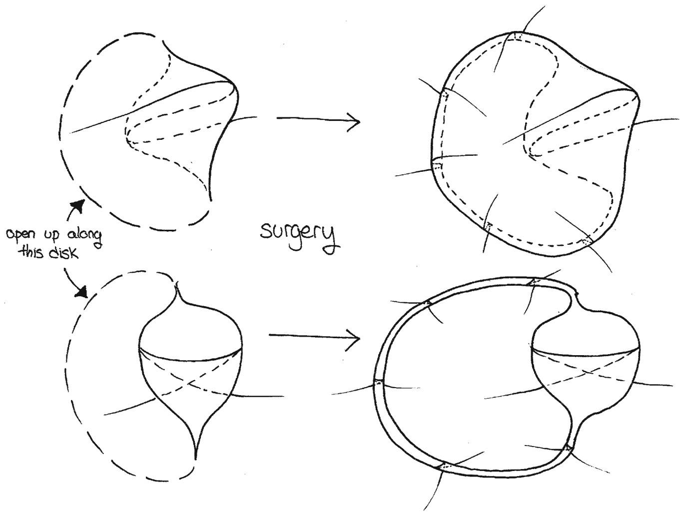

Fig. 38 Opening up a wrinkle into a double fold. The upper picture corresponds to the Lagrangian projection and the lower picture to the front projection

The precise statement that we will need is the following. Given a regular Lagrangian or Legendrian embedding $g : L \to M$ and given $S \subset \Sigma(g, \mathcal{F})$ a regularized wrinkle, there exists a C^0-small ambient Hamiltonian isotopy $\varphi_t : M \to M$ such that $\varphi_t = id_M$ outside of an arbitrarily small neighborhood of $g(S)$ in M and such that inside this neighborhood the regularized wrinkle of g is replaced by a double fold of $\varphi_1 \circ g$. If $g = \widetilde{f}$ is the regularization of a wrinkled Lagrangian or Legendrian embedding $f : L \to M$, then the wrinkles S of f will typically be nested. By this we mean that the ball $B \subset L$ bounded by any wrinkle S of f may contain other wrinkles of f. Hence when we apply the surgery construction on each regularized wrinkle of $g = \widetilde{f}$, we obtain a regular Lagrangian or Legendrian embedding $\varphi_1 \circ g$ whose singularity locus consists of a disjoint union of double folds which are nested in the sense of Sect. 1.3.

Remark 6.4 We could of course have worked with double folds all along without ever mentioning wrinkles. Instead of defining wrinkled Lagrangian and Legendrian embeddings as we did, we could have defined ‘doubly cusped’ Lagrangian and Legendrian embeddings to be topological embeddings which are smooth Lagrangian or Legendrian embeddings away from a finite union of pairs of parallel spheres, where the embedding has cusps of opposite Maslov co-orientation (the cusps are semi-quintic in the ambient symplectic or contact manifold and semi-cubic in the front projection). Our C^0-approximation result for a tangential rotation G_t would also hold for the class of doubly cusped Lagrangian and Legendrian embeddings. Moreover, the regularization of a doubly cusped Lagrangian or Legendrian embedding which is transverse to a foliation $\mathcal{F}$ is a regular Lagrangian or Legendrian embedding whose singularities of tangency with respect to $\mathcal{F}$ consist of double folds. The h-principle for the simplification of singularities proved below then follows with the same proof. We have chosen to work with wrinkles instead to draw the parallel with the smooth wrinkled embeddings theorem [15].

Suppose next that $f^z : L \to M$ is a family of wrinkled Lagrangian or Legendrian embeddings parametrized by a compact manifold Z. We can also in this case regularize and obtain a family of regular Lagrangian or Legendrian embeddings $\widetilde{f}^z : L \to M$. If f^z is transverse to $\mathcal{F}$, then the singularities of tangency of the family $\widetilde{f}^z$ with respect to $\mathcal{F}$ consist of fibered regularized wrinkles. In particular, for some values of the parameter $z \in Z$ the regular Lagrangian or Legendrian embedding $\widetilde{f}^z$ will have regularized embryos in addition to regularized wrinkles. Regularized embryos are non-generic Σ^1-type singularities of tangency which occur at the instance of birth/death of a regularized wrinkle. One can of course give a concrete local model for the regularized embryo, however it is simpler to think about families as a single object using the fibered terminology. For a concrete local model, one can take the standard fibered Lagrangian or Legendrian wrinkle defined in Sect. 2.3, after regularizing as described in Sect. 2.5. The foliation $\mathcal{F}$ is given as in the non-parametric case.

Remark 6.5 If the foliation $\mathcal{F}$ is induced by a Lagrangian fibration $\pi : M^{2n} \to B^n$, then for any family of regular Lagrangian embeddings $g : Z^m \times L^n \to M^{2n}$ the following two conditions are equivalent.

- the singularities of tangency of g with respect to $\mathcal{F}$ consist of a union of fibered regularized wrinkles.
- the fibered front $p \circ g : Z^m \times L^n \to Z^m \times B^n$ is a fibered generalized wrinkled mapping in the sense of [15].

In the Legendrian case one can of course reinterpret what fibered regularized wrinkles mean in the front projection in terms of cusps and swallowtails. Note that one can also use the Entov surgery in families to replace fibered regularized wrinkles with fibered double fold singularities. The embryos of regularized wrinkles will become embryos of double folds. An embryo of a double fold is a non-generic locus of Σ^1-type singularities of tangency consisting of a single codimension 1 sphere from which the two parallel spheres of folds can either be born or die, see Fig. 39.

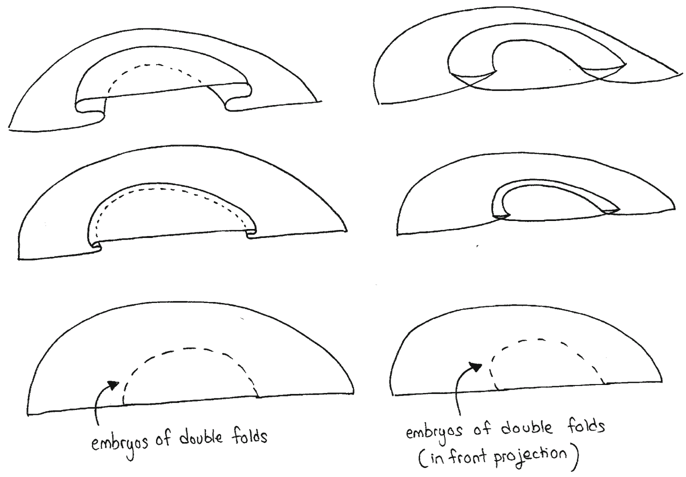

Fig. 39 One-half of the birth/death of a double fold. The picture on the left corresponds to the Lagrangian projection and the picture on the right corresponds to the front projection

6.2 The *h*-principle for the simplification of singularities

We are now ready to establish the flexibility of singularities of Lagrangian and Legendrian fronts. As above, $\mathcal{F}$ denotes a foliation by Lagrangian or Legendrian leaves of a symplectic or contact manifold M.

Theorem 6.6 *Suppose that there exists a tangential rotation* $G_t : L \to \Lambda(M)$ *of a regular Lagrangian or Legendrian embedding* $f : L \to M$ *such that* $G_1 \pitchfork \mathcal{F}$*. Then there exists a compactly supported ambient Hamiltonian isotopy* $\varphi_t : M \to M$ *such that the singularities of* $\varphi_1 \circ f$ *consist of a union of nested regularized wrinkles.*

Proof Apply the wrinkling Theorem 5.1 to G_t and f. We obtain a compactly supported exact homotopy of wrinkled Lagrangian or Legendrian embeddings $f_t : L \to M$ such that $G(df_1) \pitchfork \mathcal{F}$. Next, apply the regularization process described in Sect. 2.5 to the homotopy f_t. We obtain a compactly supported exact homotopy of regular Lagrangian or Legendrian embeddings $\widetilde{f_t} : L \to M$ such that the singularity locus $\Sigma(\widetilde{f_1}, \mathcal{F}) \subset L$ consists of a disjoint union of regularized wrinkles. Finally, since the homotopy $\widetilde{f_t}$ is exact and compactly supported, we can write $\widetilde{f_t} = \varphi_t \circ f$ for some compactly supported ambient Hamiltonian isotopy $\varphi_t : M \to M$. □

To deduce the version with double folds stated in Theorem 1.11, we simply apply the Entov surgery construction of [20] to open up each of the wrinkles as described in the previous section.

Remark 6.7 At each stage of the proof, when we apply Theorem 5.1, the regularization of Sect. 2.5 and the Entov surgery, we can always ensure that the resulting homotopy of embeddings is C^0-close to f. Hence Theorem 6.6 also holds in C^0-close form, where we demand that the Hamiltonian isotopy φ_t is C^0-close to the identity id_M. Moreover, we can also ensure that $\varphi_t = id_M$ outside of a neighborhood of $f(L)$ in M.

Remark 6.8 Suppose that $G_t = G(df)$ on $Op(A)$ for some closed subset $A \subset L$. At each stage of the proof, when we apply Theorem 5.1, the regularization of Sect. 2.5 and the Entov surgery, we can always ensure that the resulting homotopy of embeddings agrees with f on $Op(A)$. Hence Theorem 6.6 also holds in relative form. More precisely, we can demand that $\varphi_t = id_M$ on $Op\big(f(A)\big) \subset M$.

The parametric version reads as follows, and is proved in exactly the same way. At each stage we just need to invoke the parametric versions of each of the ingredients of the proof. The corresponding C^0-close and relative versions also hold, for the same reasons as in the non-parametric case.

Theorem 6.9 *Suppose that there exists a family of tangential rotations $G_t^z : L \to \Lambda(M)$ of regular Lagrangian or Legendrian embeddings $f^z : L \to M$ parametrized by a compact manifold Z such that $G_1^z \pitchfork \mathcal{F}$ for all $z \in Z$ and such that $G_t^z = G(df^z)$ for $z \in Op(\partial Z)$. Then there exists a family of compactly supported ambient Hamiltonian isotopies $\varphi_t^z : M \to M$ such that the singularities of $\varphi_1^z \circ f^z$ consist of a union of fibered nested regularized wrinkles and such that $\varphi_t^z = id_M$ for $z \in Op(\partial Z)$.*

As in the non-parametric case we can open up the fibered regularized wrinkles into fibered double folds using the Entov surgery construction [20].

Remark 6.10 Observe that in the case $n = 1$ there is no need to resolve a wrinkle into a double fold. Indeed a 1-dimensional regularized wrinkle consists of nothing more than a pair of points where the embedding has folds of opposite Maslov co-orientation. For fibered regularized wrinkles the two folds die as in the Legendrian Reidemeister I move. We explore the case $n = 1$ further in Sect. 6.5 below.

6.3 The *h*-principle for the prescription of singularities

We next prove a strengthened version of Entov's Theorem 1.16. More precisely, we apply our *h*-principle Theorem 6.6 to drop the Σ^2-nonsingularity restriction from his result. As an application we establish some concrete results for the simplification of the caustics of spheres in Sect. 6.4 below.

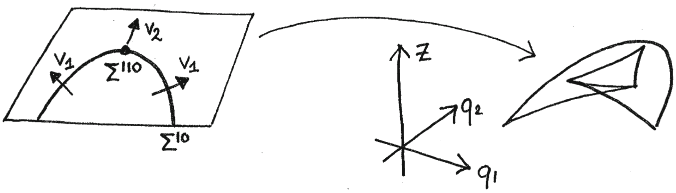

Fig. 40 The chain of singularities associated to the Σ^{110} pleat, which is a swallowtail in the front projection. A flip of the Legendrian front in the z direction would reverse the Maslov co-orientation v_1 and fix v_2

Consider $f : L \to M$ a Lagrangian or Legendrian embedding and let D be a Lagrangian distribution in TM defined along $f(L)$. In the symplectic case, D consists of linear Lagrangian subspaces of (TM, ω) and in the contact case D consists of linear Lagrangian subspaces of $(\xi, d\alpha)$, where locally $\xi = \ker(\alpha) \subset TM$.

When $\dim(df(T_qL) \cap D_{f(q)}) < 2$ for all $q \in L$ we say that D is Σ^2-nonsingular. In this case, the structure of the singularity locus $\Sigma = \{q \in L : df(T_qL) \cap D_{f(q)} \neq 0\}$ is quite simple. Indeed, for generic Σ^2−nonsingular D the locus Σ is a codimension 1 submanifold which is naturally stratified as a flag $\Sigma = \Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$ as described in Sect. 1.3. Moreover, the flag comes equipped with certain co-orientation data which we hinted about in Sect. 1.9 and which more precisely consist of the following.

