Formalising the h-principle and sphere eversion

Patrick Massot Floris van Doorn Oliver Nash

Abstract

In differential topology and geometry, the h-principle is a property enjoyed by certain construction problems. Roughly speaking, it states that the only obstructions to the existence of a solution come from algebraic topology.

We describe a formalisation in Lean of the local h-principle for first-order, open, ample partial differential relations. This is a significant result in differential topology, originally proven by Gromov in 1973 as part of his sweeping effort which greatly generalised many previous flexibility results in topology and geometry. In particular it reproves Smale's celebrated sphere eversion theorem, a visually striking and counter-intuitive construction. Our formalisation uses Theillière's implementation of convex integration from 2018.

This paper is the first part of the sphere eversion project, aiming to formalise the *global* version of the h-principle for open and ample first order differential relations, for maps between smooth manifolds. Our current local version for vector spaces is the main ingredient of this proof, and is sufficient to prove the titular corollary of the project. From a broader perspective, the goal of this project is to show that one can formalise advanced mathematics with a strongly geometric flavour and not only algebraically-flavoured mathematics.

1 Introduction

1.1 Formal proofs and geometric intuition

Mathematical arguments that rely on human geometric intuition can be seen as challenges to formalisation. Logic, discrete mathematics and algebra have certainly been formalised much more often. Two notable exceptions are an elementary proof of Jordan's curve theorem in [Hal07] and the Poincaré–Bendixson theorem in [IT20]. The goal of our project is to make a strong case that even differential topology can be formalised, including Smale's sphere eversion theorem.

The context of this theorem is Gromov's h-principle. He says that a geometric construction problem satisfies the h-principle, or is flexible, if the only obstructions to the existence of a solution come from algebraic topology (the letter h stands for "homotopy").

1.2 A toy example

The easiest example of a flexible construction problem which is not trivial and which 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 means 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[1]) homotopy of immersions from f_0 to f_1, i.e. a smooth map $F\colon \mathbb{S}^1 \times [0,1] \to \mathbb{R}^2$ 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 normalised map $f_i'/\|f_i'\|\colon \mathbb{S}^1 \to \mathbb{S}^1$. 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.

The lesson to take away from the example above is that a necessary condition imposed by algebraic topology may turn out to be sufficient. 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, e.g. via the straight-line homotopy $p \mapsto f_p = (1-p)f_0 + pf_1$. But a priori there is no 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 immersions 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 parentheses in the previous sentence indicate that this condition is always satisfied, but it is important to keep in mind for generalisations. Gromov says that such a homotopy of uncoupled pairs (f,g) is a formal solution of the original problem. Note an unfortunate terminology clash here: the word formal isn't used in the same sense as in "formalised mathematics".

[1] In this paper, smooth always means infinitely differentiable.

1.3 The h-principle

We now generalise this discussion of formal solutions to the case of maps between two vector spaces E, F with totally arbitrary pointwise constraints on their derivatives. A first-order partial differential relation for functions $E \to F$ is any subset

$$\mathcal{R} \subseteq E \times F \times L(E, F),$$

where $L(E, F)$ denotes the linear maps from E to F. A function $f : E \to F$ satisfies $\mathcal{R}$ at a point x in E if

$$(x, f(x), Df(x)) \in \mathcal{R},$$

where $Df(x)$ is the derivative of f at x.

We say that a pair $(f, \varphi) : E \to F \times L(E, F)$ is a *formal solution* if for all $x \in E$ we have $(x, f(x), \varphi(x)) \in \mathcal{R}$. Any formal solution that is obtained from a function f (i.e. $\varphi = Df$) is called *holonomic*, showing that f satisfies $\mathcal{R}$ everywhere. The product $E \times F \times L(E, F)$ is called the space of 1-jets of maps from E to F and denoted by $J^1(E, F)$.

We say that $\mathcal{R}$ satisfies the *h-principle* if any formal solution can be deformed into a holonomic one inside $\mathcal{R}$.

