Chapter 1

Intrigue

1.1 The h-principle dichotomy

Gromov observed that it's often fruitful to distinguish two kinds of geometric construction problems. He says that a geometric construction problem satisfies the h-principle if the only obstructions to the existence of a solution come from algebraic topology. In this case, the construction is called flexible, otherwise it is called rigid. This definition is purposely vague. We will see a rather general way to give it a precise meaning, but one must keep in mind that such a precise meaning will fail to encompass a number of situations that can be illuminated by the h-principle dichotomy point of view. The goal of this course is to explain several general techniques proving geometric construction flexibility results.

The easiest example of flexible construction problem which is not totally trivial and is algebraically obstructed is the deformation of immersions of circles into planes. Let f_0 and f_1 be two maps from $\mathbb{S}^1$ to $\mathbb{R}^2$ that are immersions. Since $\mathbb{S}^1$ has dimension one, this mean that both derivatives f_0' and f_1' are nowhere vanishing maps from $\mathbb{S}^1$ to $\mathbb{R}^2$. The geometric object we want to construct is a (smooth) homotopy of immersions from f_0 to f_1, ie a smooth map $F : \mathbb{S}^1 \times [0,1] \to \mathbb{R}^2$ such that $F|_{\mathbb{S}^1\times\{0\}} = f_0$, $F|_{\mathbb{S}^1\times\{1\}} = f_1$, and each $f_p := F|_{\mathbb{S}^1\times\{p\}}$ is an immersion. If such a homotopy exists then, $(t,p) \mapsto f_p'(t)$ is a homotopy from f_0' to f_1' among maps from $\mathbb{S}^1$ to $\mathbb{R}^2 \setminus \{0\}$. Such maps have a well defined winding number $w(f_i') \in \mathbb{Z}$ around the origin, the degree of the normalized map $f_i'/\|f_i'\| : \mathbb{S}^1 \to \mathbb{S}^1$. So $w(f_0') = w(f_1')$ is a necessary condition for the existence of F, which comes from algebraic topology. The Whitney-Graustein theorem states that this necessary condition is also sufficient. Hence this geometric construction problem is flexible. One can give a direct proof of this result, but of course it will also follow from very general results proved in this course.

An important lesson from the above example is that algebraic topology can give us more than a necessary condition. Indeed the (one-dimensional) Hopf degree theorem ensures that, provided $w(f_0') = w(f_1')$, there exists a homotopy g_p of nowhere vanishing maps relating f_0' and f_1'. We also know from the topology of $\mathbb{R}^2$ that f_0 and f_1 are homotopic, say using the straight-line homotopy $p \mapsto f_p = (1-p)f_0 + pf_0$. But of course there is no a priori relation between g_p and the derivative of f_p for $p \notin \{0,1\}$. So we can restate the crucial part of the Whitney-Graustein theorem as: there is a homotopy of immersion from f_0 to f_1 as soon as there is (a homotopy from f_0 to f_1) and a homotopy from f_0' to f_1' among nowhere vanishing maps. The parenthesis in the previous sentence indicated that this condition is always satisfied, but it is important to keep in mind for generalizations. Gromov says that such a homotopy of uncoupled pairs (f, g) is a formal solution of the original problem.

One can generalize this discussion of uncoupled maps replacing a map and its derivative. This is pretty easy for maps from a manifold M to a manifold N. The so called 1-jet space $J^1(M, N)$ is the space of triples (m, n, φ) with $m \in M$, $n \in N$, and $\varphi \in \mathrm{Hom}(T_m M, T_n N)$, the space of linear maps from $T_m M$ to $T_n N$. One can define a smooth manifold structure on $J^1(M, N)$, of dimension $\dim(M) + \dim(N) + \dim(M)\dim(N)$ which fibers over M, N and their product $J^0(M, N) := M \times N$. Beware that the notation (m, n, φ) does not mean that $J^1(M, N)$ is a product of three manifolds, the space where φ lives depends on m and n. Any smooth map $f : M \to N$ gives rise to a section $j^1 f$ of $J^1(M, N) \to M$ defined by $j^1 f(m) = (m, f(m), T_m f)$. Such a section is called a *holonomic section* of $J^1(M, N)$. In the Whitney Graustein example, we use the canonical trivialization of $T\mathbb{S}^1$ and $T\mathbb{R}^2$ to represent $j^1 f$ has a pair of maps (f, f'). The role played by (f, g) in this example is played in general by sections of $J^1(M, N) \to M$ which are not necessarily holonomic.