- Unit vector fields v_k, $k > 1$, where each v_k is defined on $\Sigma^{1^k} \backslash \Sigma^{1^{k+1}}$, is normal to $\Sigma^{1^{k-1}}$ in $\Sigma^{1^{k-2}}$ and cannot be extended (as such a unit normal vector field) to any subset $C \subset \Sigma^{1^k}$ which has a nontrivial intersection with $\Sigma^{1^{k+1}}$.
- An additional unit vector field v_1 defined on the whole of Σ which is normal to Σ in L. This vector field is called the Maslov co-orientation.

Adapting Eliashberg's terminology from [9], Entov defined in [20] the chain of singularities associated to f and D to consist of the flag $\Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$ together with vector fields v_k as above. The v_k are uniquely determined by the geometry of the singularity. See Fig. 40 for an illustration. Two chains of singularities are said to be equivalent if there exists an isotopy of L that transforms one into the other, including the co-orientation data. We can now state and prove an h-principle which allows for the prescription of any homotopically allowable chain of singularities. The result also holds in C^0-close and relative forms.

Theorem 6.11 *Let $f : L \to M$ be a regular Lagrangian or Legendrian embedding into a symplectic or contact manifold M equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. Let D_t be a homotopy of Lagrangian distributions defined along $f(L)$, fixed outside of a compact subset, such* *that* $D_0 = T\mathcal{F}|_{f(L)}$ *and such that* f *is* Σ^2*-nonsingular with respect to the distribution* D_1*. We moreover assume that* $f \pitchfork \mathcal{F}$ *outside of that compact subset. Then there exists a* C^0*-small compactly supported Hamiltonian isotopy* $\varphi_t : M \to M$ *such that* $\varphi_1 \circ f$ *is* Σ^2*-nonsingular with respect to* $\mathcal{F}$ *and moreover such that the chain of singularities of* $\varphi_1 \circ f$ *with respect to* $\mathcal{F}$ *is equivalent to the chain of singularities of* f *with respect to* D_1*, together with a union of nested double folds.*

Proof We restrict our attention to the Lagrangian case for concreteness, the Legendrian analogue is no different. Let $\Sigma \subset L$ be the singularity locus of f with respect to D_1. By abusing notation, we will also denote by Σ the chain of singularities which encodes the flag $\Sigma = \Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$ and the corresponding co-orientation data. Let Φ_t be a homotopy of linear symplectic isomorphisms of TM defined along $f(L)$ such that $\Phi_0 = id$ and $\Phi_t \cdot D_0 = D_t$. Set $G_t = (\Phi_t)^{-1} \cdot G(df)$, a tangential rotation of f.

Our plan will be the following. We will first apply our holonomic approximation lemma for 1−holonomic sections to G_t to make f transverse to $\mathcal{F}$ near a parallel copy $\Sigma_{1/2}$ of Σ. Then we will introduce by hand a cancelling pair of singularity loci Σ_1 and Σ_2 in $Op(\Sigma_{1/2})$ such that Σ_2 is equivalent to Σ and such that $\Sigma \cup \Sigma_1$ bounds an embedded annulus which is disjoint from Σ_2. Formally, Σ and Σ_1 can be cancelled via a rotation R_t which is fixed on Σ_2 and hence by our relative h−principle for the simplification of singularities we are able to keep the singularity locus Σ_2 and fill in the rest of the Lagrangian submanifold with double folds. See Fig. 41 for an illustration of the strategy.

Let $l = (df)^{-1}(D_1)$, which is a line field on TL defined along Σ. Extend l to a tubular neighborhood $\mathcal{N} \simeq \Sigma \times (-1, 1)$ of Σ in L so that with respect to this parametrization l is constant in the $(-1, 1)$ direction. Denote by $\Sigma_{1/2}$ the parallel copy $\Sigma \times \frac{1}{2}$ of Σ in $\mathcal{N}$. Apply Theorem 4.2 to the tangential rotation G_t and the stratified subset $K = \Sigma_{1/2}$. We obtain an exact homotopy of regular Lagrangian embeddings $f_t : L \to M$ such that $G(df_t)$ is C^0-close to G_t on $Op(\Sigma_{1/2})$. In particular, $f_1 \pitchfork \mathcal{F}$ on a neighborhood $U = \Sigma \times (1/2 - \varepsilon,\ 1/2 + \varepsilon)$ of $\Sigma_{1/2}$.

Along $f_1(U)$, the Lagrangian distributions $df_1(TU)$ and $T\mathcal{F}|_{f_1(U)}$ are transverse. We can therefore choose a symplectic isomorphism $T(T^*U)|_U \simeq TM|_{f_1(U)}$ such that the horizontal distribution TU (which is tangent to the zero section) is mapped to $df_1(TU)$ using df_1 and such that the vertical distribution VU (which is tangent to the cotangent fibres) is mapped to $T\mathcal{F}|_{f_1(U)}$. Choose an $(n-1)$-dimensional complement P for l in TU. Set $l^* = P^\perp \cap VU$ and $P^* = l^\perp \cap VU$, where $\perp$ denotes orthogonality with respect to the symplectic form $dp \wedge dq$. Let $\phi : [1/2 - \varepsilon, 1/2 + \varepsilon] \to \mathbb{R}$ be a function satisfying the following properties.

- $\phi(s) = 0$ for s near $1/2 \pm \varepsilon$.

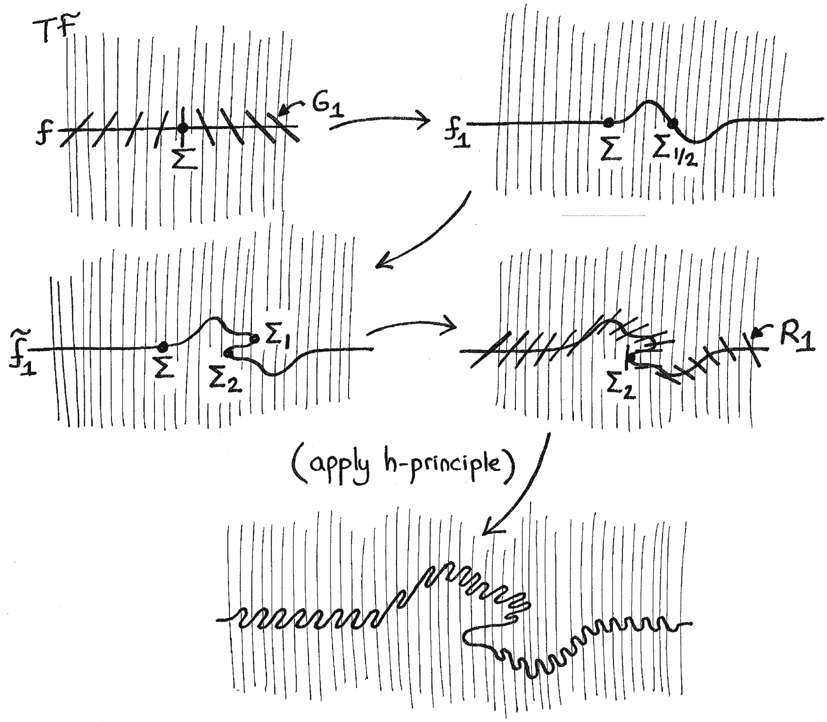

Fig. 41 The plan for our proof of Theorem 6.11

- $\phi(s) = \pi$ for s near $1/2$.
- $\phi'(s) \geq 0$ for $s \in (1/2 - \varepsilon, 1/2]$ and $\phi'(s) \leq 0$ for $s \in [1/2, 1/2 + \varepsilon)$.

Fix nonzero vector fields $v \in l$ and $w \in l^*$. Define a homotopy of Lagrangian distributions $V_t \subset T(T^*L)$ defined along $U = \Sigma \times (1/2 - \varepsilon,\ 1/2 + \varepsilon)$ by the formula

$$V_t(e, s) = \text{span}\big(\sin\big(t\phi(s)\big)v + \cos\big(t\phi(s)\big)w\big) \oplus P^*,$$
$$(e, s) \in \Sigma \times (1/2 - \varepsilon,\ 1/2 + \varepsilon).$$

Note that $V_0 = VU$, that $V_t = VU$ on ∂U and that $\dim(V_t \cap TU) \leq 1$ for all $t \in [0, 1]$. The singularities of tangency of V_1 with respect to the zero section $U \hookrightarrow T^*U$ consist of two parallel copies Σ' and Σ'' of Σ, for concreteness say Σ' is between Σ and Σ''. Along these singularitiy loci we have $V_1 \cap TU = l$. The two corresponding chains of singularities, which we also denote by Σ' and Σ'', have opposite Maslov co-orientations but are otherwise equivalent. Replacing the function ϕ by the function $-\phi$ if necessary, we may assume that the chain of singularities Σ'' is equivalent to the chain Σ.

At this point we wish to use V_t to insert by hand a cancelling pair of singularities modelled on Σ. The explicit formulas that we need are written down in Entov's paper [20]. We could use these formulas to write down a concrete model for the creation of the cancelling pair, but we can make our life even easier by directly applying Entov's Theorem 1.16 to V_t. The output of Entov's theorem is an exact homotopy of regular Lagrangian embeddings $g_t : U \to T^*U$ such that g_0 is the inclusion of the zero section $U \hookrightarrow T^*U$, such that g_t is fixed on $Op(\partial U)$ and such that the singularities of tangency of g_1 with respect to VU are equivalent to those of g_0 with respect to V_1, together with a union of nested double folds. Furthermore, the homotopy g_t can be assumed to be C^0-small, so by taking an appropriate Weinstein neighborhood we can think of this homotopy as happening inside M. The result is an exact regular homotopy $\widetilde{f}_t : L \to M$ of $\widetilde{f}_0 = f_1$ such that along $U \subset L$ the singularities of tangency of $\widetilde{f}_1$ with respect to $\mathcal{F}$ consist of a union $\Sigma_1 \cup \Sigma_2 \cup F$, where the chain Σ_1 is equivalent to Σ', the chain Σ_2 is equivalent to Σ'' and F is a union of nested double folds. Moreoever, $\Sigma \cup \Sigma_1$ bounds an annulus $A \subset L$ which is disjoint from Σ_2.

Claim 6.12 *There exists a tangential rotation $R_t : L \to \Lambda(M)$ of $\widetilde{f}_1$ which is fixed on $Op(\Sigma_2)$ and such that $R_1 \pitchfork \mathcal{F}$ away from Σ_2.*

Once this claim is established we are done, since we can apply the relative version of Theorem 6.6 to construct an exact homotopy of regular Lagrangian embeddings which is fixed on $Op(\Sigma_2)$ and such that at the end of the homotopy the singularities of tangency away from Σ_2 consist of a union of nested double folds, which is exactly what we wanted to prove.

To justify the claim, we first observe that there exists a tangential rotation $S_t : L \to \Lambda(M)$ of $\widetilde{f}_1$ such that S_t is fixed on $Op(\Sigma_1 \cup \Sigma_2)$, such that $S_1 = G_1$ outside of U and such that $S_1 \pitchfork \mathcal{F}$ away from $\Sigma_1 \cup \Sigma_2 \cup \Sigma$. To define S_t, choose $\delta_1 < \delta_2 < \varepsilon$ such that the annuli $U_i = \Sigma \times (1/2 - \delta_i, 1/2 + \delta_i) \subset U$ contain $\Sigma_1 \cup \Sigma_2 \cup F$. Inside of U_1, we let S_t kill the double folds of F so that the only remaining singularities are $\Sigma_1 \cup \Sigma_2$. On the rest of L (where we may assume that $\widetilde{f}_1 = f_1$ provided that δ_1 and δ_2 are close enough to ε), we construct S_t in three steps.

- First, rotate $G(d\widetilde{f}_1) = G(df_1)$ to a distribution W which equals $G(df_0)$ away from U and which interpolates between $G(df_0)$ and $G(df_1)$ on $U \setminus U_2$ by means of $G(df_t)$.
- Since $G(df_t)$ is C^0-close to G_t on U, we can then rotate W to a distribution W' which equals $G_0 = G(df_0)$ away from U, which interpolates between G_0 and G_1 on $U \setminus U_2$ by means of G_t and which then interpolates between G_1 and $G(df_1)$ on $U_2 \setminus U_1$.
- We can then rotate W' to a distribution W'' which equals G_1 outside of U_2 and which interpolates between G_1 and $G(df_1)$ on $U_2 \setminus U_1$. The distri-

bution $W'' = S_1$ satisfies the required properties and the rotation S_t is the concatenation of the three steps.

Consider now the annulus $A \subset L$ with boundary $\partial A = \Sigma \cup \Sigma_1$. The intersection $\lambda = \mathrm{im}(S_1) \cap T\mathcal{F} \subset TM$ consists of two line fields defined over the images of Σ and Σ_1. We claim that they extend to a line field $\lambda \subset \mathrm{im}(S_1)$ defined over the image of the whole annulus A.

