Holonomic Approximation

May 10, 2025

Kevin Tang

1 Acknowledgements

Most statements will be copied verbatim from [CEM24] and [AG24]. A lot of my understanding and inspiration also comes from watching Alvarez-Gavela's lecture on holonomic approximation on YouTube [AGY24], so many thanks to him and his notes in particular.

2 Review

Recall the following definitions from last lecture.

Definition 2.1. Given a (smooth) map $f : \mathbb{R}^n \to \mathbb{R}^q$ and a point $x \in \mathbb{R}^n$, the string of derivatives of f up to order r

$$J^r_f(x) = \left(f(x), f'(x), \dots, f^{(r)}(x)\right) \in \mathbb{R}^{qN_r}$$

where $N_r = \frac{(n+r)!}{n!r!}$ is called the *r-jet* of f at x. Here $f^{(s)}$ consists of all partial derivatives $D^\alpha f$, $\alpha = (\alpha_1, \dots, \alpha_n)$, $|\alpha| = \alpha_1 + \cdots + \alpha_n = s$, written lexicographically.

Remark 2.2. Note that the r-jet of f contains the same information as the rth order Taylor polynomial of f.

Definition 2.3. If we allow x to vary, we can view J^r_f as a section of $\mathbb{R}^n \times \mathbb{R}^{qN_r} := J^r(\mathbb{R}^n, \mathbb{R}^q)$, which we'll call the *space of r-jets* of sections $\mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^q$.

Remark 2.4. Observe that arbitrary sections of $J(\mathbb{R}^n, \mathbb{R}^q)$ need not be realizable as the r-jet of a map $f : \mathbb{R}^n \to \mathbb{R}^q$.

In the Euclidean case, we saw that jets of functions (which we can view as sections $\mathbb{R}^n \times \mathbb{R}^q \to \mathbb{R}^n$) can be uniquely characterized by their rth order Taylor polynomial expansion at a given point. If we want to generalize the notion of a jet to a section of an arbitrary fibration $X \to V$, then, a natural way to do it would be to identify sections with their rth order Taylor approximation in a given trivialization.

Definition 2.5. Let $v \in V$. Two local sections $f : \mathcal{O}p\, v \to X$ and $g : \mathcal{O}p\, v \to X$ of the fibration $X \to V$ are called *r-tangent* at the point v if $f(v) = g(v)$ and

$$J^r_{\varphi_* f}(\varphi(v)) = J^r_{\varphi_* g}(\varphi(v))$$

for a local trivialization $\varphi : U \to \mathbb{R}^n \times \mathbb{R}^q$ of X in a neighborhood U of the point $x = f(v) = g(v)$. Here $\varphi_* f$ and $\varphi_* g$ are the images of the sections f and g.

Definition 2.6. Let $X \to V$ be a fibration. The r-tangency class of a section $f : \mathcal{O}p\, v \to X$ at a point $v \in V$ is called the *r-jet of f* at v and denoted by $J^r_f(v)$. The set of all r-jets of sections of the fibration $p : X \to V$ is denoted by $X^{(r)}$, and comes equipped with a projection $p^r : X^{(r)} \to V$.

Remark 2.7. A helpful way to visualize a section $\sigma : V \to X^{(r)}$ is to imagine for each $v \in V$, the graph of the rth order Taylor polynomial in an infinitesimal neighborhood U above v. In particular, a section of $X^{(1)}$ can be visualized as a smooth surface with a non-vertical (and not necessarily tangent) plane at each point.[1]

We commented earlier that arbitrary sections of the jet space need not arise as jets of sections on their domain, even if locally this is the case. As such, we give a special name to the sections that *can* be realized in such a way.

Definition 2.8. Given a section $\sigma : V \to X^{(r)}$, we will denote by bs σ the underlying section $p_0^r \circ \sigma : V \to X$. A section $\sigma : V \to X^{(r)}$ is called holonomic if $\sigma = J^r_{\text{bs}\,\sigma}$. In particular, holonomic sections $\mathbb{R}^n \to J^r(\mathbb{R}^n, \mathbb{R}^q)$ have the form

$$x \mapsto \left(x, f(x), f'(x), \dots, f^{(r)}(x)\right)$$

so that holonomic sections of $X^{(r)} \to V$ are in one-to-one correspondence with sections of $X \to V$.

Definition 2.9. A *differential relation* of order r imposed on sections $f : V \to X$ of a fibration $X \to V$ is a subset $\mathcal{R}$ of the jet space $X^{(r)}$. It is *open/closed* if it is open/closed as a subset of $\mathcal{R}$. A section $\sigma : V \to \mathcal{R} \subset X^{(r)}$ is called a *formal solution* of the differential relation $\mathcal{R}$. A *(genuine) solution* of a differential relation $\mathcal{R} \subset X^{(r)}$ is a holonomic section $\sigma : V \to \mathcal{R}$.