One can generalize this discussion to $J^r(M, N)$ which remembers derivatives of maps up to order r for some give $r \geq 0$. One can also consider section of an arbitrary bundle $E \to M$ instead of functions from M to N, which are sections of the trivial bundle $M \times N \to N$. The jet-space in this general bundle case is denoted by $E^{(r)}$. The case of $J^1(M, N)$ will be sufficient for most of this course, but we still use the more general case in the following definition:

Definition 1.1. *A partial differential relation $\mathcal{R}$ of order r for sections of a bundle $E \to M$ is a subset of $E^{(r)}$. A solution of $\mathcal{R}$ is a section f of E such that $j^r f(m)$ is in $\mathcal{R}$ for all m. A* formal solution *of $\mathcal{R}$ is a non-necessarily holonomic section of $E^{(r)} \to M$ which takes value in $\mathcal{R}$. The* base section $\mathrm{bs}\, F$ *of a formal solution F is the section $\pi^{(r)} \circ F$ of E, where $\pi^{(r)}$ is the projection $E^{(r)} \to E$. The space of formal solutions is denoted by* $\mathrm{Sec}(\mathcal{R})$*, and the subspace of holonomic ones is denoted by* $\mathrm{Hol}(\mathcal{R})$.

The partial differential relation $\mathcal{R}$ satisfies the h-principle if any formal solution σ of $\mathcal{R}$ is homotopic, among formal solutions, to some holonomic one $j^r f$.

For instance, an immersion of M into N is a solution of $\mathcal{R} = \{(m, n, \varphi) \in J^1(M, N) \mid \varphi \text{ is injective}\}$. As we saw with the Whitney-Graustein problem, we are not only interested to individual solutions, but also in families of solutions. Remember that, in differential topology, a smooth family of maps between manifolds X and Y is a smooth map $h : P \times X \to Y$ seen as the collection of maps $h_p : x \mapsto h(p, x)$. Here P stands for "parameter space". A smooth family of sections of $E \to X$ is a smooth family of maps $\sigma : P \times X \to E$ such that each σ_p is a section.

Note that, for every $(m, n, \varphi) \in J^1(M, N)$, there is some $f : \mathrm{Op}(m) \to N$ such that $f(m) = n$ and $T_m f = \varphi$. Here the used Gromov Op notation: $\mathrm{Op}(A)$ means "an unspecified open neighborhood of the subset A, which may change from line to line". We will also need to know a version of this observation for families of sections. More precisely, what we need to know about the smooth structure on $J^1(M,N)$ is first that a smooth $f : M \to N$ gives rise a smooth $j^1 f$. But we also need to know that, for every family of sections $\sigma : P \times M \to J^1(M,N)$, there is, at least locally on $P \times M$, a family of maps $f : (P \times M) \times M \to N$ such that $\sigma_p(m) = j^1 f_{p,m}(m)$. Of course σ is a family of holonomic section if and only if one can choose $f_{p,m}$ independent of m.

When the parameter space P has boundary, we will typically assume that formal solutions σ_p are holonomic for p in $\mathrm{Op}(\partial P)$.

Definition 1.2. *A partial differential relation $\mathcal{R} \subset E^{(r)}$ satisfies the parametric h-principle if every family of formal solutions $\sigma : M \times P \to E^{(r)}$ which are holonomic for p near ∂P is homotopic, relative to* $\mathrm{Op}(\partial P)$*, to a family of holonomic sections.*

There are other variations on this definition. For instance a formal solution could be holonomic on $\mathrm{Op}(A)$ for some subset A of M, and we say that $\mathcal{R}$ satisfies the relative h-principle if σ can be deformed to a holonomic solution without changing it on $\mathrm{Op}(A)$. Note a common instance of the $\mathrm{Op}(A)$ notation trick: the second neighbordhood in the previous sentence is typically smaller than the first one.

One can also insist on the deformed solution to be C^0-close to the original one. In this case one talks about a C^0-dense h-principle. We are now ready to state a first spectacular h-principle result.

Theorem 1.3. *The relation of immersions in positive codimension (ie immersions of M into N with* $\dim(N) > \dim(M)$*) satisfies all forms of h-principles.*

Of course this theorem covers the Whitney-Graustein theorem (in its second form, assuming the existence of a homotopy between derivatives). But there are much less intuitive applications. The most famous one is the existence of sphere eversions: one can "turn $\mathbb{S}^2$ inside-out among immersions of $\mathbb{S}^2$ into $\mathbb{R}^3$).