Not all relations satisfy the h-principle. For instance the relation of immersions of circles into a line rather than a plane fails to satisfy it. A key insight of Gromov was to identify a *geometric* condition on the relation which ensures that the h-principle holds. In the special case of immersions of circles into planes, at each point x the derivative $f'(x)$ is required to belong to the complement of $\{0\} \subseteq \mathbb{R}^2$ and checking Gromov's condition boils down to the fact that this complement is open and *ample*. A set in a real affine space is called ample if all its connected components have the full space as their convex hull. Note how this condition fails in the case of immersions of circles into lines: in that case the complement of the origin has two connected components, neither of whose convex hulls is the whole line.

Returning to the general case with $\mathcal{R} \subseteq J^1(E, F)$, one says that $\mathcal{R}$ is ample if, for every $(x, y, \varphi) \in J^1(E, F)$ and every hyperplane H in E, the set:

$$\{\psi \in L(E, F) \mid \psi|_H = \varphi|_H \text{ and } (x, y, \psi) \in \mathcal{R}\}$$

is ample in the affine space of linear maps that coincide with φ on H. This set is called the slice of $\mathcal{R}$ associated to (x, y, φ) and H. In fact, it is more convenient to use a variant of this definition which assumes more data than just H, and so obtain slices that are subsets of F rather than $L(E, F)$. Instead of considering hyperplanes in E, we consider what we call *dual pairs*. This is a pair $p = (\pi, v) \in E^* \times E$ such that $\pi(v) = 1$. Using p, we define the update map $\Upsilon_p \colon L(E,F) \to F \to L(E,F)$ sending (φ, w) to the linear map $\Upsilon_p(\varphi, w)$ which coincides with φ on $\ker \pi$ and sends v to w.

The hyperplane associated to p is $\ker \pi$ and the corresponding slice is affine-isomorphic to

$$\mathcal{R}((x,y,\varphi),p) := \{w \in F \mid (x, y, \Upsilon_p(\varphi, w)) \in \mathcal{R}\}.$$

Before stating the main theorem, we note that what we called the h-principle above is only the weakest possible variation on this theme. One can add constraints along various sets, proximity constraints, and parameters. In the statement below, P is the parameter space, it would be one-dimensional in the example of homotopies of immersions. A family of formal solutions of $\mathcal{R}$ parametrised by P is a smooth map $\mathcal{F} : P \times E \to F \times L(E,F)$ such that each $\mathcal{F}_p := \mathcal{F}(p, \cdot)$ is a formal solution.

Theorem 1.1 (Local version of Gromov's h-principle for open and ample first order relations)**.** *Let E, F and P be finite dimensional real normed vector spaces. Let $\mathcal{R}$ be an open and ample set in $J^1(E,F)$. Let C and K be sets in $P \times E$ such that C is closed and K is compact. Let $\mathcal{F}_0$ be a family of formal solutions of $\mathcal{R}$ parametrised by P such that $(\mathcal{F}_0)_p$ is holonomic at x for all (p,x) near C.*

For every positive real number ε, there exists a family of formal solutions $\mathcal{F}$ parametrised by $\mathbb{R} \times P$ such that:

- $\forall p\, x, \mathcal{F}((0,p),x) = \mathcal{F}_0(p,x)$,
- $\forall (p,x)$ *near* $C, \forall t, \mathcal{F}((t,p),x) = \mathcal{F}_0(p,x)$,
- $\forall t\, p\, x, \left\| \mathrm{pr}_F \left(\mathcal{F}((t,p),x) - \mathcal{F}_0(p,x) \right) \right\| \le \varepsilon$,
- $\forall (p,x)$ *near* $K, \mathcal{F}_{(1,p)}$ *is holonomic at* x.

where $\mathrm{pr}_F : F \times L(E,F) \to F$ *is the projection to the first factor.*

This paper describes our formalisation of the above theorem. We also explain how to apply it to obtain Smale's theorem whose formalised statement follows (the notation `I` refers to the unit interval in $\mathbb{R}$ and we omit the declaration of the fact that `E` is a 3-dimensional real vector space equipped with an inner product).

```
theorem sphere_eversion_of_loc :
  ∃ f : ℝ → E → E,
    (𝒞 ∞ ↿f) ∧
    (∀ x ∈ 𝕊², f 0 x = x) ∧
    (∀ x ∈ 𝕊², f 1 x = -x) ∧
    ∀ t ∈ I, sphere_immersion (f t)
```

1.4 Convex integration

The proof of Theorem 1.1 uses convex integration, a technique invented by Gromov c.1970, inspired by the C^1 isometric embedding results of Nash and the original proof of flexibility of immersions. This term is pretty vague however, and there are several different implementations. The most recent, and by far the most efficient, is Mélanie Theillière's corrugation process from [The22]. We opted for this implementation and describe it briefly below.