Now the reason we introduced this language of jets was because many geometric problems can be phrased in terms of finding genuine solutions to a system of differential relations in the jet space. Finding formal solutions to these equations is typically an easier algebraic-topological obstruction problem, whereas finding genuine solutions directly is much harder. But for certain types of differential relations, *this is enough*, and the existence of a formal solution implies the existence of a genuine solution. More specifically, for such relations, every formal solution is homotopic to a genuine one. In this scenario, we say that the differential relation satisfies the *h-principle*.

3 Motivation

So how do we begin showing that a relation satisfies the *h*-principle? Well, recall that a formal solution is simply a section of the jet space with image contained in the relation, and that such a solution is genuine if it is holonomic. A first step, then, might be to ask the question if an arbitrary section into the jet space can at least be *approximated* by a holonomic one. To explore this question, let's consider the following illustrative example.

Example 3.1. Suppose $A = [0,1] \times \{0\}$. Consider the section

$$\sigma : \mathcal{O}p\, A \to J^1\left(\mathbb{R}^2, \mathbb{R}\right), \quad (x_1, x_2) \mapsto (x_1, x_2, x_1, 0, 0)\,.$$

Can we find a holonomic approximation of this section on A? Well, suppose there was such an approximation $\hat{f}$. Then we would have a holonomic section

$$J^1_{\hat{f}}(x_1, x_2) = \left(x_1, x_2, \hat{f}(x_1, x_2), \frac{\partial \hat{f}}{\partial x_1}(x_1, x_2), \frac{\partial \hat{f}}{\partial x_2}(x_1, x_2)\right)$$

that is C^0-close to σ.

[1]Let $x = \sigma(v)$. Note that we can view the plane above v as $P_x := d_v f(T_v V)$ for some localized section $f : U \to X$ that is 1-tangent to σ at v. Now if the plane was vertical, then for every $w \in P_x$, we would have $d_x p(w) = 0$. But $p(f(v)) = v$, so that $d_x p \circ d_v f = d_v(p \circ f) = d_v \mathrm{id}_U = \mathrm{id}_{T_v V}$ - a contradiction, since we can't simultaneously have $(d_x p \circ d_v f)(T_v V) = d_x p(P_x) = 0$ and $(d_x p \circ d_v f)(T_v V) = \mathrm{id}_{T_v V}(T_v V) = T_v V$. So the planes must be non-vertical. A similar non-verticality argument allows us to identify sections of the fibration $p_0^1 : X^{(1)} \to X$ with connections on X.

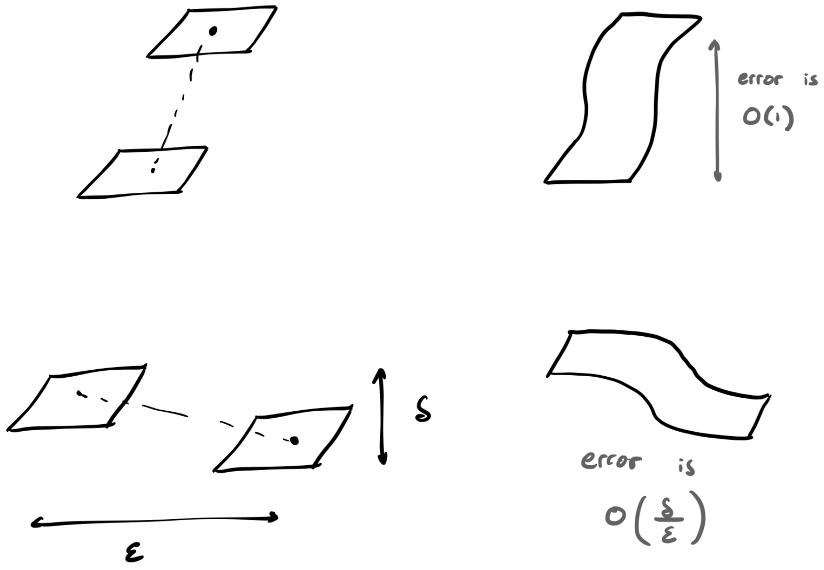

Figure 1: We can get a good approximation if we allow ourselves to move *sideways*. Then the error goes from order $O(1)$ to order $O(\delta/\epsilon)$.

Since we're looking only on A, we consider only points of the form $(x_1, 0)$. Therefore, it suffices to consider the single variable function $f(x_1) := \hat{f}(x_1, 0)$. Now $f'(x_1) = \frac{\partial \hat{f}}{\partial x_1}(x_1, 0)$ being C^0-close to 0 means that $|f(x_1)| < \epsilon$ for some small $\epsilon > 0$ for all $x_1 \in [0, 1]$. Now by the mean value theorem, there is a $c \in (0, 1)$ such that

$$|f(1) - f(0)| = f'(c)(1 - 0) = f'(c) < \epsilon.$$

But if $f(0) \approx \sigma_0(0, 0) = 0$ and $f(1) \approx \sigma_0(1, 0) = 1$ (where σ_0 represents the 0-derivative part of σ), then $|f(1) - f(0)| \approx 1 - 0 = 1$, a contradiction. The change of f must be $O(1)$ yet simulataneously $O(\epsilon)$. So there cannot be a holonomic approximation on this domain.