Corollary 1.4 (Smale 1958)**.** *There is a homotopy of immersion of $\mathbb{S}^2$ into $\mathbb{R}^3$ from the inclusion map to the antipodal map $a : q \mapsto -q$.*

The reason why this is turning the sphere inside-out is that a extends as a map from $\mathbb{R}^3 \setminus \{0\} \to \mathbb{R}^3 \setminus \{0\}$ by

$$\hat{a} : q \mapsto -\frac{1}{\|q\|^2} q$$

which exchanges the interior and exterior of $\mathbb{S}^2$. More abstractly, one can say the normal bundle of $\mathbb{S}^2$ is trivial, hence one can extend a to a tubular neighborhood of $\mathbb{S}^2$ as an orientation preserving map. Since a is orientation reversing, any such extension will be reversing coorientation.

Proof of the sphere eversion corollary. We denote by ι the inclusion of $\mathbb{S}^2$ into $\mathbb{R}^3$. We set $j_t = (1-t)\iota + ta$. This is a homotopy from ι to a (but not an immersion for $t = 1/2$). Using the canonical trivialization of the tangent bundle of $\mathbb{R}^3$, we can set, for $(q,v) \in T\mathbb{S}^2$, $G_t(q,v) = \mathrm{Rot}^{\pi t}_{Oq}(v)$, the rotation around axis Oq with angle πt. The family $\sigma : t \mapsto (j_t, G_t)$ is a homotopy of formal immersions relating $j^1\iota$ to $j^1 a$. The above theorem ensures this family is homotopic, relative to $t = 0$ and $t = 1$, to a family of holonomic formal immersions, ie a family $t \mapsto j^1 f_t$ with $f_0 = \iota$, $f_1 = a$, and each f_t is an immersion. Stricto sensu, our family σ is not relative to $\mathrm{Op}(\partial[0,1])$ but only to $\partial[0,1]$, but this easily fixed by using instead $t \mapsto \sigma_{\rho(t)}$ for some cut-off function $\rho : [0,1] \to [0,1]$ which vanished near 0 and equals 1 near 1. $\square$

1.2 About the parametric h-principle

1.2.1 Homotopy groups and homotopy equivalences

The parametric h-principle for $\mathcal{R} \subset E^{(r)}$, as defined in the previous section imply in particular that, in any dimension k, every family of sections $\sigma : M \times \mathbb{B}^k \to E$ which are holonomic for parameters near $\partial\mathbb{B}^k$ is homotopic, relative to $\partial\mathbb{B}^k$, to a family whose members are all holonomic. This is (almost) equivalent to saying that the k-th relative homotopy group $\pi_k(\mathrm{Sec}(\mathcal{R}), \mathrm{Hol}(\mathcal{R}))$ vanishes. Geometrically, it means that every ball of formal solution whose boundary lies in the subspace of actual solutions can be pushed into the subspace of actual solutions without moving its boundary. The only reason why the above reformulation is not completely accurate is that we only consider smooth families, whereas homotopy groups are defined in terms of continuous maps. If $\mathcal{R}$ is an open subset of $E^{(r)}$ then this difference is not important, but in general it could matter.

Thanks to the long exact sequence for homotopy groups of pairs, vanishing of all $\pi_k(\mathrm{Sec}(\mathcal{R}), \mathrm{Hol}(\mathcal{R}))$ is equivalent to the fact that the inclusion of $\mathrm{Hol}(\mathcal{R})$ into $\mathrm{Sec}(\mathcal{R})$ induces isomorphisms $\pi_k(\mathrm{Hol}(\mathcal{R})) \to \pi_k(\mathrm{Sec}(\mathcal{R}))$ for all k: one says that this inclusion is a weak homotopy equivalence. In this discussion, we discared base points, but the h-principle without parameter says that each connected component of $\mathrm{Sec}(\mathcal{R})$ contains a point of $\mathrm{Hol}(\mathcal{R})$, and the parametric version then gives homotopy equivalences for all connected components.

In the other direction, if we assume the parametric h-principle only when the parameter space is a ball (of any dimension), then we can deduce it for every parameter space (which is a manifold with boundary). This holds because manifolds can be orderly covered by balls, say using a triangulation or a handle decomposition.

The vocabulary of homotopy theory will not be used in the remainder of those notes. The geometric version with a manifold of parameters is what will really be used. But we will indeed use the preceding paragraph whenever assuming that the parameter space is a ball can make things easier.

1.2.2 Parametricity for free

In many cases, relative parametric h-principles can be deduced from relative non-parametric ones with a larger source manifold. We now explain this strategy for the case of relations in $J^1(X, Y)$.

Let X, P and Y be manifolds, with P seen a parameter space. Denote by Ψ the map from $J^1(X \times P, Y)$ to $J^1(X, Y)$ sending (x, p, y, ψ) to $(x, y, \psi \circ \iota_{x,p})$ where $\iota_{x,p} : T_xX \to T_xX \times T_pP$ sends v to $(v, 0)$.