Let f be a smooth map from $\mathbb{R}^n$ to $\mathbb{R}^m$ and suppose we want to turn f into a immersion. We tackle each partial derivative in turn. We first make sure $\partial_1 f(x) := \partial f(x)/\partial x_1$ is non-zero (for all x). Then we make sure $\partial_2 f(x)$ is not collinear with $\partial_1 f(x)$. Then we make sure $\partial_3 f(x)$ is not in the plane spanned by the two previous derivatives, etc until all n partial derivatives are everywhere linearly independent.

More generally, if we want f to satisfy some relation (not necessarily the relation of immersions) then for each j between 1 and n, we want $\partial_j f(x)$ to live in some open subset $\Omega_x \subseteq \mathbb{R}^m$. The brilliant idea is to assume there is a smooth family of loops $\gamma\colon \mathbb{R}^n \times \mathbb{S}^1 \to \mathbb{R}^m$ such that each γ_x takes values in Ω_x, and has average value $\overline{\gamma}_x = \partial_j f(x)$. Obviously such loops can exist only if $\partial_j f(x)$ is in the convex hull of Ω_x, hence the name convex integration, and we will see this condition is almost sufficient. This is where the ampleness condition enters the story. In the immersion case, this condition holds as soon as $m > n$ because, from the above description, each Ω_x is the complement of a subspace with codimension at least two.

For some large positive N, we replace f by the new map

$$x \mapsto f(x) + \frac{1}{N}\int_0^{Nx_j} [\gamma_x(s) - \overline{\gamma}_x]\, ds.$$

The key is that, provided N is large enough, we have achieved $\partial_j f(x) \in \Omega_x$, almost without modifying derivatives $\partial_i f(x)$ for $i \neq j$, and almost without moving $f(x)$.

This is a very local construction, and it isn't obvious how the absence of homotopical obstruction, embodied by the existence of a formal solution, should enter the discussion. The answer is that it essentially provides a way to coherently choose base points for the γ_x loops.

1.5 Lean and `mathlib`

We formalised all this using the Lean 3 theorem prover [dMKA+15], developed principally by Leonardo de Moura since 2013 at Microsoft Research. Lean implements a version of the calculus of inductive constructions [CP88], with quotient types, non-cumulative universes, and proof irrelevance. Our formalisation is built on top of `mathlib` [Com20], a library of formalised mathematics with more than 250 contributors and almost one million lines of code as of September 2022. In return, about two thirds of the code we wrote for this project has been, or will be, added to `mathlib`. As with many recent projects using `mathlib`, the most crucial property is that library is extremely tightly integrated. We needed to make simultaneous use of its theories of topology, affine geometry, linear algebra, calculus, and integration. We also emphasise that `mathlib` follows the standard mathematical practice where nothing is made more complicated than it needs to be for the sake of avoiding the law of excluded middle or the axiom of choice.

Paper outline: Section 2 discusses topics that are pervasive and not specific to a part of the project. Section 3 provides the supply of loops. Section 4 then discusses the proof of the main theorem, including the key corrugation construction, and Section 5 deduces Smale's sphere eversion theorem.

We will link to specific results in `mathlib` and the sphere eversion project using this icon. We will use static links to the version of these repositories when we wrote this paper.

2 Preliminaries

2.1 Proximity and filters

The kind of geometrical constructions that we consider in this paper frequently require consideration of properties which hold near a point or near a set. Informally, this means that there is an unspecified neighbourhood of that point or set where the property holds, and that this neighbourhood can shrink finitely many times during the construction. Formally, this idea is perfectly captured by the theory of filters. This was already noted in [HIH13] but we will review it here since it is relevant to the future development of differential topology and since it gives us the opportunity to advertise a slightly different point of view.