But what if we allow ourselves to deform our domain slightly via a diffeotopy? Then maybe I can move in the x_2 direction so that the rise is gentler, and less than $O(1)$

Think of how roads are designed on mountains - by doing a lot of switchbacks, we can climb the mountain while staying relatively flat. Holonomic approximation tells us that on such a wiggled domain, we *do* have a holonomic approximation.

Definition 3.2. A closed subset $K \subset M$ is called a *polyhedron* if it is a subcomplex of some smooth triangulation of M.

Remark 3.3. We assume that the manifold V is endowed with a Riemannian metric and the bundle $X^{(r)}$ is endowed with a Euclidean structure in a neighborhood U of the section $\sigma(V) \subset X^{(r)}$, where a Euclidean structure means that we have a smoothly varying inner product defined on the fibers of the bundle in a neighborhood.

Theorem 3.4 (Holonomic approximation). *Let $\sigma : V \to X^{(r)}$ be a section of the r-jet bundle of a fibre bundle $p : X \to V$ and let $K \subset V$ be a polyhedron of positive codimension. Then there exists an isotopy $F_t : V \to V$ and a holonomic section $\hat{\sigma} : \mathcal{O}p\, F_1(K) \to X^{(r)}$ such that the following properties hold:*

- *$\hat{\sigma}$ is C^0-close to σ on $\mathcal{O}p\, F_1(K)$.*
- *F_t is C^0-close to the identity.*

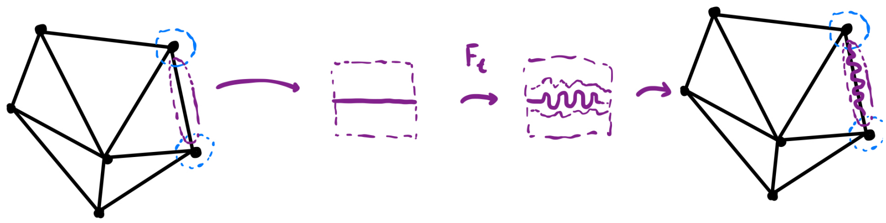

Figure 2: Reducing to the localized version

Furthermore, if σ is holonomic on $\mathcal{O}p\, A$ for some closed subset $A \subset V$, then we may take $F_t \equiv \mathrm{id}_V$

In the following, we'll denote $I = [-1, 1]$.

Theorem 3.5 (Localized holonomic approximation)**.** *Fix $k < m$. Let $\sigma : I^m \to J^r(\mathbb{R}^m, \mathbb{R}^n)$ be a section such that $\sigma = 0$ on $\mathcal{O}p\, \partial I^m$. Then there exists an isotopy $F_t : I^m \to I^m$ and a holonomic section $\hat{\sigma} : I^m \to J^r(\mathbb{R}^m, \mathbb{R}^n)$ such that the following properties hold:*

1. *$\hat{\sigma}$ is C^0-close to σ on $\mathcal{O}p\, F_1(I^k)$.*
2. *F_t is C^0-close to the identity.*
3. *$F_t \equiv \mathrm{id}_{I^m}$ and $\hat{\sigma} = 0$ on $\mathcal{O}p\, \partial I^m$.*

Note that localized holonomic approximation proves holonomic approximation. Indeed, in a local trivialization of a contractible neighborhood, one can view the jet space $X^{(r)}$ as $J^r(\mathbb{R}^m, \mathbb{R}^n)$ if $\dim V = m$ and $\dim X = m + n$.

Working in this Euclidean picture, we can begin with the Taylor approximation on zero simplices. Locally, this is C^0-close, and we can cut it off so it becomes 0 right when it begins to deviate from your desired closeness.

Now suppose we have a one simplex. If we look at a neighborhood of the interior of the simplex, one can trivialize in such a way that locally, the neighborhood looks like I^m and the 1-simplex looks like I in I^m. If we cut off the section near the boundaries, we can use Theorem 2.5 to give us a holonomic approximation along a wiggled neighborhood with respect to that trivialization that respects the section being 0 near the boundaries. Now we can pull this back to a holonomic section in the one simplex viewed in $X^{(r)}$.

Let's consider a 1-simplex and its boundary points. How do we combine our approximations? Well since we constructed our approximations so that they would become 0 near the boundary, on the intersection of the Taylor neighborhoods and the neighborhood of the interior of the one simplex, one can glue the holonomic approximations in a way that doesn't affect C^0-closeness and pull it back to a holonomic section defined on the 1-simplex and its boundary points that's C^0-close to the original section. I'm not too sure about this statement, but I imagine we can glue things together because we can consider a trivialization that maps a holonomic section near the original section to 0, so when they glue along this 0, the pullback to the jet bundle still gives you something holonomic and C^0-close. This needs clarification, of course, but I realized my original argument that I presented in lecture of "adding sections" does not work, since adding holonomic sections in a trivialization does not necessarily give you a holonomic section.