Indeed, the chain of singularities of Σ_1 is equivalent to that of Σ up to Maslov co-orientation. But the isotopy class of the line field which arises from a Σ^1-type singularity locus is completely dictated by the flag $\Sigma^1 \supset \Sigma^{11} \supset \cdots \supset \Sigma^{1^n}$ together with the non-Maslov co-orientation data. Hence the line fields are isotopic in TL. It follows that we can find a line field $\widetilde{l} \subset TL$ defined along A such that $\widetilde{l}|_{\Sigma_1} = d\widetilde{f}_1^{-1}(\lambda)$ and such that $\widetilde{l}|_{\Sigma} = df^{-1} \circ \Phi_1(\lambda)$.

Suppose that there exists a family of symplectic isomorphisms Ψ_t of TM such that $\Psi_0 = id$, such that $\Psi_t \cdot G(df) = S_t$, such that $\Psi_1 \circ df = d\widetilde{f}_1$ near Σ_1 and such that $\Psi_1 = \Phi_1^{-1}$ near Σ. Then the line field $\lambda = \Psi_1 \circ df(\widetilde{l})$ is the required extension. It remains to confirm that the family Ψ_t exists. We need to define Ψ_t over $A \times [0, 1]$, where $t \in [0, 1]$ and we have prescribed Ψ_t over $A \times 0 \cup (\partial A \times [0, 1])$. Furthermore, we also have prescribed the image of Ψ_t under the map $\Psi_t \mapsto \Psi_t \cdot G(df)$ over all of $A \times [0, 1]$. Since this map is a Serre fibration, it follows that we can find a lift to all of $A \times [0, 1]$. This completes the proof of the existence of the line field $\lambda \subset \mathrm{im}(S^1)$.

Next we observe that the distribution $S_1 : A \to \Lambda(M)$ satisfies $S_1(\partial A) \subset \Sigma^1(M, \mathcal{F}) = \bigcup_{x \in M} \{W \in \Lambda(M)_x : \dim(W \cap T_x\mathcal{F}) = 1\}$ and $S_1\big(\mathrm{int}(A)\big) \subset \Lambda^{\pitchfork}(M, \mathcal{F}) = \bigcup_{x \in M} \{W \in \Lambda(M)_x : W \cap T_x\mathcal{F} = 0\}$. Pick a complement $Q \subset \mathrm{im}(S^1)$ to λ. Set $\lambda^* = Q^{\perp} \cap T\mathcal{F}$ and $Q^* = \lambda^{\perp} \cap T\mathcal{F}$. Pick nonzero vector fields $v \in \lambda$ and $w \in \lambda^*$ such that $\omega(v, w) > 0$ on $\mathrm{int}(A)$ and define a rotation $R_t : A \to \Lambda(M)$ starting at $R_0 = S_1$ by the formula

$$R_t = \mathrm{span}\big(\cos(\pi t/2)v + \sin(\pi t/2)w\big) \oplus Q, \qquad t \in [0, 1].$$

Observe that on ∂A we have $\lambda^* = \lambda$ and hence $R_t = S_1$ for all $t \in [0, 1]$. Hence we can extend R_t outside of A by letting it equal S_1 elsewhere. Observe also that $R_1 \cap T\mathcal{F} = \lambda^*$ along A and hence $\mathrm{im}(R_1)|_A \subset \Sigma^1(M, \mathcal{F})$. Recall that $\Sigma^1(M, \mathcal{F})$ is a two-sided hypersurface of $\Lambda(M)$, so that if $\mathcal{O} \subset \Lambda(M, \mathcal{F})$ is a small enough neighborhood of $\mathrm{im}(R_1)|_A$, then $\mathcal{O} \backslash \Sigma^1(M, \mathcal{F})$ has exactly two connected components. The fact that the Maslov co-orientations of Σ_1 and Σ are opposite means precisely that $\mathrm{im}(S_1)|_{Op(A) \backslash A}$ lies in the same connected component of $\mathcal{O} \backslash \Sigma^1(M, \mathcal{F})$. Hence we can push the image of R_1 entirely off of $\Sigma^1(M, \mathcal{F})$ by a small deformation which is fixed outside of $Op(A)$. The result is a rotation $\widetilde{R}_t : L \to \Lambda(M)$ starting at $\widetilde{R}_0 = S_1$ such that $\widetilde{R}_1 = S_1$ on $Op(\Sigma_2)$ and such that $\widetilde{R}_1 \pitchfork \mathcal{F}$ away from Σ_2. This completes the proof of Claim 6.12, hence also of Theorem 6.11. □

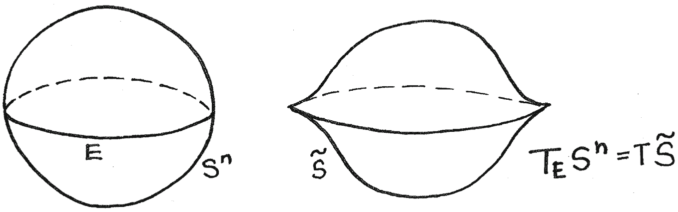

Fig. 42 Take Σ to be the equatorial sphere $E \subset S^n$. The vector bundle $T_E S^n$ can be visualized as the tangent bundle of a singular surface $\widetilde{S}$ as illustrated above

6.4 Application: the caustics of spheres

We now return to the first example considered in Sect. 1.1. Our goal is to study the extent to which it is possible to simply the caustic of an embedded Lagrangian or Legendrian sphere $S \subset M$, where M is a symplectic or contact manifold equipped with a foliation $\mathcal{F}$ by Lagrangian or Legendrian leaves. For greater clarity of the exposition we will restrict our discussion to the Lagrangian version of the problem, but the Legendrian analogue is no different.

First we observe that by the Weinstein neighborhood theorem we can immediately reduce to the case where $M = T^*S^n$ and S is the image of the zero section $S^n \hookrightarrow T^*S^n$, which we will also denote by S^n. Note that for $n = 1$ the problem is uninteresting because the generic caustic consists only of folds, so the simplification of singularities can be trivially achieved. We assume $n > 1$ in what follows. Let V be the restriction to S^n of the distribution $T\mathcal{F}$ of Lagrangian planes tangent to $\mathcal{F}$. We begin with the following topological obstruction to the simplification of singularities.

Proposition 6.13 *If S^n is Hamiltonian isotopic to a Lagrangian sphere whose singularities of tangency with respect to $\mathcal{F}$ consist only of folds, then V is stably trivial as a real vector bundle over the sphere.*

We precede the proof with some notation. Let $\Sigma \subset S^n$ be any compact hypersurface. Following [17], it is conceptually useful to introduce a real n-dimensional vector bundle $T_\Sigma S^n$ which is obtained from TS^n by regluing along Σ with a fold. More precisely, write $S^n = X \cup Y$ for $X, Y \subset S^n$ two compact n-dimensional submanifolds whose common boundary $\partial X = X \cap Y = \partial Y$ is the hypersurface Σ. Fix also an identification $TS^n|_\Sigma \simeq T\Sigma \oplus \varepsilon$, where ε denotes the trivial line bundle. Define $T_\Sigma S^n$ to be the real n-dimensional vector bundle over S^n given by gluing the disjoint union $TX \coprod TY$ over the intersection $X \cap Y = \Sigma$ via the isomorphism (Fig. 42)

$$\mu = id \oplus (-1) : T\Sigma \oplus \varepsilon \to T\Sigma \oplus \varepsilon.$$

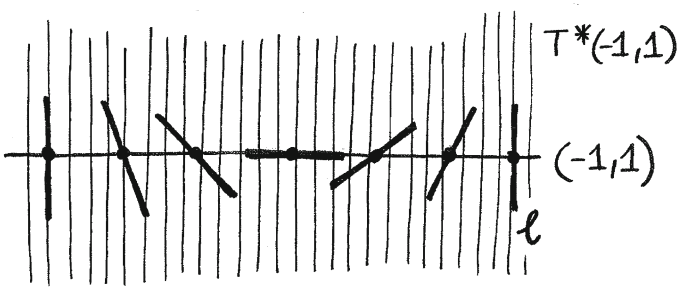

Fig. 43 The rotating line field $l \subset T\big(T^*(-1,1)\big)$

The bundle $T_\Sigma S^n$ can be realized as a distribution of Lagrangian planes V_Σ in T^*S^n defined along the zero section $S^n \hookrightarrow T^*S^n$ whose singularities of tangency with respect to the zero section S^n consist of folds along Σ. In order to do this, we fix a co-orientation of Σ, which will agree with the Maslov co-orientation induced by V_Σ. Let $\Sigma \times (-1,1) \simeq \mathcal{N} \subset S^n$ be a tubular neighborhood of Σ such that the canonical orientation of the interval $(-1,1)$ induces the chosen co-orientation of Σ. The Lagrangian Grassmannian $\Lambda\big(T^*(-1,1)\big)$ is the trivial circle bundle $T^*(-1,1) \times S^1$. We use the canonical coordinates $(q,p) \in (-1,1) \times \mathbb{R} = T^*(-1,1)$. Let $l : (-1,1) \to \Lambda\big(T^*(-1,1)\big)$ be the rotating line field defined over the zero section $(-1,1) \hookrightarrow T^*(-1.1)$ by the formula (Fig. 43)

$$l_q = \operatorname{span}\left(\cos\left(\frac{\pi i q}{2}\right)\frac{\partial}{\partial q} + \sin\left(\frac{\pi i q}{2}\right)\frac{\partial}{\partial p}\right) \subset T_{(q.0)}\big(T^*(-1,1)\big).$$

Define $V_\Sigma : \mathcal{N} \to \Lambda(T^*\mathcal{N})$ to be the distribution of Lagrangian planes defined over the zero section $\mathcal{N} \hookrightarrow T^*\mathcal{N}$ which corresponds to the product of the cotangent fibres of $T^*\Sigma$ and the line field l under the isomorphism $T^*\mathcal{N} \simeq T^*\Sigma \times T^*(-1,1)$. The distribution V_Σ extends to the complement of $\mathcal{N}$ in S^n by letting it consist of the cotagent fibres of T^*S^n on $S^n \setminus \mathcal{N}$. The real vector bundle underlying V_Σ is isomorphic to $T_\Sigma S^n$.

Proof of Proposition 6.13 We first consider the special case where S^n itself has only fold singularities with respect to $\mathcal{F}$. Then the caustic $\Sigma = \Sigma(S^n, \mathcal{F})$ is an embedded hypersurface in S^n co-oriented by the Maslov co-orientation. A direct consequence of the local model for the Σ^{10} fold is that V and V_Σ are homotopic in the space of Lagrangian distributions. Since the real vector bundle underlying V_Σ is isomorphic to $T_\Sigma S^n$, it remains to show that $T_\Sigma S^n$ is stably trivial. To see this, observe that $T_\Sigma S^n \oplus \varepsilon$ is obtained from $TX \oplus \varepsilon$ and $TY \oplus \varepsilon$ by using the gluing $\mu \oplus (1) = id \oplus (-1) \oplus (1)$ along $X \cap Y = \Sigma$, where we still think of $TS^n|_\Sigma$ as $T\Sigma \oplus \varepsilon$. Nothing changes if instead we use the gluing $\eta = id \oplus (1) \oplus (-1)$, since the two linear isomorphims of $\mathbb{R}^2$ given by

$(x, y) \mapsto (-x, y)$ and $(x, y) \mapsto (x, -y)$ are in the same connected component of $GL(2, \mathbb{R})$. We can therefore define a bundle map $T_\Sigma S^n \oplus \varepsilon \to TS^n \oplus \varepsilon$ by sending $TX \oplus \varepsilon \to TS^n \oplus \varepsilon$ via the inclusion $id \oplus (1)$, by sending $TY \oplus \varepsilon \to TS^n \oplus \varepsilon$ via the map $id \oplus (-1)$ and by gluing the two pieces into a global map $T_\Sigma S^n \oplus \varepsilon \to TS^n \oplus \varepsilon$ using η. This glued up map is an isomorphism, hence $T_\Sigma S^n \oplus \varepsilon \simeq TS^n \oplus \varepsilon \simeq \varepsilon^{n+1}$, as claimed.

Consider now the general case where $\varphi_t : T^*S^n \to T^*S^n$ is a Hamiltonian isotopy such that $\varphi_1(S^n)$ only has fold singularities with respect to $\mathcal{F}$. Equivalently, S^n only has fold singularities with respect to the pullback foliation $\varphi_1^*\mathcal{F}$. From the special case already considered it follows that the restriction V' to S^n of the distribution $T(\varphi_1^*\mathcal{F})$ must be stably trivial as a real vector bundle over the sphere. But V and V' are homotopic as distributions of Lagrangian planes and therefore isomorphic as real vector bundles. Hence V is also stably trivial. □

We now use our h-principle for the prescription of singularities to show that for n even, the necessary condition for the simplification of singularities provided by Proposition 6.13 is also sufficient.