To any family of sections $F_p : x \mapsto (f_p(x), \varphi_{p,x})$ of $J^1(X, Y)$, we associate the section $\bar{F}$ of $J^1(X \times P, Y)$ sending (x, p) to $\bar{F}(x, p) := (f_p(x), \varphi_{p,x} \oplus \partial f/\partial p(x, p))$.

Lemma 1.5. *In the above setup, we have:*

- *$\bar{F}$ is holonomic at (x, p) if and only if F_p is holonomic at x.*

- *F is a family of formal solutions of some $\mathcal{R} \subset J^1(X, Y)$ if and only if $\bar{F}$ is a formal solution of $\mathcal{R}^P := \Psi^{-1}(\mathcal{R})$.*

In the above setup, let $\mathcal{R}$ be a subset of $J^1(X, Y)$, P' a subset of P, and X' a subset of X such that F_p is holonomic when p belongs to a neighborhood of P', and every F_p is holonomic on a neighborhood of X' (note that both P' and X' could be empty). The goal is to construct a homotopy of such families F^t, starting at $F^0 = F$, such that all F^1_p are everywhere holonomic, and $F^t_p(x) = F_p(x)$ whenever x is near X' or p is near P'. The above lemma ensures this is possible as soon as we have the relative h-principle for $\mathcal{R}^P$. Hence the non-parametric relative h-principle for $\mathcal{R}^P$ implies the parametric relative h-principle for $\mathcal{R}$.

Chapter 2

Convex integration

2.1 Introduction

The goal of this chapter is to explain (Theillière's implementation of) convex integration, the next chapters will give applications. In this chapter $\mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$ and maps defined on $\mathbb{S}^1$ will also freely be seen as 1-periodic maps defined on $\mathbb{R}$.

The elementary step of convex integration modifies the derivative of a map in one direction. Let $f : \mathbb{R}^m \to \mathbb{R}^n$ be a smooth map with compact support. Say we wish $\partial_j f(x)$ could live in some open subset $\Omega_x \subset \mathbb{R}^n$. Assume there is a smooth compactly supported family of loops $\gamma\colon \mathbb{R}^m \times \mathbb{S}^1 \to \mathbb{R}^n$ such that each γ_x takes values in Ω_x, and has average value $\int_{\mathbb{S}^1} \gamma_x = \partial_j f(x)$. Obviously such loops can exist only if $\partial_j f(x)$ is in the convex hull of Ω_x, and we will see this is almost sufficient (hence the word convex in the title of this chapter). For some large positive integer N, we set

$$f'(x) = f(x) + \frac{1}{N} \int_0^{Nx_j} \big[\gamma_x(s) - \partial_j f(x)\big]\, ds.$$

The next proposition implies that, provided N is large enough, we have achieved $\partial_j f'(x) \in \Omega_x$, almost without modifying derivatives in the other directions, and almost without moving $f(x)$. In addition, if we assume that γ_x is constant (necessarily with value $\partial_j f(x)$) for x in some closed subset K where $\partial_j f$ was already good, then the modification is relative to K.

Lemma 2.1 (Theillière 2018)**.** *The new function f' satisfies, uniformly in x:*

1. $\partial_j f'(x) = \gamma(x, Nx_j) + O\left(\frac{1}{N}\right)$
2. $\partial_i f'(x) = \partial_i f'(x) + O\left(\frac{1}{N}\right)$ *for all* $i \neq j$
3. $f'(x) = f(x) + O\left(\frac{1}{N}\right)$
4. $f'(x) = f(x)$ *whenever* γ_x *is constant.*

Proof. We set $\Gamma_x(t) = \int_0^t \big(\gamma_x(s) - \partial_j f(x)\big)\, ds$, so that $f'(x) = f(x) + \Gamma_x(Nx_j)/N$. Because each Γ_x is 1-periodic, and everything has compact support in $\mathbb{R}^m$, all derivatives of Γ are uniformly bounded. Item 3 in the statement is then obvious.

Item 2 also follows since $\partial_i f'(x) = \partial_i f(x) + \partial_i \Gamma(x, Nx_j)/N$. In order to prove Item 1, we compute:

$$\begin{aligned} \partial_j f'(x) &= \partial_j f(x) + \frac{1}{N}\partial_j \Gamma(x, Nx_j) + \frac{N}{N}\partial_t \Gamma(x, Nx_j) \\ &= \partial_j f(x) + O\left(\frac{1}{N}\right) + \gamma(x, Nx_j) - \partial_j f(x) \\ &= \gamma(x, Nx_j) + O\left(\frac{1}{N}\right). \end{aligned}$$

Item 4 is obvious since Γ_x vanishes identically when γ_x is constant. □

Of course the above proposition can be usefull only if we have an efficient way to build families of loops with prescribed average values. The key observation, which will be subsumed by Proposition 2.5, is:

Lemma 2.2 (Gromov 1969)**.** *Let v be a vector in some finite dimensional $\mathbb{R}$-vector space V. Let Ω be a connected open subset of V. If v is in $\operatorname{IntConv}\Omega$, the interior of the convex hull of Ω, then there exists a loop $\gamma : \mathbb{S}^1 \to \Omega$ whose average is v.*

Sketch of proof. By assumption, there is a finite collection of points p_i in Ω and $\lambda_i \in [0, 1]$ such that v is the barycenter $\sum \lambda_i p_i$. Since Ω is open and connected, there is a smooth loop γ_0 which goes through each p_i. The claim is that v is the average value of $\gamma = \gamma_0 \circ h$ for some self-diffeomorphism h of $\mathbb{S}^1$. The idea is of course to choose h such that γ rushes to p_1, stays there during a time roughly λ_1, rushes to p_2, etc. But, in order to achieve average exactly v, it seems like h needs to be a discontinuous piecewise constant map. The assumption that v is in the *interior* of the convex hull gives enough slack to get away with a smooth h, see the proof of Proposition 2.5. Actually the conclusion would be false without this interior assumption (exercise). □

In the previous proof sketch, there is a lot of freedom in constructing γ, which is problematic when trying to do it in families, so we will need a more specific construction. However those two lemmas still contain the essential ideas of convex integration. We will now state precisely the geometric context we will use, and the corresponding fundamental result.

In this chapter, X and Y are smooth manifolds and $\mathcal{R}$ is a first order differential relation on maps from X to Y: $\mathcal{R} \subset J^1(X, Y)$. For any $\sigma = (x, y, \varphi)$ in $\mathcal{R}$ and $(\lambda, v) \in T_x^*V \times T_xV$ such that $\lambda(v) = 1$ (like (dx_j, ∂_j) in Lemma 2.1), we set:

$$\mathcal{R}_{\sigma,\lambda,v} = \operatorname{Conn}_{\varphi(v)} \left\{ w \in T_yY \;;\; (x,\ y,\ \varphi + (w - \varphi(v)) \otimes \lambda) \in \mathcal{R} \right\}$$

where $\operatorname{Conn}_a A$ is the connected component of A containing a. In order to decipher this definition, it suffices to notice that $\varphi + (w - \varphi(v)) \otimes \lambda$ is the unique linear map from T_xX to T_yY which coincides with φ on $\ker \lambda$ and sends v to w. In particular, $w = \varphi(v)$ gives back φ.

Of course we will want to deal with more that one point, so we will consider a vector field V and a 1–form λ such that $\lambda(V) = 1$ on some subset U of X, a formal solution F (defined at least on U), and get the corresponding $\mathcal{R}_{F,\lambda,v}$ over U.

Definition 2.3. *A formal solution F of $\mathcal{R}$ over U is (λ, V)–short if $T_x \operatorname{bs} F(V_x) \in \operatorname{IntConv} \mathcal{R}_{F(x),\lambda_x,V_x}$ for all x. It is short if it is (λ, V)–short for all (λ, V).*

One easily checks that $\mathcal{R}_{\sigma,\kappa^{-1}\lambda,\kappa v} = \kappa\mathcal{R}_{\sigma,\lambda,v}$ hence the above definition only depends on $\ker \lambda$ and the direction $\mathbb{R}V$.

The context of the elementary step of convex integration is given by two manifolds X and Y, a function $\pi : U \to \mathbb{R}$ on some open subset $U \subset X$, a vector field V such that $d\pi(V) = 1$, and a sub-bundle $H \subset \ker d\pi$ (which could be $\{0\}$).

Lemma 2.4 (Fundamental lemma of convex integration)**.** *In the above context, let $\mathcal{R} \subset J^1(X, Y)$ be an open differential relation. Let F be a formal solution of $\mathcal{R}$ on X which is $(d\pi, V)$–short and H-holonomic on U, and V–holonomic near some closed subset C. For any compact $K \subset U$ and any positive ε, there is a homotopy of global formal solutions F^t, starting at $F^0 = F$, such that:*

- *$F^t(x) = F(x)$ for all t whenever x is outside U or near C*
- *$\mathrm{d}_{C^0}(\operatorname{bs} F^t, \operatorname{bs} F) \leq \varepsilon$ for all t,*
- *F^1 is $H \oplus \mathbb{R}V$–holonomic near K*

Everything until the end of this chapter is dedicated to the proof of the above lemma, and historical remarks about it. Applications in the next chapters will only use this statement.