We can continue using a similar argumentative approach to obtain the general statement of holonomic approximation from the localized statement.

Lemma 3.6 (Inductive holonomic approximation)**.** *Fix $j < k < m$. Let $\sigma : I^m \to J^r(\mathbb{R}^m, \mathbb{R}^n)$ be a section such that*

1. *$\sigma = 0$ on $\mathcal{O}p\, \partial I^m$.*

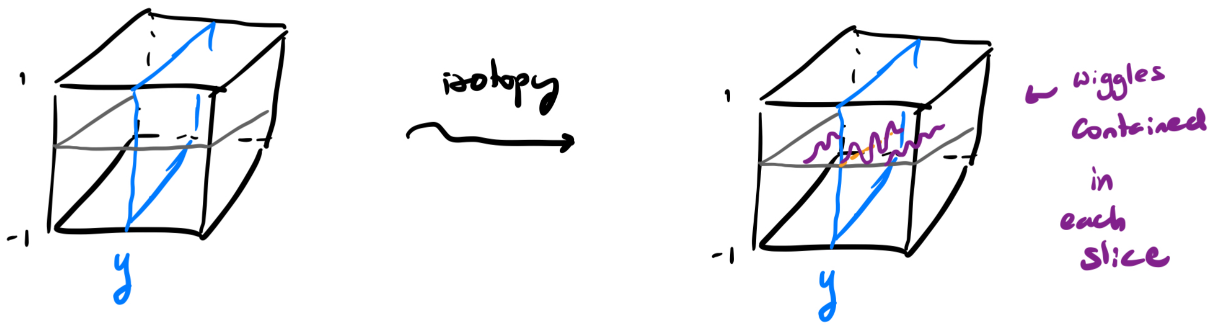

Figure 3: The isotopy acts independently on each slice.

2. *σ is holonomic along the cubes $I^j \times y \times 0, y \in I^{m-j-1}$.*

Then there exists an isotopy $F_t : I^m \to I^m$ and a section $\hat{\sigma} : I^m \to J^r(\mathbb{R}^m, \mathbb{R}^n)$ such that the following properties hold:

1. *$\hat{\sigma}$ is C^0-close to σ on $\mathcal{O}p\, F_1\left(I^k\right)$.*
2. *$\hat{\sigma}$ is holonomic along the cubes $F_1\left(I^{j+1} \times y \times 0\right), y \in I^{m-j-2}$.*
3. *F_t is C^0-close to the identity.*
4. *$F_t \equiv \mathrm{id}_{I^m}$ and $\hat{\sigma} = 0$ on $\mathcal{O}p\, \partial I^m$.*
5. *F_t is fibered over the projection $\left(x_1, \dots, x_m\right) \mapsto \left(x_{j+2}, \dots, x_{m-1}\right)$, hence can be viewed as a family of isotopies $F_t^q : I^{j+1} \times q \times I \to I^{j+1} \times q \times I$, where $q \in I^{m-j-2}$.*

Remark 3.7. Let's think about what this means for $m = 3, k = 2, j = 0$. In this case, we have a section $\sigma : I^3 \to J^r(\mathbb{R}^3, \mathbb{R}^n)$ that vanishes near the boundary of I^3. If $j = 0$, then the second condition tells us that we want σ to be holonomic along the points $\{(y_1, y_2, 0)\}$, where $(y_1, y_2) \in I^2$. In other words, it's holonomic at every point of the middle square. But since sections are *always* holonomic along points (as we can always construct the Taylor polynomial in a neighborhood of a point), this condition is trivially satisfied for any σ for $j = 0$.

The lemma then gives us an isotopy $F_t^1 : I^3 \to I^3$ and a section $\hat{\sigma}^{(1)} : I^3 \to J^r(\mathbb{R}^3, \mathbb{R}^n)$ satisfying properties 1-5. We have, then, that $\hat{\sigma}^{(1)}$ is C^0-close to σ (property 1) on $\mathcal{O}p\, F_1^1(I^2 \times 0)$ and that $\hat{\sigma}^{(1)}$ is holonomic along the wiggled lines (parametrized left to right)[2] $F_1^1(I \times y \times 0)$ where $y \in I$ (property 2). By properties 3 and 4, this deformation is really small, in the sense that it's C^0-close to the identity, and also doesn't affect the boundary at all (it's supported in the interior), and finally, property 5 tells us that the isotopy is fibered over the projection $(x_1, x_2, x_3) \mapsto x_2$, so that the deformation happens independently on the squares going left to right.

Now note $\hat{\sigma}^{(1)}$ is holonomic along each leaf (wiggled line) $L_y := F_1(I \times y \times 0), y \in I$. So we have a neighborhood $\mathcal{O}p\, L_y$ such that the restriction to the leaf of $\hat{\sigma}^{(1)}$ can be extended to a holonomic section $\hat{\sigma}_y^{(1)} : \mathcal{O}p\, L_y \to J^r(\mathbb{R}^3, \mathbb{R}^n)$.