Corollary 6.14 *Assume that n is even and that $V = T\mathcal{F}|_S$ is stably trivial as a real vector bundle over the sphere. Then there exists a compactly supported Hamiltonian isotopy $\varphi_t : T^*S^n \to T^*S^n$ such that the singularities of tangency of $\varphi_1(S^n)$ with respect to $\mathcal{F}$ consist only of folds. Moreover, we can take φ_t to be C^0-close to the identity and supported on an arbitrarily small neighborhood of the zero section.*

Remark 6.15 From the proof we can also extract a precise description of the permissible fold loci $\Sigma = \Sigma(\varphi_1(S^n), \mathcal{F})$ as hypersurfaces of S^n in terms of the Euler number $e(V)$ of V. The locus Σ can be arranged to consist of the boundary ∂Y of any n-dimensional compact submanifold $Y \subset S^n$ of Euler characteristic $\chi(Y) = 1 \pm \frac{1}{2}e(V)$, together with a disjoint union of nested double folds.

Proof If $B \subset S^n$ is a closed embedded n−ball, it is readily seen that $T_{\partial B}S^n$ is the trivial bundle. Fix a trivialization $V_{\partial B} \simeq S^n \times \mathbb{R}^n$. We obtain a trivialization $T(T^*S^n)|_{S^n} \simeq S^n \times \mathbb{C}^n$ by identifying both bundles with $V_{\partial B} \otimes \mathbb{C}$. Suppose that B is chosen so that $\mathcal{F}$ is transverse to S^n along $Op(B)$. Then with respect to this trivialization the distribution V determines a class $\alpha \in \pi_n(\Lambda_n)$, where $\Lambda_n = U_n/O_n$ is the Grassmannian of linear Lagrangian subspaces of $\mathbb{C}^n$ and we choose any $b \in \text{int}(B)$ as a basepoint. Let $\beta \in \pi_{n-1}(O_n)$ be the image of α under the map $\pi_n(\Lambda_n) \to \pi_{n-1}(O_n)$ given by long exact sequence in homotopy groups associated to the fibration $O_n \to U_n \to \Lambda_n$. Observe that β is the clutching function corresponding to the real vector bundle underlying the distribution V. Note that the choice of ball B induces a choice of orientation on V, which is encoded in the class β.

The stable triviality of V means that β is in the kernel of the map $\pi_{n-1}(O_n) \to \pi_{n-1}(O)$, where $O = \lim_k O_k$ is the stable orthogonal group. However, $\pi_{n-1}(O_k) \to \pi_{n-1}(O_{k+1})$ is an isomorphism as soon as $k > n$, and therefore $\beta \in \ker\big(\pi_{n-1}(O_n) \to \pi_{n-1}(O_{n+1})\big) = \operatorname{im}\big(\pi_n(S^n) \to \pi_{n-1}(O_n)\big)$, where the map is given by the long exact sequence in homotopy groups associated to the fibration $O_n \to O_{n+1} \to S^n$. Recall that under this map the fundamental class $1 \in \mathbb{Z} \simeq \pi_n(S^n)$ is sent to the clutching function $\gamma \in \pi_{n-1}(O_n)$ corresponding to the tangent bundle TS^n. We can therefore write $\beta = k\gamma$ for some $k \in \mathbb{Z}$.

Let $E \subset S^n$ by any compact hypersurface disjoint from B. Let X and Y be as in the construction of $T_E S^n$, so that $S^n = X \cup Y$ and $X \cap Y = E$. We choose the labels so that $B \subset X$, and then we agree to orient $T_E S^n$ so that the inclusion $TX \hookrightarrow T_E S^n$ is orientation preserving. It is straightforward to compute the Euler class $e(T_E S^n) = 2-2\chi(Y)$ using for example the Poincaré-Hopf index theorem. Since $e(V) = 2k$, if we choose the hypersurface E so that $\chi(Y) = 1-k$, then it follows that $T_E S^n$ and V are isomorphic as oriented real vector bundles.

Using the same construction as above, we can exhibit $T_E S^n$ as a distribution V_E of Lagrangian planes in T^*S^n defined along the zero section $S^n \hookrightarrow T^*S^n$. Observe that the singularities of tangency of the zero section S^n with respect to the distribution V_E consist of Σ^{10} folds along E.

Since n is even, $\pi_n(U_n) = 0$ and hence we have an injection $\pi_n(\Lambda_n) \hookrightarrow \pi_{n-1}(O_n)$. Observe that the homotopy classes in $\pi_n(\Lambda_n)$ determined by the distributions V_E and V have the same image β under this map. It follows that V_E and V are homotopic in the space of Lagrangian distributions. The h-principle for the prescription of singularities Theorem 6.11 applies to produce a C^0-small Hamiltonian isotopy $\varphi_t : T^*S^n \to T^*S^n$ supported in a neighborhood of the zero section such that the singularities of tangency of $\varphi_1(S^n)$ with respect to $\mathcal{F}$ are equivalent to those of S^n with V together with a union of nested double folds, which completes the proof. □

In fact, the assumption that V is stably trivial is automatically satisfied for all even n such that $n \not\equiv 2 \bmod 8$. One can argue in the following way. Choose a class $\beta \in \pi_{n-1}(O_n)$, which we think of as the clutching function of a real vector bundle. By exactness of the long exact sequence in homotopy groups associated to the fibration $O_n \to U_n \to \Lambda_n$, it is equivalent to ask that β is in the image of the map $\pi_n(\Lambda_n) \to \pi_{n-1}(O_n)$ or to ask that it is in the kernel of the map $\pi_{n-1}(O_n) \to \pi_{n-1}(U_n)$. The first condition says that the vector bundle can be realized as a distribution of Lagrangian planes in T^*S^n defined along the zero section $S^n \hookrightarrow T^*S^n$, while the second condition says that the complexification of the vector bundle is trivial. Suppose that β is such a class and let $S(\beta) \in \pi_{n-1}(O_{n+1})$ be the image of β under the stabilization map S

induced by the inclusion $O_n \subset O_{n+1}$. By commutativity of the diagram below, observe that $S(\beta)$ lies in the kernel of the map $\pi_{n-1}(O_{n+1}) \to \pi_{n-1}(U_{n+1})$.

$$\begin{array}{ccc} \pi_{n-1}\big(O_n\big) & \longrightarrow & \pi_{n-1}\big(U_n\big) \\ \big\downarrow & & \big\downarrow \\ \pi_{n-1}\big(O_{n+1}\big) & \longrightarrow & \pi_{n-1}\big(U_{n+1}\big) \end{array}$$

However, $\ker\big(\pi_{n-1}(O_{n+1}) \to \pi_{n-1}(U_{n+1})\big) \simeq \ker\big(\pi_{n-1}(O) \to \pi_{n-1}(U)\big)$, since both homotopy groups lie in the stable range. This kernel can be computed from Bott periodicity. Indeed, $\Omega(U/O) \simeq \mathbb{Z} \times BO$ implies that $\pi_k(U/O) \simeq \pi_{k-2}(O)$ and therefore the groups appearing in the exact sequence $\pi_n(U/O) \to \pi_{n-1}(O) \to \pi_{n-1}(U)$ depend on the residue class of n mod 8 as follows.

n mod 8	$\pi_n(U/O)$	$\pi_{n-1}(O)$	$\pi_{n-1}(U)$
0	0	$\mathbb{Z}$	$\mathbb{Z}$
1	$\mathbb{Z}$	$\mathbb{Z}/2$	0
2	$\mathbb{Z}/2$	$\mathbb{Z}/2$	$\mathbb{Z}$
3	$\mathbb{Z}/2$	0	0
4	0	$\mathbb{Z}$	$\mathbb{Z}$
5	$\mathbb{Z}$	0	0
6	0	0	$\mathbb{Z}$
7	0	0	0

From the table we deduce that $\ker\big(\pi_{n-1}(O) \to \pi_{n-1}(U)\big) = 0$ except if $n \equiv 1$ or 2 mod 8 (in which case the kernel is isomorphic to $\mathbb{Z}/2$). It follows that if n is even and $n \not\equiv 2$ mod 8, then we necessarily have $S(\beta) = 0$, as claimed.

Remark 6.16 The simplest example of a caustic that cannot be simplified to consist only of folds occurs when $n = 2$ and V is the Hopf bundle on S^2. It is easy to check that in this case a Σ^{110} pleat is unavoidable, in addition to the Σ^{10} folds.

When n is odd, the same reasoning still shows that a necessary and sufficient condition for the simplification of singularities to be possible is that V is homotopic to one of the standard models V_Σ in the space of Lagrangian distributions. However, stable triviality of the underlying real vector bundle is not sufficient to guarantee that this condition is satisfied because $\pi_n(U_n) \neq 0$ and hence the map $\pi_n(\Lambda_n) \to \pi_{n-1}(O_n)$ need not be an injection.

We have only touched the surface of the homotopy theoretic calculations which are necessary to understand the formal condition obstructing the simplification of caustics. In the very concrete example of spheres considered above we were able to reason in a fairly hands-on manner. We believe that it should be possible to carry out a more systematic approach in the spirit of obstruction theory to study the general case.

6.5 Application: families of 1-dimensional Legendrians

We now turn to the second application discussed in Sect. 1.1. Our goal is to establish that higher singularities are unnecessary for the homotopy theoretic study of the space of Legendrian knots in the standard contact Euclidean $\mathbb{R}^3$. Recall that we think of $\mathbb{R}^3$ as the jet space $J^1(\mathbb{R}, \mathbb{R}) = \mathbb{R}(q) \times \mathbb{R}(p) \times \mathbb{R}(z)$ which comes equipped with the contact form $dz - pdq$. The Lagrangian projection is the map $\mathbb{R}^3 \to \mathbb{R}^2$, $(q, p, z) \mapsto (q, p)$ which corresponds to the forgetful map $J^1(\mathbb{R}, \mathbb{R}) \to T^*\mathbb{R}$. The front projection is the map $\mathbb{R}^3 \to \mathbb{R}^2$, $(q, p, z) \mapsto (q, z)$ which corresponds to the forgetful map $J^1(\mathbb{R}, \mathbb{R}) \to J^0(\mathbb{R}, \mathbb{R})$. The Reeb direction is $\partial/\partial z$ and it will also be useful to think of the projection along the Reeb direction $\mathbb{R}^2 \to \mathbb{R}$ which is the map $(q, z) \mapsto q$.

The fibres of the front projection form a Legendrian foliation $\mathcal{F}$ of $\mathbb{R}^3$. Recall that a Legendrian knot $f : S^1 \to \mathbb{R}^3$ is said to have mild singularities when the only singularities tangency of f with respect to $\mathcal{F}$ are folds and embryos. Folds are the generic Σ^{10} singularities of a single Legendrian knot and in the front projection correspond to cusps, see Fig. 44. Embryos are the generic Σ^{110} singularities of a 1-parametric family of Legendrian knots and in the front projection correspond to Type I Reidemeister moves, namely the instances of birth/death of two cusps. See Fig. 45.

Generically, a Legendrian knot only has folds and a 1-parametric family of Legendrian knots only has folds and embryos. However, the caustic of a family of Legendrian knots parametrized by a space of high dimension will generically be very complicated. It is therefore not a priori clear how the topology of the space of Legendrian knots $\mathcal{L}$ is related to that of the subspace $\mathcal{M} \subset \mathcal{L}$ consisting of those Legendrian knots whose singularities are mild. In Sect. 1.1 we defined a space of decorations $\widetilde{\mathrm{C}}(S^1)$ and a space $\mathcal{D}$ of pairs (f, D) consisting of a Legendrian knot with mild singularities $f \in \mathcal{M}$ together with a decoration $D \in \widetilde{\mathrm{C}}(S^1)$ of the singularities of f. See Fig. 46 for an example of a decoration D compatible with the standard front projection of the figure eight knot.

By composing the forgetful map $\mathcal{D} \to \mathcal{M}$ given by $(f, D) \mapsto f$ with the inclusion $\mathcal{M} \hookrightarrow \mathcal{L}$ we obtain a map $\mathcal{D} \to \mathcal{L}$. In this section we will prove the following result, which is a consequence of our parametric h-principle for the simplification of caustics.