2.2 Constructing loops

In this section, we explain how to construct families of loops to feed into the corrugation process of Proposition 2.1. More precisely, all this section is devoted to proving the following proposition.

Proposition 2.5. *Let $g : \operatorname{Op} I^m \to \mathbb{R}^n$ be a smooth map, and let Ω be an open set in $\mathbb{R}^m \times \mathbb{R}^n$. For each x in I^m, assume that $\Omega_x := \Omega \cap (\{x\} \times \mathbb{R}^n)$ is connected. If, for each x in I^m, $g(x)$ is in the convex hull of Ω_x, then there exists a smooth family of loops $\gamma : \operatorname{Op} I^m \times \mathbb{S}^1 \to \mathbb{R}^n$ such that $\gamma_x(t) \in \Omega_x$ for all t, and $\bar{\gamma}_x = g(x)$.*

Let K be a closed subset of I^m. If $g(x)$ is in Ω_x for x in $\operatorname{Op} K$, and $g|_K$ extends to a section of $\Omega \to I^m$, then one can additionally ensure that, for all $x \in \operatorname{Op} K$, γ_x is the constant loop $t \mapsto g(x)$.

Note that the extension condition is the second part is not automatic. For instance, one could have $n = m$, $\Omega_x = \{1/2 < \|y\|_\infty < 2\}$, $K = \partial I^m$, $g(x) = x$.

2.2.1 Surrounding families

It will be convenient to introduce some more vocabulary. We say a loop γ surrounds (resp. strictly surrounds) a vector v if v is in the convex hull of $\gamma(\mathbb{S}^1)$ (resp. the interior of this convex hull). Also, we fix a base point in $\mathbb{S}^1$ and say a loop is based at some y if the base point is sent to y.

The first main task in proving Proposition 2.5 is to construct suitable families of loops γ_x surrounding $g(x)$, by assembling local families of loops. The key here is to look for loops that are concatenations of round-trips. A round-trip is a loop γ such that $\gamma(1-t) = \gamma(t)$ for all t, so that γ follows some path until $t = 1/2$ and then backtracks. We will see this conditions ensure a canonical contraction to a constant loop. The goal of this subsection it to prove the following lemma.

Lemma 2.6. *Under the assumptions of Proposition 2.5, let β be a section of Ω over I^m extending $g|_K$. For any smooth family of concatenations of round-trip loops γ_x based at $\beta(x)$, defined for x in $\operatorname{Op} K$, and strictly surrounding $g(x)$, there exists such a family γ'_x defined for all x, and satisfying $\gamma'_x = \gamma_x$ for x in $\operatorname{Op} K$.*

The crucial property of surrounding round-trips is that one can smoothly interpolate between them in families.

Lemma 2.7. *Let U be an open subset of $\mathbb{R}^m$, Ω an open subset in $U \times \mathbb{R}^n$, $\beta : U \to \Omega$ a section, and $g : U \to \mathbb{R}^n$. Let $\gamma^0, \gamma^1 : U \times \mathbb{S}^1 \to \mathbb{R}^n$ be two smooth families of concatenation of round-trips where each γ^i_x is based at $\beta(x)$ and strictly surrounds $g(x)$. There is a smooth map $\gamma : [0,1] \times U \times \mathbb{S}^1 \to \mathbb{R}^n$ such that $\gamma^t := \gamma(t, \cdot, \cdot)$ interpolates between γ^0 and γ^1, and each γ^t_x is a loop based at $\beta(x)$, with values in Ω_x, and strictly surrounding $g(x)$.*

Proof. The key is that any round-trip γ can be canonically contracted to a constant loop through the family $H_u\gamma$, $u \in [0,1]$, whose image starts backtracking at $s = 1/2 - u/2$ instead of $s = 1/2$. In formula:

$$H_u\gamma(s) = \begin{cases} \gamma(s) & \text{if } s \leq \frac{1}{2} - \frac{u}{2} \text{ or } s \geq \frac{1}{2} + \frac{u}{2} \\ \gamma(\frac{1}{2} - \frac{u}{2}) & \text{if } s \in [\frac{1}{2} - \frac{u}{2}, \frac{1}{2} + \frac{u}{2}] \end{cases}$$

The round-trip condition ensures each $H_u\gamma$ is a continuous loop (only continuity at $s = 1/2 + u/2$ needs some attention). This is not a family of smooth loops but this can be corrected by reparametrization (without leaving the class of round-trips). Note that reparametrization does not change the convex hull of a loop (it does change its average though). By construction, $\operatorname{im} H_u\gamma \subset \operatorname{im}\gamma$ for all u, $H_0\gamma = \gamma$ and $H_1\gamma$ is constant. The same procedure can be used for concatenation of round-trips, contracting all pieces at the same time. The result is still a concatenation of round-trips for each u.