For each $y \in I$, we can intersect $\mathcal{O}p\, L_y$ with the plane $I \times \{y\} \times I$ to get $U_y = L_y \times (-\varepsilon, \varepsilon)_{x_3}$, and define a smooth map $f_y = \mathrm{bs}\, \hat{\sigma}_y^{(1)}|_{U_y} : U_y \to \mathbb{R}^n$ where $J^r\left(f_y\right)|_{U_y} = \hat{\sigma}^{(1)}\Big|_{U_y}$. Since F_t^1 is fibered over these leaves, we can glue these to define a smooth function $\hat{f} : W := \bigcup_{y \in I} U_y \to \mathbb{R}^n : \hat{f}(x_1, x_2, x_3) = f_{x_2}(x_1, x_3)$. This is smooth because the local holonomic extensions $\hat{\sigma}_y^{(1)}$ vary smoothly on the parameter y. Finally, extend this

[2]Left to right means in the x_2 direction. Front and back refer to the x_1 direction, and up and down refer to the x_3 direction.

using a cutoff function to a function $f : I^3 \to \mathbb{R}^n$ on the whole cube .

Now we have enough to continue the induction. Let $\sigma^{(1)} := J^r_{f\circ(F^1_1)^{-1}} : I^3 \to J^r(\mathbb{R}^3, \mathbb{R}^n)$. Note that $\sigma^{(1)}$ is holonomic along $I \times y \times 0$ with $y \in I$ (straight lines going left to right) by construction (it is holonomic everywhere , by definition!), so it satisfies condition 2 of the hypothesis. Note condition 1 is unaffected as near the boundary, we'll have $f_y = 0$. Finally, observe that since $J^r_f = \hat{\sigma}^{(1)}$ on some $\mathcal{O}p\, F^1_1(I^2 \times 0)$ (more specifically, on W), which is C^0-close to σ on some $\mathcal{O}p\, F^1_1(I^2 \times 0)$ with F^1_1 is itself C^0-small, $\sigma^{(1)} := J^r_{f\circ(F^1_1)^{-1}}$ will also be C^0-close to σ on $\mathcal{O}p\, F^1_1(I^2 \times 0)$.

So we can apply the lemma again to get a section $\hat{\sigma}^{(2)} : I^3 \to J^r(\mathbb{R}^3, \mathbb{R}^n)$ and isotopy $F^2_t : I^3 \to I^3$ satisfying properties 1-5. Now $\hat{\sigma}^{(2)}$ is C^0-close to $\sigma^{(1)}$ (property 1) on $\mathcal{O}p\, F^2_1(I^2 \times 0)$, and since $\sigma^{(1)}$ is C^0-close to σ on $\mathcal{O}p\, F^1_1(I^2 \times 0)$ and F^2_1 is C^0-close to the identity, we have that $\hat{\sigma}^{(2)}$ is C^0-close to σ as well on some $\mathcal{O}p\, (I^2 \times 0)$ which is the intersection of these two open neighborhoods. Property 2 then tells us that $\hat{\sigma}^{(2)}$ is holonomic along $F^2_1(I^2 \times 0)$, so that it's now holonomic over a deformed version of the bottom square of the cube. Like before, property 3 and 4 preserve the boundary conditions and make it so that the isotopy can be made small and property 5 becomes irrelevant (it's not fibered over any coordinate).

Since $\hat{\sigma}^{(2)}$ is holonomic along $F^2_1(I^2\times 0)$, there is a holonomic extension $\sigma^{(2)} : \mathcal{O}p\, F^2_1(I^2\times 0) \to J^r(\mathbb{R}^3, \mathbb{R}^n)$ such that $\sigma^{(2)}|_{I^2\times 0} = \hat{\sigma}^{(2)}|_{I^2\times 0}$. As before, take bs $\sigma^{(2)} : \mathcal{O}p\, F^2_1(I^2 \times 0) \to \mathbb{R}^n$ and extend it to a map $g : I^3 \to \mathbb{R}^n$. Consider $\hat{\sigma} = J^r_g : I^3 \to J^r(\mathbb{R}^3, \mathbb{R}^n)$. This is a holonomic section where $F_t := F^2_t \circ F^1_t$ is C^0-close to the identity, $\hat{\sigma}$ is C^0-close to σ on $\mathcal{O}p\, F_1(I^2 \times 0)$ (by making neighborhoods small enough and composing the two isotopies, we can get this contained in the intersection of neighborhoods $\mathcal{O}p\, F^1_1(I^2 \times 0)$ and $\mathcal{O}p\, F^2_1(I^2 \times 0)$), and $F_t \equiv \mathrm{id}_{I^3}$ and $\hat{\sigma} = 0$ on $\mathcal{O}p\, \partial I^3$. Hence, we've proved using the inductive holonomic approximation the localized holonomic approximation theorem for $m = 3, k = 2$.

Clearly, the same argument works for any m and any k. So the inductive holonomic approximation proves the localized holonomic approximation.