Corollary 6.17 *The map $\mathcal{D} \to \mathcal{L}$ is a weak homotopy equivalence on each connected component.*

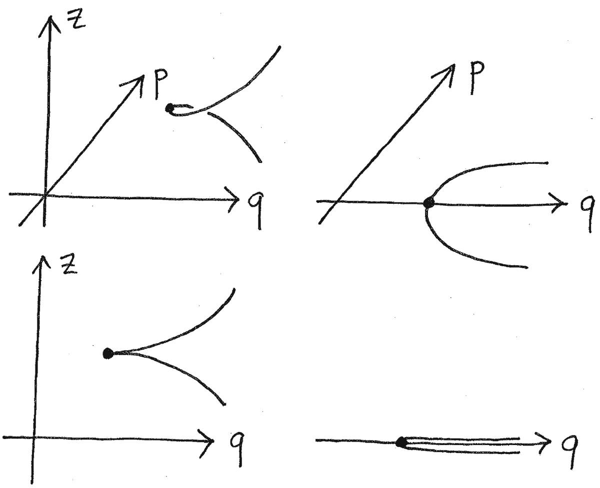

Fig. 44 The standard fold as seen from the Lagrangian and front projections (top right and bottom left respectively). If we project all the way down to $\mathbb{R} = \mathbb{R}(q)$ (bottom right), the germ of the resulting map is equivalent to that of $x \mapsto x^2$

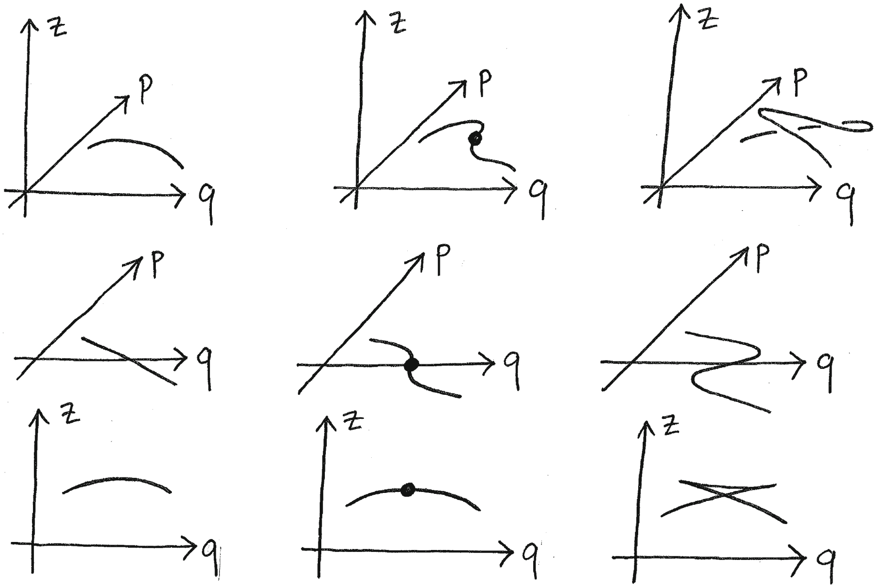

Fig. 45 The embryo singularity is illustrated in the middle column. We can picture it in the ambient contact $\mathbb{R}^3$ (top), in the Lagrangian projection (middle) and in the front projection (bottom). An embryo is a generically isolated singularity of a 1−parametric family of Legendrian knots, which we exhibit from left to right. The bottom row (which takes place in the front projection) gives us the familiar Reidemeister I move for Legendrian fronts

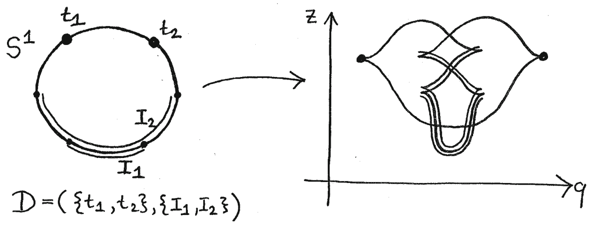

Fig. 46 An example of a decoration which consists of two points t_1, t_2 and two nested intervals $I_1 \subset I_2$

Remark 6.18 The decoration D is necessary because the inclusion $\mathcal{M} \hookrightarrow \mathcal{L}$ is not a homotopy equivalence, indeed $\pi_2(\mathcal{L}, \mathcal{M}) \neq 0$. To see this, let f^z be a family of Legendrian knots parametrized the closed unit 2-disk D^2 which has mild singularities everywhere except for a single Σ^{1110} singularity appearing in the interior. Then it is easy to see that the family $\{f^z\}_{z \in \partial D^2}$ represents a nontrivial element of $\pi_2(\mathcal{L}, \mathcal{M})$. The decoration D is designed to kill this homotopy group.

Remark 6.19 For an explicit example of the D^2 family mentioned in the previous remark, take a Legendrian front with only cusp singularities. The family will be localized near a single cusp. On the boundary ∂D^2, the family does the following. Start with your front, apply a Reidemesiter I move near the cusp, slide the result of the move over the cusp (as in the first four pictures of Fig. 47) and then eliminate it on the other side of the cusp with another Reidemeister I move to end back where you started. It is an instructive exercise to understand why we cannot assign compatible decorations to this S^1 family in a continuous way. Note also that the S^1-family can then be coned off to obtain a D^2 family by taking the distances between the fixed cusp point and the points where the Reidemesiter I moves are applied to be proportional to the radial coordinate r of $D^2 = \{re^{i\theta} : 0 \leq r \leq 1\}$. Everywhere except at the origin $0 \in D^2$ the fronts have mild singularities, while at the origin two arcs of embryos meet in a single Σ^{1110} singularity.

Remark 6.20 If $f \in \mathcal{L}$ is any Legendrian knot, then by a generic perturbation we may assume that the singularities $\Sigma \subset S^1$ of f consist only of a finite number of folds. Then f is compatible with the trivial decoration $D = (\{t_i\}, \{I_j\})$ consisting of $\{t_i\} = \Sigma$ and $\{I_j\} = \varnothing$. It follows that $\pi_0(\mathcal{D}) \to \pi_0(\mathcal{L})$ is surjective. However, is it easy to see that $\pi_0(\mathcal{D}) \to \pi_0(\mathcal{L})$ is not injective, since in the space $\mathcal{D}$ we are keeping track of the decoration D.

To prove Corollary 6.17 it suffices to show that $\pi_n(\mathcal{L}, \mathcal{D}) = 0$ for $n > 1$ and that $\pi_1(\mathcal{D}) \to \pi_1(\mathcal{L})$ is surjective. We deal with each of the statements separately.

Fig. 47 The family L_η. The parameter η runs from 0 to 1

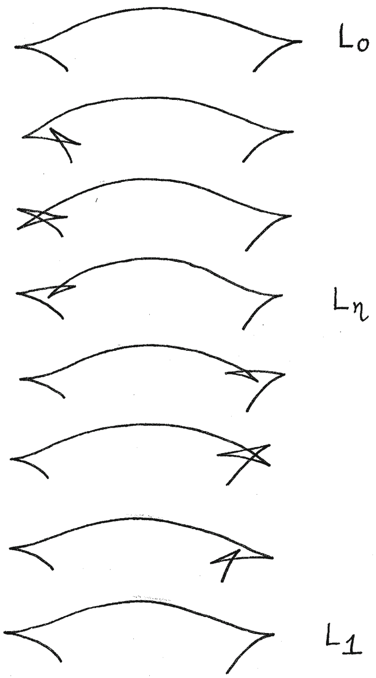

Proof that $\pi_n(\mathcal{L}, \mathcal{D}) = 0$ *for* $n > 1$ Let $\alpha \in \pi_n(\mathcal{L}, \mathcal{D})$ be any class. We can represent α by a map $F : D^n \to \mathcal{L}$ such that $F|_{\partial D^n}$ lifts to a map $\widetilde{F} : \partial D^n \to \mathcal{D}$. To conclude that $\alpha = 0$ we must show that there exist a homotopy $F_t : D^n \to \mathcal{L}$ which is fixed on $Op(\partial D^n)$ and such that $\widetilde{F} : \partial D^n \to \mathcal{D}$ extends to a lift $\widetilde{F}_1 : D^n \to \mathcal{D}$ of F_1.

We begin by examining the singularity locus of F on the boundary, which is the subset $\Sigma(F|_{\partial D^n}) \subset \partial D^n \times S^1$ consisting of all pairs $(z, s) \in S^{n-1} \times S^1$ such that the front of the Legendrian knot $F(z) : S^1 \to \mathbb{R}^3$ has a fold or embryo singularity at the point $s \in S^1$. Denote the map $(f, D) \mapsto D$ which forgets the knot but remembers the decoration by dec : $\mathcal{D} \to \widetilde{\mathrm{C}}(S^1)$. The family of decorations dec $\circ\, \widetilde{F} : \partial D^n \to \widetilde{\mathrm{C}}(S^1)$ induces a decomposition of the singularity locus $\Sigma(F|_{\partial D^n}) = \mathcal{C} \cup \mathcal{W}$, where $\mathcal{C}$ consists of folds and $\mathcal{W}$ consists of pairs of folds with opposite Maslov co-orientations together with the embryos that give rise to the birth/death of such pairs. The folds of $\mathcal{C}$ correspond to the points $t_1, \ldots, t_k$ and the pairs of folds or embryos of $\mathcal{W}$ correspond to the endpoints of the intervals $I_1, \ldots, I_m$. Note that the number m of intervals may vary with the parameter z but the number k of points is fixed since $n > 1$. After a generic perturbation we may assume that $\mathcal{C}$ and $\mathcal{W}$ are smooth codimension 1 submanifolds of $S^{n-1} \times S^1$ and moreover that the set of embryos $\mathcal{E}$ is a smooth codimension 1 submanifold of $\mathcal{W}$.

Our strategy is the following. The first step is to construct F_t near the boundary of the parameter space ∂D^n. This involves manually killing all the pairs of folds in $\mathcal{W}$. The next step is to extend the folds in $\mathcal{C}$ to the interior of the parameter space $\mathrm{int}(D^n)$. After these two preparatory steps we can apply the relative form of our parametric h-principle to construct F_t everywhere else so that the only additional singularities of the deformed family F_1 are the folds and embryos resulting from the wrinkling process. By construction the resulting map $F_1 : D^n \to \mathcal{M}$ will have an obvious lift to $\mathcal{D}$, completing the proof.

We now perform the first of these preparatory steps. The key idea, which appears repeatedly throughout the literature of the wrinkling philosophy, is that to kill a zig-zag one may create a very small new zig-zag near one end of the old zig-zag and then slowly let the new zig-zag take over, eventually killing the old zig-zag and replacing it. The newly created zig-zag does not bother us because it will end up completely contained in the interior of the parameter space D^n.

Fix a collar neighborhood $A \subset D^n$ of S^{n-1}, which we parametrize radially as $A = [0, 1) \times S^{n-1}$ with $0 \times S^{n-1}$ corresponding to ∂D^n. It will be convenient to assume that F is radially invariant on A, and indeed by means of an initial homotopy of F fixed on the boundary we can arrange it so that $F(\lambda, z) = F(0, z)$ for all $\lambda \in [0, 1)$ and all $z \in S^{n-1}$. Note then that $F(A) \subset \mathcal{M}$ and moreover $\Sigma(F|_A) = [0, 1) \times \Sigma(F_{\partial D^n})$. For an F satisfying this condition we establish the following preparatory result.

Lemma 6.21 (Preliminary arrangement near the boundary) *There exists a homotopy $F_t : D^n \to \mathcal{L}$ of $F = F_0$ such that the following properties hold.*

- *F_t is fixed on $Op\big(\partial D^n \cup (D^n \backslash A)\big)$.*
- *$F_t(A) \subset \mathcal{M}$.*
- *The folds in $\mathcal{C}$ are left untouched throughout the homotopy. To be more precise, the subset $[0, 1) \times \mathcal{C} \subset \Sigma(F_t|_A)$ does not vary with time.*
- *The pairs of folds in $\mathcal{W}$ are killed at the end of the homotopy. To be more precise, over each cylinder $[0, 1) \times z \times S^1 \subset A \times S^1$ the singularity locus $\Sigma(F_1|_A)$ contains arcs $a_1, \ldots, a_m$ whose interiors lie in $(0, 1) \times z \times S^1$ and whose endpoints lie in $0 \times z \times S^1$ and in fact consist precisely of the endpoints of the intervals $I_1, \ldots, I_m$. Moreover, each arc a_j consists everywhere of folds except at a single point in its interior, which is an embryo.*

Proof To construct the homotopy F_t we will use the $1-$parameter family of Legendrian fronts L_η exhibited in Fig. 47. Suppose that I_j is a non-degenerate interval appearing in the decoration $D = \mathrm{dec}\big(\widetilde{F}(z)\big)$ for some $z \in \partial D^n$. Assume moreover that I_j is isolated, meaning that there are no other intervals I_k contained inside I_j or containing I_j. In a neighborhood of $I_j \subset S^1$ the front

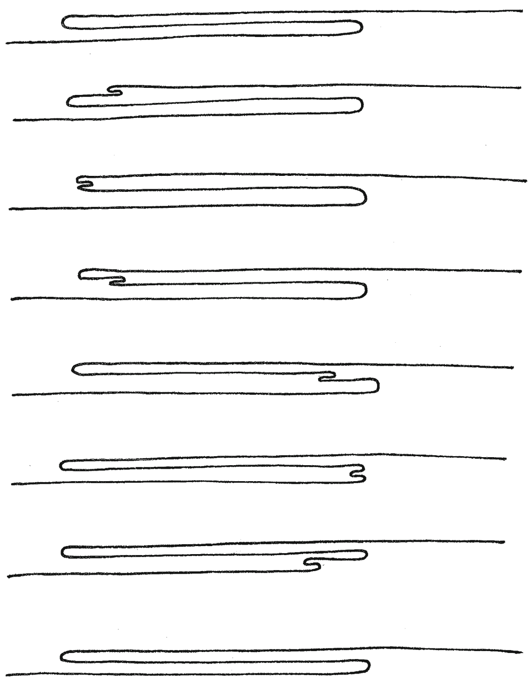

Fig. 48 The projection of Fig. 47 along the Reeb direction

of the knot $F(z)$ is equivalent to either the local model L_0 or to a flip of L_0 in the vertical direction, depending on the Maslov co-orientations. By replacing L_η by the vertical flip of L_η whenever this is needed, we may assume without loss of generality that the former case holds.