We now turn to our two families γ^0 and γ^1. After precomposition by a fixed smooth map, we can assume all loops with a given base point can be smoothly concatenated. Let $\rho \colon [0,1] \to [0,1]$ be a smooth cut-off function which vanishes on $[0, 1/2]$ and equals one near $t = 1$. The symmetric function $t \mapsto \rho(1-t)$ vanishes near $t = 0$ and equals one on $[1/2, 1]$. We set

$$\delta^t = \left(H_{\rho(t)}\gamma^0\right) \# \left(H_{\rho(1-t)}\gamma^1\right).$$

Note that the convex hull of δ^t contains the convex hull of γ^0 (resp. γ^1) when t is less (resp. more) than $1/2$. Hence each δ^t_x strictly surrounds $g(x)$. In addition, for $i = 0, 1$, δ^i is γ^i up to reparametrization. We don't get exactly γ^i because we needed to reparametrize in order to smoothly concatenate, and because δ^i spends half its time at the base point. Nothing about this depends on x, and everything can be fixed by some initial and final deformations of γ^0 and γ^1 respectively, in order to get our final interpolation γ. □

Corollary 2.8. *Under the assumption of the preceding lemma, if γ_x^0 and γ_x^1 are defined for x in neighborhoods of compact subsets K_0 and K_1 respectively, based at β and strictly surrounding g, there is a family γ of concatenation of round-trips defined on* $\mathrm{Op}\,(K_0 \cup K_1)$*, based at β and strictly surrounding g, such that $\gamma = \gamma^0$ on* $\mathrm{Op}\, K_0$.

Proof. For $i = 0, 1$, let U_i be an open neighborhood of K_i where γ^i is defined. Let U_0' be another open neighborhood of K_0 whose closure $\bar{U}_0'$ is compact in U_0. Since $\bar{U}_0'$ and $K_1' := K_1 \setminus (K_1 \cap U_0)$ are disjoint compact subsets of $\mathbb{R}^m$, there is some smooth cut-off $\rho \colon \mathbb{R}^m \to [0, 1]$ which vanishes on U_0' and equals one on some neighborhood U_1' of K_1'.

Lemma 2.7 gives a homotopy of loops γ^t from γ^0 to γ^1 on $U_0 \cap U_1$. On $U_0' \cup (U_0 \cap U_1) \cup U_1'$, which is a neighborhood of $K_0 \cup K_1$, we set

$$\gamma_x = \begin{cases} \gamma_x^0 & \text{for } x \in U_0' \\ \gamma_x^{\rho(x)} & \text{for } x \in U_0 \cap U_1 \\ \gamma_x^1 & \text{for } x \in U_1' \end{cases}$$

which has the required properties. $\square$

Proof of Lemma 2.6. We begin with two general observations, and then prove the lemma. First observe that for each x there is a round-trip γ based at $\beta(x)$ and strictly surrounding $g(x)$. Indeed, by assumption, $g(x)$ is a strict convex combination of points p_i all in the connected component of Ω_x containing $\beta(x)$. Then we chose any path $\lambda \colon [0, 1] \to \Omega_x$ going through these points and form the corresponding round-trip by smooth reparametrization of the concatenation $\lambda \# \lambda^{-1}$, say $\gamma(s) = \lambda((1 - \cos 2\pi s)/2)$.

Next observation is that $s \mapsto \beta(x') + (\gamma(s) - \beta(x))$ is a round-trip based at $\beta(x')$ which strictly surrounds $g(x')$ for x' in $\mathrm{Op}\,\{x\}$ because Ω is open and $g(x)$ is in the interior of the convex hull of $\mathrm{im}\,\gamma$.

In the setup of the lemma, let U_0 be an open set containing K and contained in the domain of γ. Let U_0' be an open neighborhood of K with compact closure in U_0. Let U_i, $1 \le i \le N$ be a covering of $I^m \setminus U_0'$ by open subsets not intersecting K and where the preceding observations gives families of loops γ^i. We also set $\gamma^0 = \gamma|_{U_0}$. In particular the open sets U_i, $i \ge 0$ cover the whole of I^m, and only U_0 intersects K. Let K_i, $0 \le i \le N$, be a family of compact sets with $K_i \subset U_i$ which covers I^m. We repeatedly apply Corollary 2.8 to K_i and K_{i+1}, in this order, to get a family γ' defined over all I^m. Since each step preserves the family on $\mathrm{Op}\, K_i$ and only U_0 intersects (in fact contains) K, we do have $\gamma' = \gamma$ on $\mathrm{Op}\, K$. $\square$

2.2.2 The reparametrization lemma

The second ingredient needed to prove Proposition 2.5 is a parametric version of Lemma 1.5.