Remark 3.8. For the proof of the inductive lemma, note that we will essentially state the proof in [AG24] verbatim, with some mild annotation/elaboration.

Proof. We begin with the case $m = 2, j = 0, k = 1, r = 1, n = 1$. In this case, we have a section $\sigma : I^2 \to J^1\left(\mathbb{R}^2, \mathbb{R}^1\right)$ that is 0 on an open neighborhood of the boundary. Note that since $j = 0$, condition 2 is trivially met like it was in the example above since we are asking σ to be holonomic along points $y \times 0, y \in I$. For each point $x = (x_1, x_2) \in I^2$ we have a 1-jet $\sigma(x) \in J^1\left(\mathbb{R}^2, \mathbb{R}\right)\big|_x$ which consists of a 0 -jet part $h(x) \in \mathbb{R}$ and the two first order formal derivatives $\sigma_1(x), \sigma_2(x) \in \mathbb{R}$, so $\sigma(x) = (x, h(x), \sigma_1(x), \sigma_2(x))$. For fixed $x \in I^2$, let $h_x(y)$ be the unique degree 1 polynomial in the variables $y = \left(y_1, y_2\right)$ whose 2 -jet at x is $\sigma(x)$. We write $h = h\left(\left(u_1, u_2\right), \left(w_1, w_2\right)\right)$ for the function $h_{(u_1,u_2)}\left(w_1, w_2\right)$, (where the first pair is the base point and the second point is the evaluation of the polynomial) so we have

$$\sigma\left(x_1, x_2\right) = \left(h\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right), \frac{\partial h}{\partial w_1}\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right), \frac{\partial h}{\partial w_2}\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right)\right)$$

We will build up our approximation on the wiggled domain by interpolating between these Taylor approximations, which work in neighborhoods of a point.

Given $\epsilon, \delta > 0$ small such that δ/ϵ is also small (the extent will be quantified below), consider the curve

$$w(u) = \frac{\epsilon}{2}\sin\left(\frac{\pi u}{2\delta}\right), \quad u \in \mathbb{R}$$

The idea behind this curve is to make precise the wiggling we talked about earlier in the mountain example. The amplitude $\epsilon/2$ is there to make the total open neighborhood of our wiggled domain (which will be of size

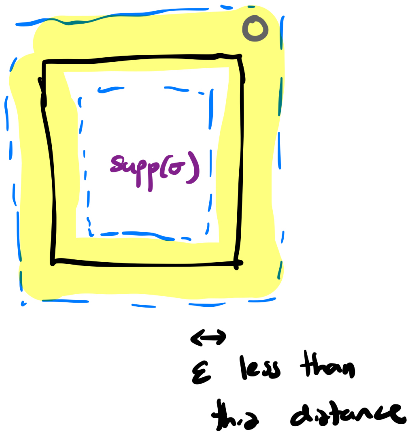

Figure 4: The condition $\operatorname{supp}(\sigma) \subset [-1+\varepsilon, 1-\varepsilon]^m$

$\epsilon/2$) within $[-\epsilon, \epsilon]$, and the factor of $\pi u/2\delta$ is to make the period of the oscillation on the order of δ. The factor of 4 is simply to make things look a bit more aesthetic.

Fix a cutoff function $\psi : \mathbb{R} \to \mathbb{R}$ such that $\psi(u) = 0$ for $|u| < \frac{1}{2}$ and $\psi(u) = 1$ for $|u| > 3/4$. If $\varepsilon > 0$ is small enough so that $\operatorname{supp}(\sigma) \subset [-1+\varepsilon, 1-\varepsilon]^m$, define an isotopy $F_t : I^2 \to I^2$ by the formula

$$F_t(x_1, x_2) = (x_1, x_2 + \varphi_t(x_1, x_2))$$
$$\varphi_t(x_1, x_2) = t\psi\left(\frac{1-|x_1|}{\varepsilon}\right)\psi\left(\frac{1-|x_2|}{\varepsilon}\right)w(x_1).$$

The cutoff functions ensure that the isotopy does not affect the boundary whatsoever, and note that we also have $|\varphi_t(x_1, x_2)|$ bounded by the maximum value of $w(x_1)$, which is $\epsilon/2$. So this wiggling being of order ϵ is necessary to make the isotopy C^0-close to the identity.

Note that

$$dF_t = \begin{pmatrix} 1 & 0 \\ \partial_{x_1}\varphi_t & 1 + \partial_{x_2}\varphi_t \end{pmatrix}$$

Through standard calculus, one gets that the $\partial_{x_1}\varphi_t$ term has maximum magnitude $\pi\epsilon/4\delta$, so is of order $O(\epsilon/\delta)$. Later, we will see that we must have δ/ϵ very small, so this constant is actually quite large, and will dominate the other terms, which are of $O(1)$. Hence, we have the estimate

$$\|dF_t\|_{C^0} \le C\frac{\epsilon}{\delta}$$

Note that though this estimate is quite large, it scales uniformly with ϵ and δ. In other words, the proportion of the wiggling is maintained across different scales, and does not become infinite. It is controlled by the parameters. We will now construct a holonomic approximation $\hat{\sigma} = J^1_g$ of the section σ. The domain of definition of the function g will be the wiggled neighborhood of $I \times 0 \subset I^2$ given $U = F_1\left(\left\{|x_2| < \frac{1}{4}\varepsilon\right\}\right) \subset I^2$. Essentially, we should imagine U as a uniform neighborhood around the oscillation that goes between $(-3\epsilon/4, 3\epsilon/4)$.