Note that the family of fronts L_η can be made to be C^0-close to the constant family L_0 and moreover we can arrange that the field of tangent lines to L_η is C^0-close to the field of tangent lines to L_0 (when both of these C^0-closeness properties hold for two given fronts we say that the fronts are C^1-close). Hence the resulting Legendrian isotopy can be made C^0-small. We can therefore think of the $1-$parameter family L_η as a Legendrian isotopy of $F(z)$ supported on $Op(I_j)$.

It is conceptually useful to understand the projection of the family L_η along the Reeb direction. The front L_0 projects down to a zig-zag. As the parameter η increases from 0 to 1, a new zig-zag is created just outside of I_j. We then make this new zig-zag bigger and bigger, until it takes over and replaces the old zig-zag, which has died by the time that η is close to 1. This process is illustrated in Fig. 48

To define F_t formally, let $\varphi : [0, 1] \to [0, 1]$ be a function such that the following properties hold.

- $\varphi = 0$ on $Op(\partial[0, 1])$.
- $\varphi = 1$ on $Op(\frac{1}{2})$.

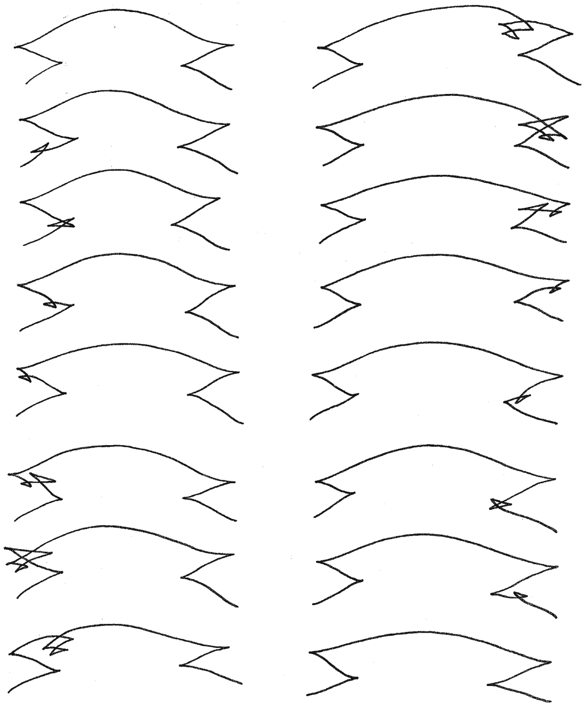

Fig. 49 A nesting of Fig. 47 for two intervals $I_k \subset I_j$

- φ is non-decreasing on $[0, \frac{1}{2}]$ and non-increasing on $[\frac{1}{2}, 1]$.

We define the homotopy F_t on $[0, 1) \times z \times Op(I_j)$ by the formula $F_t(\lambda, z, s) = L_{t\varphi(\lambda)}(s)$. Suppose next that there are two nested intervals $I_k \subset I_j$ with no other interval either contained or containing I_k or I_j. Then we define the homotopy F_t just like we did before, but using a nested version of the family L_η which we exhibit in Fig. 49. For more complicated configurations of intervals I_j we repeat this strategy but using the obvious model which is obtained by nesting the $1-$parameter family L_η (or its flip in the vertical direction) according to the nesting of the configuration of intervals.

The construction described above can be realized parametrically as $z \in S^{n-1}$ varies, as long as no interval I_j degenerates to a point. However, in a neighborhood of the locus $\mathcal{E} \subset \mathcal{W}$ of embryos we need a different local model so that the family L_η does not degenerate into a higher singularity. The 2-parametric family $L_{\eta,\tau}$ exhibited in Fig. 50 gets the job done. Let us first understand what the locus $\mathcal{W}$ looks like in a neighborhood of $\mathcal{E}$. Fix a connected component $\mathcal{W}_0 \subset \mathcal{W}$ and set $\mathcal{E}_0 = \mathcal{E} \cap \mathcal{W}_0$. Consider the image $\widehat{\mathcal{W}}_0$ of $\mathcal{W}_0$ under the projection $S^{n-1} \times S^1 \to S^{n-1}$. Note that $\widehat{\mathcal{W}}_0 \subset S^{n-1}$ is a smooth codimension 0 submanifold with boundary, that $\widehat{\mathcal{E}}_0 = \partial\widehat{\mathcal{W}}_0$ is the

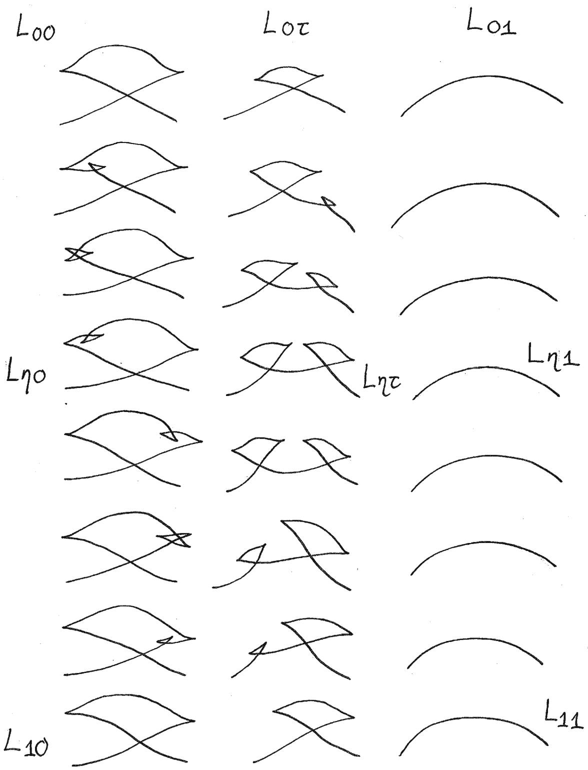

Fig. 50 The 2-parametric family $L_{\eta,\tau}$. The parameters η and τ both run from 0 to 1. To visualize $L_{\eta,\tau}$, start with the (constant) 1-parametric family $L_{\eta,1}$ which is the rightmost column of the figure. As you move towards the left the family undergoes Reidemesiter Type I moves at two separate points of the front, but towards the top one of the moves is cut off and towards the bottom the other move is cut off. As you keep moving to the left you fit the newly created pieces of the front together to obtain the previously defined 1-parametric family L_η, which in the present figure sits as the leftmost vertical column $L_{\eta,0}$

image of $\mathcal{E}_0$, that the map $\mathcal{W}_0 \to \widehat{\mathcal{W}}_0$ is a 2 to 1 cover away from $\mathcal{E}_0$ and that along $\mathcal{E}_0$ the map $\mathcal{W}_0 \to \widehat{\mathcal{W}}_0$ has folds. In particular, the restriction $\mathcal{E}_0 \to \widehat{\mathcal{E}}_0$ is an embedding, see Fig. 51.

Let $\widehat{\mathcal{E}}_0 \times (0, 1)$ be a collar neighborhood of $\widehat{\mathcal{E}}_0 = \widehat{\mathcal{E}}_0 \times \frac{1}{2}$ in S^{n-1} such that $\widehat{\mathcal{E}}_0 \times (0, \frac{1}{2}] \subset \widehat{\mathcal{W}}_0$. Given $e \in \mathcal{E}_0$, let $\widehat{e}$ be its image in $\widehat{\mathcal{E}}_0$ and let $z_t \in S^{n-1}$, $t \in (0, 1)$ correspond to the arc $\widehat{e} \times (0, 1) \subset \widehat{\mathcal{E}}_0 \times (0, 1)$. Then the 1-parametric family $F(z_t)$ is equivalent in a neighborhood of the embryo point e to the 1-parametric family $L_{0,\tau}$ exhibited in the top row of Fig. 50 (or to its flip in the vertical direction). Note that the 1-parametric family $L_{0,\tau}$ fits into the 2-parametric family $L_{\eta,\tau}$ shown in Fig. 50, corresponding to the side $0 \times [0, 1]$ of the square of parameters $(\eta, \tau) \in [0, 1] \times [0, 1]$.

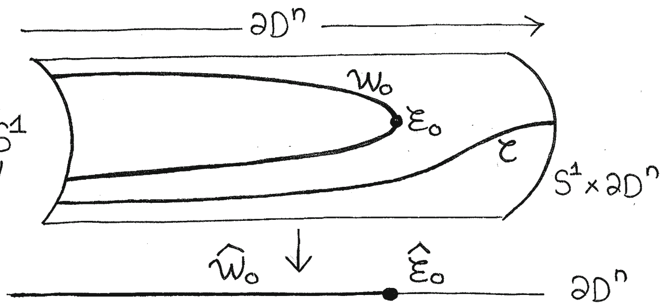

Fig. 51 The local geometry of the projection $\mathcal{W} \to \partial D^n$

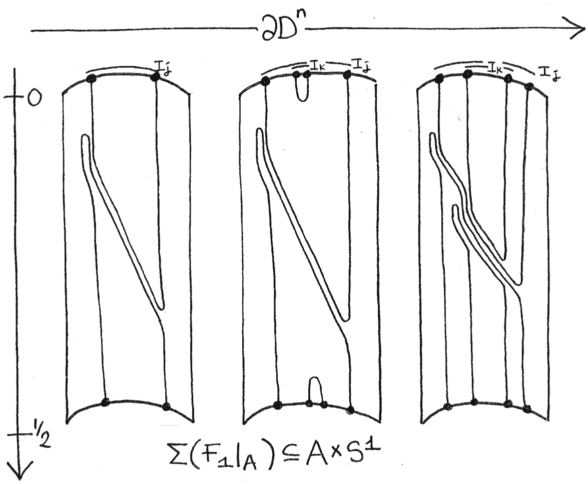

Fig. 52 The singularity locus of F_1 on $\partial D^n \times [0, \frac{1}{2}) \times S^1 \subset A \times S^1$, which is one-half of the full locus $\Sigma(F_1|_A)$

Observe that the family $L_{\eta,\tau}$ can be taken to be C^1-close to the family $L_{0,\tau}$ which is constant in η. We can therefore think of the 2-parametric family $L_{\eta,\tau}$ as a C^0-small Legendrian isotopy of the 1-parametric family $F(z_t)$ supported in a neighborhood of the embryo point. Note that $L_{\eta,0} = L_\eta$, so the isotopy is compatible with our previous isotopy. Notice also that $L_{\eta,1}$ is constant. We can therefore define the homotopy F_t by the formula $F_t(\eta, z, s) = L_{t\varphi(\eta),\tau}(s)$, where $z = (\widehat{e}, \tau) \in \widehat{\mathcal{E}}_0 \times (0, 1)$. The construction can be realized parametrically in z, see Fig. 52 for an illustration. The construction can also be realized with any configuration of intervals, by nesting the families shown in Figs. 49 and 50 according to the nesting of the intervals. This completes the proof of Lemma 6.21. □

The next step is to extend the cusp locus $\mathcal{C}$ to the interior of the parameter space D^n. This is achieved by a second preparatory lemma. For notational convenience, we now forget about our old family and use the letter F to denote the new family F_1 produced by Lemma 6.21. In particular, all of the properties listed in the conclusion of Lemma 6.21 are satisfied by F.