In [HIH13], filters are mostly seen as providing generalised bounded quantifiers. We will explain this below, but more generally we see filters on a type X as generalised sets of X. For instance, given a point x_0 in a topological space X, we may wish to consider the generalised set of "points that are close to x_0". For the types of natural or real numbers we need the generalised sets of "very large numbers". These generalised sets are called filters on X. There is an injective map from the type $\mathcal{P}(X)$ of ordinary sets whose elements have type X into the type $\mathcal{F}(X)$ of filters on X. One can extend the inclusion order relation to an order relation $\leq$ on $\mathcal{F}(X)$. Given a function $f\colon X \to Y$, we can also extend the associated direct image map from $\mathcal{P}(X) \to \mathcal{P}(Y)$ to $\mathcal{F}(X) \to \mathcal{F}(Y)$ (and extend also the inverse image map). Using this order relation and direct image operation, one can for instance say that a sequence $u\colon \mathbb{N} \to X$ converges to a point x_0 if the direct image under u of the filter of very large natural numbers is contained in the filter of points that are close to x_0.

For any predicate P and any filter F on X, one can form the statement $\forall^f x \in F, P(x)$ claiming that "$P(x)$ holds for all x in F", generalising the statement $\forall x \in A, P(x)$ that one can form for any $A \in \mathcal{P}(X)$. However, contrasting with ordinary sets, the symbols $x \in F$ don't mean anything by themselves. For instance we cannot say that a given real number is very large but we can say that for all very large real numbers x we have $1/x < 10^{-6}$. There is also an existential quantifier version which is less crucial but still nice to keep things symmetric. For instance we can prove that a set A in a topological space X is open if and only if for every x in A, every point close to x is in A. Similarly A is closed if and only if, for every x in X, if there exists a point of A close to x then x is in A.

Filters are implemented as sets of sets satisfying three conditions. Thus $F \in \mathcal{P}(\mathcal{P}(X))$ is a filter if:

- $X \in F$,
- $\forall U\, V \in \mathcal{P}(X),\ U \in F$ and $U \subseteq V \implies V \in F$,
- $\forall U\, V \in \mathcal{P}(X),\ U \in F$ and $V \in F \implies U \cap V \in F$.

For instance the filter of points that are close to a point x_0 in a topological space X is the filter $\mathcal{N}_{x_0}$ of neighbourhoods of x_0. The inclusion of $\mathcal{P}(X)$ into $\mathcal{F}(X)$ sends a set A to the set of sets that contain A. The order relation $\mathcal{F}(X)$ is opposite to the order induced by $\mathcal{P}(\mathcal{P}(X))$ (so that the inclusion $\mathcal{P}(X) \hookrightarrow \mathcal{F}(X)$ is indeed order preserving) and the direct image operation associated to $f \colon X \to Y$ sends $F \in \mathcal{F}(X)$ to the set of $U \in \mathcal{P}(Y)$ such that $f^{-1}(U) \in F$. The statement $\forall^f x \in F, P(x)$ is implemented as $\{x | P(x)\} \in F$.

The conditions in the above list translate directly into conditions on the $\forall^f \ldots \in \ldots$ generalised bounded quantifier. Given predicates P and Q on X, we get in order:

- if $\forall x \in X, P(x)$ then $\forall^f x \in F, P(x)$,
- if $\forall x \in X, P(x) \implies Q(x)$ then $\big(\forall^f x \in F, P(x)\big) \implies \big(\forall^f x \in F, Q(x)\big)$,
- if $\forall^f x \in F, P(x)$ and $\forall^f x \in F, Q(x)$ then $\forall^f x \in F, \big(P(x)$ and $Q(x)\big)$.

Of course filters seen as generalised sets do not share *all* properties of ordinary sets. Otherwise one could easily prove the inclusion to be a bijection and nothing would be gained. For instance the conjunction property listed above can obviously be extended to conjunctions involving a finite number of predicates, but not to infinite conjunctions. The infinite-conjunction extension would imply that the image of $\mathcal{P}(X)$ is all of $\mathcal{F}(X)$. It would also break the intended meaning in the case of neighbourhoods. Furthermore, generalised quantifiers do not commute: given filters F and G on X and Y and given a predicate P on $X \times Y$, the statements $\forall^f x \in F, \forall^f y \in G,\ P(x, y)$ and $\forall^f y \in G, \forall^f x \in F,\ P(x, y)$ are *not* equivalent in general.