To define g we fix a function $\eta : \mathbb{R} \to \mathbb{R}$ such that

- $\eta(u) = -1$ for $u < -1$
- $1 \le \eta(u) \le 1$ for $-1 \le u \le 1$

- $\eta(u) = 1$ for $u > 1$

This function will be a sort of interpolation parameter that'll control how we go between our Taylor approximations at certain points.

We give an explicit formula for g on each of the rectangles

$$R_j = [(2j-1)\delta, (2j+1)\delta] \times [-\epsilon, \epsilon] \subset I^2$$

such that R_j is contained in the support of σ. Now imagine these rectangles as partitioning the support into parts $[-\delta, \delta]$ where the wiggle is increasing/decreasing. Indeed, recall the period of the wiggle was 4δ, so these size 2δ intervals capture half a period.

Suppose first that $j \in \mathbb{Z}$ is even. Then this is when the wiggle is increasing. We set

$$g : R_j \to \mathbb{R}, \quad g(x_1, x_2) = h\left(p(x_1, x_2), x_1, x_2\right)$$

where

$$p(x_1, x_2) = \left((2j\delta) + \delta\eta\left(\frac{4x_2}{\varepsilon}\right), 0\right), \quad (x_1, x_2) \in R_j$$

Let $b(u) = (2j\delta) + \delta\eta(4u/\varepsilon)$, so that $p(x_1, x_2) = (b(x_2), 0)$. Let's dissect this function. Basically, we smoothly interpolate between the Taylor polynomials by moving the base point from left to right. Indeed, note that $\delta\eta(4x_2/\epsilon)$ is δ when $x_2 > \epsilon/4$ and $-\delta$ when $x_2 < -\epsilon/4$.

Using repeated applications of the chain rule, note that

$$\left|b^{(i)}\right| \leq C\frac{\delta}{\epsilon^i}$$

this estimate tells us that we are changing the base point in a way that keeps things close to the original section. In particular, there'll be error terms that involve the derivative of b that can be made arbitrarily small, as long as δ/ϵ^i is small.

For consistency, for $j \in \mathbb{Z}$ odd, we define the function

$$g : R_j \to \mathbb{R}^n, \quad g(x_1, x_2) = h\left(p(x_1, x_2), (x_1, x_2)\right)$$

where

$$p(x_1, x_2) = \left((2j\delta) - \delta\eta\left(\frac{4x_2}{\varepsilon}\right), 0\right), \quad (x_1, x_2) \in R_j$$

The reason for this can be seen with a picture. In essence the right base point for the increasing part becomes the left base point for the decreasing part. To keep the approximation the same, we flip the sign.

Note that we have designed this function to agree specifically on the wiggled neighborhood U. In general, they will not glue on the union of the rectangles. This gives us a globally defined function $g : U \to \mathbb{R}$.

Now let's show that this g gives us a good approximation. Note that the derivatives of b are crucial in showing that the error from moving further and further away from the Taylor approximation at a given point is negligible. Recall that the section $\sigma : I^2 \to J^1\left(\mathbb{R}^2, \mathbb{R}\right)$ has value at a point $q \in I^2$ given by the 2 -jet at q of a

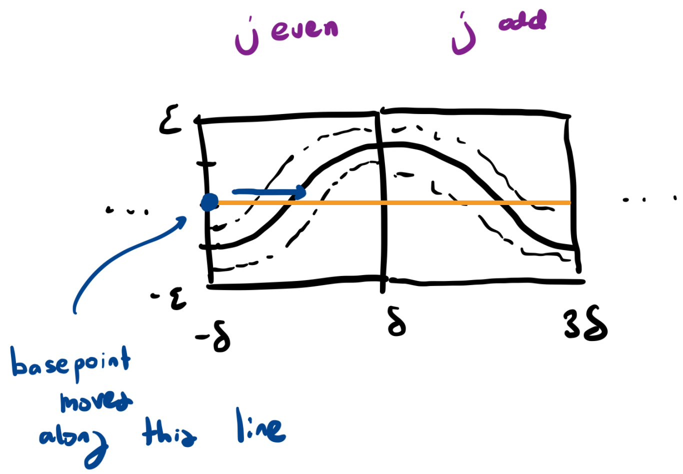

Figure 5: The moving base point.