Lemma 6.22 (Preliminary arrangement in the interior) *There exists a homotopy $F_t : D^n \to \mathcal{L}$ of $F = F_0$ such that the following properties hold.*

- *F_t is fixed on A.*
- *The singularity locus $\Sigma(F_1) \subset D^n \times S^1$ contains a properly embedded submanifold with boundary $\mathcal{I}$ of codimension 1 in $D^n \times S^1$ which consists entirely of folds and such that $\mathcal{I} \cap (A \times S^1) = \mathcal{C} \times [0, 1)$.*

Remark 6.23 Since F_t is fixed on A, $\Sigma(F_1)$ also contains the properly embedded codimension 1 submanifold with boundary $\mathcal{K}$ formed by the arcs a_j which kill $\mathcal{W}$. In addition to $\mathcal{I}$ and $\mathcal{K}$, the singularity locus $\Sigma(F_1)$ may have other components, but we will not care about them because they are all homotopically trivial and contained in $\text{int}(D^n) \times S^1$.

Proof We assume that $\mathcal{C} \neq \varnothing$, otherwise the Lemma is trivial. Recall that the space of decorations $\widetilde{\mathrm{C}}(S^1)$ is fibered over the (unordered) configuration space of points on the circle $\mathrm{C}(S^1) = \bigsqcup_k \mathrm{C}_k(S^1)$. The map is $\big(\{t_j\}, \{I_i\}\big) \mapsto \{t_j\}$ and its fibers are contractible. Denote by $\text{conf} : \mathcal{D} \to \mathrm{C}(S^1)$ the composition of $\text{dec} : \mathcal{D} \to \widetilde{\mathrm{C}}(S^1)$ with the fibration $\widetilde{\mathrm{C}}(S^1) \to \mathrm{C}(S^1)$. We claim that the map $\text{conf} \circ \widetilde{F} : \partial D^n \to \mathrm{C}(S^1)$ extends to a map $c : D^n \to \mathrm{C}(S^1)$.

First observe that each component $\mathrm{C}_k(S^1)$ of $\mathrm{C}(S^1)$ is homotopy equivalent to S^1. Hence for $n > 2$ there is nothing to prove because $\pi_{n-1}(S^1) = 0$. If $n = 2$, then we need to justify the claim. Write $H^*(\partial D^2 \times S^1; \mathbb{R}) = \mathbb{R}[x, y]/(x^2, y^2)$, where x is Poincaré dual to $\partial D^2 \times pt$ and y is Poincaré dual to $pt \times S^1$. Consider the Gauss map $G(dF) : D^2 \times S^1 \to S^1$, $(z, s) \mapsto G\big(dF(z)\big)(s)$, where $\Lambda(\mathbb{R}^3) = \mathbb{R}^3 \times S^1$ and we project away the $\mathbb{R}^3$ factor. Explicitly, an angle θ corresponds to the line field spanned by $\cos(\theta)\partial/\partial p + \sin(\theta)(\partial/\partial z + p\partial/\partial q)$. Observe that $\big(\partial D^2 \times S^1\big) \cap G(dF)^{-1}\big(\text{span}(\partial/\partial p)\big) = \mathcal{C} \cup \mathcal{W}$. Observe also that the fundamental class of $\mathcal{C}$ is Poincaré dual to $kx + ly$ for some $l \in \mathbb{Z}$, where we recall that k is the number of points $t_1, \ldots, t_k$ in the decorations $\text{dec} \circ F(z)$. If we write $i : \partial D^2 \times S^1 \hookrightarrow D^2 \times S^1$ for the inclusion and denote by $u \in H^1(S^1; \mathbb{R})$ the class which is Poincaré dual (PD) to a point, then we have

$$\begin{aligned} kx + ly &= \text{PD}[\mathcal{C}] = \text{PD}[\mathcal{C} \cup \mathcal{W}] = \text{PD}\left[\big(G(dF) \circ i\big)^{-1}\big(\text{span}(\partial/\partial p)\big)\right] \\ &= (G(dF) \circ i)^* u = i^*\big(G(dF)^* u\big). \end{aligned}$$

However, $i^* : H^*(D^2 \times S^1; \mathbb{R}) \to H^*(\partial D^2 \times S^1; \mathbb{R})$ has image generated by x. It follows that $l = 0$ and hence that $\mathcal{C}$ is an embedded curve in $\partial D^2 \times S^1$

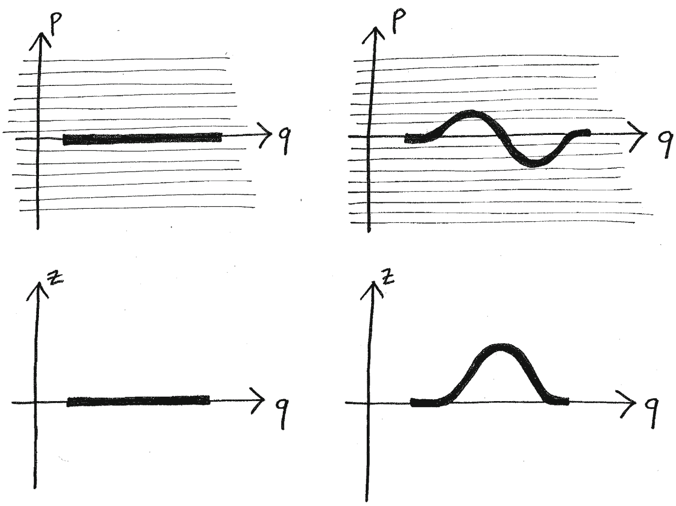

Fig. 53 The local model for a creation of double folds when the Legendrian is almost tangent to the foliation. For $\mathcal{F} = \text{span}(\partial/\partial q + p\partial/\partial z)$ the model obviously creates a double fold, hence by stability it also does so for nearby $\mathcal{F}$

which is homologous to $k[\partial D^2 \times pt]$. Note then that $\mathcal{C}$ has necessarily k components, each of which is homologous to $[\partial D^2 \times pt]$. It is now a triviality to check that $\text{conf} \circ \widetilde{F} : \partial D^2 \to \text{C}(S^1)$ extends to a map $c : D^2 \to \text{C}(S^1)$, as claimed.

Choose then such an extension c and assume without loss of generality that c is radially constant in the annulus $A \subset D^n$. Choose also a tangential rotation $G_t : D^n \times S^1 \to \Lambda(\mathbb{R}^3)$ of the family F such that the following properties hold.

- G_t is fixed on A.
- $G_1 = \partial/\partial p$ on the subset $\mathcal{I} = \{(z, t) : \; t \in c(z)\} \subset D^n \times S^1$.

Using the parametric version of theorem Theorem 4.2 (which in the 1-dimensional case is the same as Theorem 4.10 since all rotations are simple) we obtain a homotopy F_t of the family F which is fixed on A and such that $G(dF_t)$ is C^0-close to G_t on $Op(\mathcal{I})$. The family F_1 does not quite have folds along $\mathcal{I}$, but $G(dF_1)$ is almost parallel to $\partial/\partial p$ on $Op(\mathcal{I})$ and $F_1|_A$ does have folds along $\mathcal{I} \cap A$. By implanting the local model for the creation of a pair of folds exhibited in Fig. 53 into F_1 we can arrange it so that the new family does have folds precisely along $\mathcal{I}$. Moreover, we can arrange it so that the new family agrees with the old family inside A. Away from $Op(A \cup \mathcal{I})$ the singularities of F might be a mess but we don't care. The proof of Lemma 6.22 is complete. □

We can now conclude the proof that $\pi_n(\mathcal{L}, \mathcal{D}) = 0$ for $n > 2$. Given $\alpha \in \pi_{n-1}(\mathcal{L}, \mathcal{D})$ represented by a family F, we can apply Lemmas 6.21

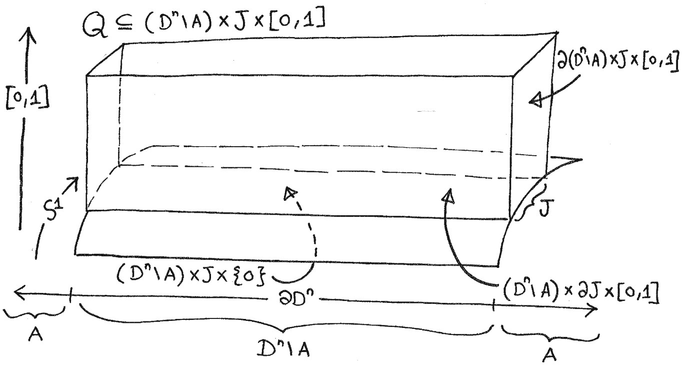

Fig. 54 The cube Q sits like an open box inside $(D^n\backslash A) \times J \times [0, 1]$

and 6.22 and replace F with the family obtained after performing the two preliminary arrangements, in that order. For the new F, we claim the existence of a family of tangential rotations $G_t : D^n \times S^1 \to S^1$ of the family F such that the following properties hold.

- G_t is fixed on $Op\big((\partial D^n \times S^1) \cup \mathcal{K} \cup \mathcal{I})\big)$.
- $G_1 \pitchfork \mathcal{F}$ away from $\mathcal{K} \cup \mathcal{I}$.

To verify the claim, we begin by considering the restriction of the Gauss map $G(dF) : D^n \times S^1 \to S^1$ to the annulus A. Note that by construction the lift $\widetilde{F} : \partial D^n \times S^1 \to \mathcal{D}$ extends to a lift $\widetilde{F} : A \times S^1 \to \mathcal{D}$, where we assign intervals to the new pair of folds created by the family L_η. The intervals $I_1, \dots, I_m$ of the decoration $\text{dec} \circ \widetilde{F} : A \to \widetilde{\mathrm{C}}(S^1)$ which do not correspond to pairs of folds in $\mathcal{K}$ give us a homotopically canonical deformation $G_t : A \times S^1 \to S^1$ of $G(dF)|_A$ such that G_t is fixed on $Op\big((\partial D^n \times S^1) \cup \mathcal{K} \cup (\mathcal{I} \cap A)\big)$ and such that $G_1 \pitchfork \mathcal{F}$ away from $\mathcal{K} \cup (\mathcal{I} \cap A)$. Together with the requirement that G_t is fixed near $\mathcal{I}$, this defines the map $(z, s, t) \mapsto G_t(z, s)$ on $(A \times S^1 \times [0, 1]) \cup (D^n \times S^1 \times 0) \cup (Op(\mathcal{I}) \times [0, 1])$. Each connected component of the complement of $Op(\mathcal{I})$ in $(D^n\backslash A) \times S^1$ is diffeomorphic to $(D^n\backslash A) \times J$, where J is a closed interval and the diffeomorphism is of the form $(z, s) \mapsto \big(z, \psi(z, s)\big)$. Consider the cube

$$Q = \partial(D^n\backslash A) \times J \times [0, 1] \ \cup \ (D^n\backslash A) \times J \times 0 \ \cup \ (D^n\backslash A) \times \partial J \times [0, 1]$$

which we think of as a subset of $(D^n\backslash A) \times S^1 \times [0, 1]$ via the above diffeomorphism. See Fig. 54. Note that Q has boundary

$$\partial Q = \partial(D^n\backslash A) \times J \times 1 \ \cup \ (D^n\backslash A) \times \partial J \times 1.$$

The homotopy G_t defined thus far gives a map of pairs $(Q, \partial Q) \to (S^1, S^1\backslash pt)$, where $pt = \text{span}(\partial/\partial p)$. Since $\pi_j(S^1, S^1\backslash pt) = 0$ for $j > 1$, there exists a homotopy of pairs relative to the boundary so that at the end of the homotopy the image is disjoint from $\text{span}(\partial/\partial p)$. This is precisely what we needed to define G_t on the remaining part of $D^n \times S^1 \times [0, 1]$ so that the required conditions are satisfied.

Now that we have established the existence of such a tangential rotation G_t, we can invoke Theorem 6.9 to construct a homotopy $F_t : D^n \to \mathcal{L}$ of F which is fixed on $Op\big((\partial D^n \times S^1) \cup \mathcal{K} \cup \mathcal{I}\big)$ and such that away from $\mathcal{K} \cup \mathcal{I}$ the singularities of the family F_1 consist of a finite union of fibered nested regularized wrinkles. It only remains to show that $\widetilde{F} : \partial D^n \to \mathcal{D}$ extends to a lift of F_1 to $\mathcal{D}$. However, this is clear because to the folds of $\mathcal{I}$ and to the pairs of folds of $\mathcal{K}$ we can assign points and intervals in the obvious way, while away from $\mathcal{K} \cup \mathcal{I}$ the singularities of F_1 consist only of the pairs of points in the fibered regularized wrinkles, to which intervals can be canonically assigned. This completes the proof that $\pi_n(\mathcal{L}, \mathcal{D}) = 0$ for $n > 1$. □

Proof that $\pi_1(\mathcal{D}) \to \pi_1(\mathcal{L})$ is surjective. Let $\alpha \in \pi_1(\mathcal{L})$ be any class. We can represent α by a map $F : [0, 1] \to \mathcal{L}$ such that $F(0) = F(1) = f_0$. Choose any decoration D_0 which is compatible with f_0. We must show that there exists a homotopy $F_t : [0, 1] \to \mathcal{L}$ of $F = F_0$ such that $F_t(0) = F_t(1) = f_0$ for all $t \in [0, 1]$ and such that $F_1 : [0, 1] \to \mathcal{L}$ lifts to a map $\widetilde{F}_1 : [0, 1] \to \mathcal{D}$ with $\widetilde{F}_1(0) = \widetilde{F}_1(1) = (f_0, D_0)$.