Lemma 2.9. *Let $\gamma : X \times \mathbb{S}^1 \to E$ be a family of loops in a vector bundle $E \to X$ (ie. each γ_x is a loop in E_x) and let g be a section of E. If γ_x strictly surrounds $g(x)$ for all x, and γ_x has average $g(x)$ when x is outside some compact subset of X, then there is a family of circle diffeomorphisms $h : X \times \mathbb{S}^1 \to \mathbb{S}^1$ such that each $\gamma_x \circ h_x$ has average $g(x)$.*

Proof. For any fixed x, since γ_x strictly surrounds $g(x)$, there are points $s_1, ..., s_{n+1}$ in $\mathbb{S}^1$ such that $g(x)$ is in the interior of the simplex spanned by the corresponding points $\gamma_x(s_j)$. The fact that $n+1$ points are enough is the statement of Caratheodory's theorem in convex geometry. Let $\mu_1, ..., \mu_{n+1}$ be smooth positive probability measures very close to the Dirac measures on s_j (ie. $\mu_j = f_j\, ds$ for some smooth positive function f_j and, for any function h, $\int h\, d\mu_j$ is almost $h(s_j)$). We set $p_j = \int \gamma_x\, d\mu_j$, which is almost $\gamma_x(s_j)$ so that $g(x) = \sum w_j p_j$ for some weights w_j in the open interval $(0,1)$ (this is where we crucially use that $g(x)$ was *strictly* surrounded by γ_x).

If x' is in a sufficiently small neighborhood of x, $g(x')$ is in the interior of the simplex spanned by $p_j(x') := \int \gamma_{x'}\, d\mu_j$. In such a neighborhood U, we then have smooth weight functions w_j such that $g(x') = \sum w_j(x')p_j(x')$. By compactness of I^m, we extract a finite cover by such open sets U^i, $1 \leq i \leq N$, with corresponding measures μ_j^i, moving points p_j^i and weight functions w_j^i. Let (ρ_i) be a partition of unity associated to this covering. For every x, we set

$$\mu_x = \sum_{i=1}^{N} \sum_{j=1}^{n+1} \rho_i(x) w_j^i(x)\, \mu_j^i$$

so that:

$$\begin{aligned} \int \gamma_x\, d\mu_x &= \sum_{i=1}^{N} \rho_i(x) \sum_{j=1}^{n+1} w_j^i(x) \int \gamma_x\, d\mu_j^i \\ &= \sum_{i=1}^{N} \rho_i(x) \sum_{j=1}^{n+1} w_j^i(x) p_j^i(x) \\ &= \sum_{i=1}^{N} \rho_i(x) g(x) = g(x). \end{aligned}$$

We now set $h_x^{-1}(t) = \int_0^t d\mu_x$ so that $g(x) = \overline{\gamma_x \circ h_x}$ for all x. □

2.2.3 Proof of the loop construction proposition

We finally assemble the ingredient from the previous two sections.

Proof of Proposition 2.5. Let β be a section of Ω extending $g|_{\mathrm{Op}\, K}$. If K is not empty, existence of β is part of the assumptions, otherwise use contractibility of I^m and openness of Ω to construct it (note that Ω intersects each fiber because the convex hull of Ω_x contains $g(x)$).

Let γ^* be a round-trip loop strictly surrounding the origin in $\mathbb{R}^n$. For x in some neighborhood U^* of K where $g = \beta$, we set $\gamma_x = g(x) + \varepsilon\gamma^*$ where $\varepsilon > 0$ is sufficiently small to ensure the image of γ_x and its convex hull are contained in Ω_x (recall Ω is open and K is compact). Lemma 2.6 extends this family to a family of strictly surrounding loops γ_x for all x (this is not yet our final γ). Then Lemma 2.9 gives a family of circle diffeomorphisms h_x such that $\gamma_x \circ h_x$ has average $g(x)$.

Finally we choose a cut-off function function χ which vanishes on $\mathrm{Op}\, K$ and equals one on $\mathrm{Op}\, I^m \setminus U^*$. In U^*, we replace $\gamma_x \circ h_x = g(x) + \gamma^* \circ h_x$ by $g(x) + \chi(x)\gamma^* \circ h_x$. This operation does not change the average values of these loops, because it rescales them around their average value, but makes them constant on $\operatorname{Op} K$. Also those loops stay in Ω thanks to our choice of ε. Note that one could easily arrange for h_x to be independent of x on $\operatorname{Op} K$, but this is not necessary in this proof. $\square$

To be continued... The proof of Lemma 2.4 is coming soon (all important pieces are already in the preceding pages).