degree 1 polynomial $h_q : \mathbb{R}^2 \to \mathbb{R}$. We write $h = h\left(\left(u_1, u_2\right), \left(w_1, w_2\right)\right)$ for the function $h_{\left(u_1, u_2\right)}\left(w_1, w_2\right)$, so the coordinates of σ are:

$$\sigma\left(x_1, x_2\right) = \left(h\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right), \frac{\partial h}{\partial w_1}\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right), \frac{\partial h}{\partial w_2}\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right)\right)$$

We now explicitly compute

$$\frac{\partial g}{\partial x_1}\left(x_1, x_2\right) = \frac{\partial h}{\partial w_1}\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right)$$
$$\frac{\partial g}{\partial x_2}\left(x_1, x_2\right) = \frac{\partial h}{\partial w_2}\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right) + b'\left(x_2\right)\frac{\partial h}{\partial u_1}\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right)$$

Observe that the terms $\partial h / \partial w_i\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right)$ are very close to the terms $\partial h / \partial w_i\left(\left(x_1, x_2\right), \left(x_1, x_2\right)\right)$ because $b\left(x_2\right)$ is very close to x_1 (it is δ-close) since $\left|b\left(x_2\right)\right| = \left|\delta\eta\left(\frac{4x_2}{\varepsilon}\right)\right| \leq \delta\left|\eta\left(\frac{4x_2}{\varepsilon}\right)\right| \leq \delta \cdot 1 = \delta$, and x_2 is very small (it is ε-small) (recall $\left|x_2\right| < \epsilon$), hence very close to 0 (ε-close). Hence $\left(b\left(x_2\right), 0\right)$ is very close to (x_1, x_2), and so continuity of $\partial h / \partial w_i$ gives the desired conclusion.

The remaining term $b'\left(x_2\right)\partial h / \partial u_1\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right)$ will be shown to be small as soon as the germ $b'\left(x_2\right)$ is shown to be small, since $\partial h / \partial u_1\left(\left(b\left(x_2\right), 0\right), \left(x_1, x_2\right)\right)$ will be uniformly bounded above by some a priori constant $C > 0$, by continuity of $\partial h / \partial u_1$ (indeed, the arguments where it is evaluated is confined to the compact set $[-\delta, \delta]$, so by the extreme value theorem, it is bounded). But we may ensure that $\left|b'\right|$ is arbitrarily small by choosing $\varepsilon, \delta > 0$ arbitrarily small such that δ / ε is arbitrarily small (for example perform the above construction with $\varepsilon = \frac{1}{N}, \delta = \frac{1}{N^2}$ and N large). Hence the 1-jet $J^1_g\left(x_1, x_2\right)$ is C^0-close to $\sigma\left(x_1, x_2\right)$ on $U \supset F_1(I)$ and the proof is complete for this case. Now let's see how this case can be generalized. If $n > 1$ is arbitrary, then we simply apply the same argument to each of the n-coordinate output functions.

If $r > 1$ is arbitrary, we do the same interpolation, but using the higher order Taylor polynomials. Analogous computations show that we end up with a bound of order $O\left(\delta / \epsilon^r\right)$. In particular, the computation involving the derivatives of b will go up to order r, so the estimate on the r-derivatives of ϵ we had above gives us the bound. So it suffices to choose $\varepsilon, \delta > 0$ arbitrarily small such that δ / ε^r is arbitrarily small (for example perform the above construction with $\varepsilon = \frac{1}{N}, \delta = \frac{1}{N^{r+1}}$ and N large).

For higher dimensional input, let's consider the example $m = 3$. Then we can simply do the construction in the coordinates x_1 and x_3, then parametrize the construction by x_2 to get something over the whole plane.

You should then be able to do a similar "base point" moving construction to get an isotopy $F_t(x_1, x_2, x_3) = (x_1, x_2, x_3 + \varphi_t(x_1, x_2))$ and an open set U. Our moving base point will be of form $b(x_1, x_2)$

Controlling the size of the approximation will involve more derivatives, but in the end, one will see a dominant error term of order δ/ϵ^r at order r. In general, take the coordinates x_1 and x_m and apply the previous construction, parametrizing with respect to the parameters $(x_2, \dots, x_{m-1})$. Then we get a section that is holonomic along the first coordinate x_1.

For the general case, we have that σ is holonomic along the first j coordinates. One apply the same process as the previous case to the coordinates x_{j+1} and x_m to construct a family of functions parametrized by $y \in I^{m-j-2}$. Analogous calculations show that we have the same error term of order δ/ϵ^r. □

References

[AG24] Alvarez-Gavela, D. (2024). *Wiggling and Wrinkling*. Lecture notes for 18.917, taught at MIT spring 2024. Available at: https://drive.google.com/file/d/1S_Rg-zXYmNS6uirPCJcrJxE1Y7jdBsDh/view?usp=sharing

[CEM24] Cieliebak, K., Eliashberg, Y., & Mishachev, N. M. (2024). *Introduction to the h-principle*. Vol. 239, American Mathematical Society.

[AGY24] Alvarez-Gavela, D. (2024). *The holonomic approximation lemma I*. YouTube lecture. Available at: https://www.youtube.com/watch?v=kFgm0EwHvEA