Write $D_0 = (\{t_i\}, \{I_j\})$ for points $t_1, \ldots, t_k \in S^1$ and non-degenerate intervals $I_1, \ldots, I_m \subset S^1$. Let $K = \{t_1, \ldots, t_k\} \cup \partial I_1 \cup \cdots \cup \partial I_m \subset S^1$. Observe that the Gauss map $G(dF) : [0, 1] \times S^1 \to \Lambda(\mathbb{R}^3)$ of the family F satisfies $G(dF) = \text{span}(\partial/\partial p)$ on $\partial[0, 1] \times K$. Let $G_t : [0, 1] \times S^1 \to \Lambda(\mathbb{R}^3)$ be a tangential rotation of the family F such that G_t is fixed on $Op(\partial[0, 1] \times S^1)$ and such that $G_1 = \text{span}(\partial/\partial p)$ on $[0, 1] \times K$. Using Theorem 4.10 as above, we can construct a homotopy $F_t : [0, 1] \to \mathcal{L}$ which is fixed near $\partial[0, 1]$ and such that $G(dF_1)$ is C^0-close to $\text{span}(\partial/\partial p)$ on $[0, 1] \times K$.

By the insertion of the local model in Fig. 53 we can assume that F_1 actually has folds along $[0, 1] \times K$. Theorem 6.9 can then be used to further homotope F_1 rel $Op\big((\partial[0, 1] \times S^1) \cup ([0, 1] \times K)\big)$ so that on the complement of $[0, 1] \times K$ the only singularities are fibered nested regularized wrinkles. This new $F_1 : [0, 1] \to \mathcal{L}$ admits a canonical lift $\widetilde{F}_1 : [0, 1] \to \mathcal{D}$ by assigning intervals to the pairs of points in the fibered regularized wrinkles. This completes the proof that $\pi_1(M) \to \pi_1(\mathcal{L})$ is surjective. Hence Corollary 6.17 is also proved. □

We conclude this section with a remark. Proving that $\pi_n(\mathcal{L}, \mathcal{D}) = 0$ for $n > 1$ amounts to solving the following lifting problem. Given a diagram of the form

$$\begin{array}{ccc} \mathcal{D} & \longrightarrow & \mathcal{L} \\ \uparrow & & \uparrow \\ S^{n-1} & \longrightarrow & D^n \end{array}$$

we must show that there exists a map $D^n \to \mathcal{D}$ such that when added to the above diagram all compositions commute up to a homotopy fixed on S^{n-1}. The proof of Corollary 6.17 achieves this, but in fact proves slightly more. Because all of the theorems invoked hold in C^0-close form and because all of the local models used are C^0-small perturbations, it follows that the composition of the lift $D^n \to \mathcal{D}$ with the map $\mathcal{D} \to \mathcal{L}$ can be taken to be C^0-close to the original map $D^n \to \mathcal{L}$. The analogous C^0-approximation result holds for the corresponding lifting property for proving that $\pi_1(\mathcal{D}) \to \pi_1(\mathcal{L})$ is surjective.

6.6 Final remarks

We conclude our discussion with a couple of remarks.

Remark 6.24 All of the results proved in this paper also hold for immersed rather than embedded Lagrangians or Legendrians $f : L \to M$. The reason is that from the onset one can replace M with T^*L or $J^1(L, \mathbb{R})$ by choosing a Weinstein neighborhood of the immersion, thereby reducing to the embedded case. The only difference in the conclusion is that the resulting exact homotopy of regular Lagrangian or Legendrian immersions will not be induced by an ambient Hamiltonian isotopy in the original manifold M.

Remark 6.25 It is worth giving the following warning. If the singularities of a regular Lagrangian or Legendrian embedding $g : L \to M$ with respect to $\mathcal{F}$ consist only of a disjoint union of regularized wrinkles (or double folds), then the singularity locus is quite simple in the source. However, in the target the image of the singularity locus is likely to be very complicated. It would be interesting to know how much of the rigidity of a Lagrangian or Legendrian embeddeding can be read from this image.

Remark 6.26 From Theorems 5.1 and 5.3 we can also deduce a full h-principle for directed embeddings of wrinkled Lagrangian or Legendrian embeddings analogous to the one deduced by Eliashberg and Mishachev from their wrinkled embeddings theorem [15]. Before we can state it, we need a definition.

Definition 6.27 For any Lagrangian or Legendrian embedding $f : L \to M$ and for any subset $A \subset \Lambda(M)$, we say that f is A-directed if $\mathrm{im}\big(G(df)\big) \subset A$.

The result is then the following.

Theorem 6.28 *Let $f : L \to M$ be a Lagrangian or Legendrian embedding, let $A \subset \Lambda(M)$ be any open subset and assume that there exists a tangential rotation G_t of f such that $im\big(G_1\big) \subset A$. Then there exists an exact homotopy of wrinkled Lagrangian or Legendrian embeddings $f_t : L \to M$ such that f_1 is A-directed.*

This theorem holds in C^0-close, relative and parametric forms and follows immediately from Theorems 5.1 and 5.3 since A is assumed to be open.

Acknowledgements I am very grateful to my advisor Yasha Eliashberg for insightful guidance throughout this project. I would also like to thank Laura Starkston for reading carefully the first draft of this paper and offering numerous remarks and corrections which have greatly improved the exposition. I am indebted to the ANR Microlocal group who held a workshop in January 2017 to dissect an early version of the paper and in particular to Sylvain Courte and Alexandre Vérine who spotted several mistakes in the proof of the local wrinkling lemma and made useful suggestions for fixing them. My gratitude also goes to Roger Casals, Sander Kupers, Emmy Murphy, Oleg Lazarev and Kyler Siegel for many helpful discussions surrounding the general notion of flexibility. Finally, many thanks to the referee for numerous helpful comments and corrections.

References

1. Alvarez-Gavela, D.: Refinements of the holonomic approximation lemma. Algebr. Geom. Topol. **18**, 2265–2303 (2018)
2. Arnold, V.I.: A characteristic class entering in quantization conditions. Funct. Anal. Appl. **1**(1), 1–13 (1967)
3. Arnold, V.I.: Singularities of Caustics and Wave Fronts. Kluwer Academic Publishers, Dordrecht (1990)
4. Arnold, V.I., Gusein-Zade, S.M., Varchenko, A.N.: Singularities of Differentiable Maps, vol. I. Springer, Berlin (1985)
5. Arnold, V.I., Gusein-Zade, S.M., Varchenko, A.N.: Singularities of Differentiable Maps, vol. II. Springer, Berlin (1988)
6. Boardman, J.M.: Singularities of differentiable maps. Publ. Mathématiques de l'I.H.É.S **33**(2), 21–57 (1967)
7. Ekholm, T.: Morse flow trees and Legendrian contact homology in 1-jet spaces. Geom. Topol. **11**, 1083–1224 (2007)
8. Eliashberg, Y.M.: On singularities of folding type. Izv. Akad. Nauk SSSR Ser. Mat. **4**(5), 1119–1134 (1970)
9. Eliashberg, Y.M.: Surgery of singularities of smooth mappings. Izv. Akad. Nauk SSSR Ser. Mat. **36**(6), 1321–1347 (1972)
10. Eliashberg, Y.M., Mishachev, N.M.: Wrinkling of smooth mappings and its applications. I. Invent. Math. **130**(2), 345–369 (1997)
11. Eliashberg, Y.M., Mishachev, N.M.: Wrinkling of smooth mappings. III. Foliations of codimension greater than one. Topol. Methods Nonlinear Anal. **11**(2), 321–350 (1998)
12. Eliashberg, Y.M., Mishachev, N.M.: Wrinkling of smooth mappings. II. Wrinkling of embeddings and K. Igusa's theorem. Topology **39**(4), 711–732 (2000)
13. Eliashberg, Y.M., Mishachev, N.M.: Holonomic approximation and Gromov's h-principle. arXiv:math/0101196
14. Eliashberg, Y.M., Mishachev, N.M.: Introduction to the h-principle. Graduate Studies in Mathematics, vol. 48. American Mathematical Society, Providence (2002)

15. Eliashberg, Y.M., Mishachev, N.M.: Wrinkled embeddings, foliations, geometry, and topology. Contemp. Math. **498**, 207–232 (2009)
16. Eliashberg, Y.M., Mishachev, N.M.: The space of framed functions is contractible. In: Pardalos, P., Rassias, T. (eds.) Essays in Mathematics and its Applications. Springer, Berlin, Heidelberg (2012)
17. Eliashberg, Y.M., Mishachev, N.M.: Topology of spaces of S-immersions. In: Perspectives in Analysis, Geometry, and Topology, Progress in Mathematics, vol. 296, Birkhäuser, Boston, MA (2012)
18. Eliashberg, Y.M., Galatius, S., Mishachev, N.M.: Madsen–Weiss for geometrically minded topologists. Geom. Topol. **15**(1), 411–472 (2011)
19. Eliashberg, Y.M.: Weinstein manifolds revisited, arXiv:1707.03442
20. Entov, M.: Surgery on Lagrangian and Legendrian singularities. Geom. Funct. Anal. **9**(2), 298–352 (1999)
21. Entov, M.: On the necessity of Legendrian fold singularities. IMRN International Mathematics Research Notices, No. 20 (1998)
22. Emmanuel, F., Pushkar', E.P.: Non cancellation of singularities on wave fronts. C. R. Acad. Sci. Paris. Ser. I Math. **327**(8), 827–831 (1998)
23. Emmanuel, F., Pushkar', E.P.: Morse theory and global coexistence of singularities on wave fronts. J. Lond. Math. Soc. **74**(2), 527–544 (2006)
24. Gromov, M.L.: Stable maps of foliations into manifolds. Izv. Akad. Nauk SSSR Ser. Mat. **33**(4), 671 (1969)
25. Gromov, M.L.: Convex integration of partial differential relations. Izv. Akad. Nauk SSSR Ser. Mat. **37**, 329–343 (1973)
26. Gromov, M.L.: Partial Differential Relations. Springer, Berlin (1986)
27. Hirsch, M.: Immersions of manifolds. Trans. Am. Math. Soc. **93**, 242–276 (1959)
28. Igusa, K.: Higher singularities are unnecessary. Ann. Math. **119**, 1–58 (1984)
29. Igusa, K.: The space of framed functions. Trans. Am. Math. Soc. **301**(2), 431–477 (1987)
30. Kragh, T.: Parametrized ring-spectra and the nearby Lagrangian conjecture, arXiv:1107.4674 (2011)
31. Kupers, S.: Three applications of delooping applied to H-principles, arXiv:1701.06788
32. Lurie, J.: On the classification of topological field theories, arXiv:0905.0465
33. Morin, B.: Formes canoniques des singularites d'une application differetiable. C. R. Acad. Sci. Paris **260**, 5662–5665 (1965)
34. Murphy, E.: Loose Legendrian embeddings in high dimensional contact manifolds, arXiv:1201.2245 (2012)
35. Nadler, D.: Arboreal singularities, arXiv:1309.4122
36. Nadler, D.: Non-characteristic expansions of Lagrangian singularities, arxiv:1507.01513
37. Phillips, A.: Submersions of open manifolds. Topology **6**, 171–206 (1967)
38. Rourke, C., Sanderson, B.: The compression theorem I–II. Geom. Topol. **5**, 399–429 (2001). (431–440)
39. Rourke, C., Sanderson, B.: The compression theorem III: applications. Algebr. Geom. Topol. **3**, 857–872 (2003)
40. Smale, S.: The classification of immersions of spheres in Euclidean spaces. Ann. Math. **2**(69), 327–344 (1959)
41. Spring, D.: Convex Integration Theory. Birkhäuser, Basel (1998)
42. Spring, D.: Directed embeddings and the simplification of singularities. Commun. Contemp. Math. **4**, 107–144 (2002)
43. Spring, D.: Directed embeddings of closed manifolds. Commun. Contemp. Math. **7**, 707–725 (2005)
44. Starkston, L.: Arboreal singularities in Weinstein Skeleta, arXiv:1707.03446
45. Thom, R.: Les singularités des applications différentiables. Annales de l'Institut Fourier **6**(6), 43–87 (1956)

46. Thurston, W.: The theory of foliations of codimension greater than one. Commentarii Mathematici Helvetici **49**, 214–231 (1974)