The moral of this story is that even in formal mathematics one may write things like "for all x close to x_0, ...", which might naïvely sound a bit sloppy. Better yet, formalising such statements yields a neat theory, together with precise lemmas saying how to manipulate its statements. In our opinion, this theory provides an elegant means of capturing our informal intuition.

2.2 Continuity and differentiability lemmas

We often had to prove that functions were continuous or smooth in a certain region. In `mathlib` there is a `continuity` tactic that is able to do this for simple functions, but this tactic can be slow and is often unable to prove more complicated continuity conditions. Therefore we usually proved continuity and smoothness conditions manually.

We noticed that the way one formulates continuity and smoothness lemmas greatly affects the convenience of using them. For example, consider the statement that addition on $\mathbb{R}$ is continuous. We can write this as `continuous_add : continuous (λ p : ℝ × ℝ, p.1 + p.2)`. We would like to be able to use this as follows:

```
example : continuous (λ x : ℝ, x + 3) :=
continuous_add.comp
  (continuous_id.prod_mk continuous_const)
```

but Lean rejects this.

To understand why, we remind the reader about Lean's projection notation (discussed further in [Lew19]). E.g. `continuous_id.prod_mk continuous_const` is short for `continuous.prod_mk continuous_id continuous_const`. We also recall some relevant lemmas from `mathlib`:

```
continuous_id : continuous (λ x, x)
continuous_const : continuous (λ x, a)
continuous.comp : continuous g →
  continuous f → continuous (g ∘ f)
continuous.prod_mk : continuous f →
  continuous g → continuous (λ x, (f x, g x))
```

The problem is that the elaborator gets confused when unifying the function `λ x : ℝ, x + 3` in the statement with the partially elaborated function `(λ p : ℝ × ℝ, p.1 + p.2) ∘ λ x : ℝ, (x, 3)` in the proof. We can fix this by giving more information. The following example is accepted, but is cumbersome:

```
example : continuous (λ x : ℝ, x + 3) :=
continuous_add.comp
  (continuous_id.prod_mk continuous_const :
  continuous (λ x : ℝ, (x, (3 : ℝ))))
```

The situation is much nicer if we replace our original continuity statement with:

```
continuous.add : continuous f →
  continuous g → continuous (λ x, f x + g x)
```

Using this, the following slick proof succeeds:

```
example : continuous (λ x : ℝ, x + 3) :=
continuous_id.add continuous_const
```

In general, instead of stating that a certain function (like addition) is continuous, we prefer to state that if we apply the function to continuous arguments, then the resulting function is continuous. This becomes even more important when the source or target spaces become (products of) spaces of continuous linear maps and the relevant map is evaluating or composing continuous linear maps since such cases require a lot more work from the elaborator. The same holds for many other properties of functions, like differentiability, smoothness, measurability, and continuity at a point or in a set. We use this technique throughout our formalisation and mostly also in `mathlib`.

We had to prove various lemmas showing that functions defined piecewise were continuous. A general version of such a lemma (that follows the preceding principle) is the following:

```
lemma continuous_if {p : α → Prop}
  (hp : ∀ a ∈ frontier {x | p x}, f a = g a)
  (hf : continuous_on f (closure {x | p x}))
  (hg : continuous_on g (closure {x | ¬p x})) :
  continuous (λ a, if p a then f a else g a)
```

We will use a special case of this lemma where `p x` is an inequality between two continuous functions in Theorem 3.4.

2.3 Convexity

Affine spaces over ordered scalars have a much richer structure than those over general scalars. An especially important concept that arises in the presence of a scalar ordering is that of convexity.

For our application, we were concerned with convexity of subsets of vector spaces over the real numbers. Given a real vector space F together with a subset $S \subseteq F$, one says S is convex if for all x, y in S, the line segment joining x and y is contained in S.

`mathlib` contains a substantial library of results about convexity which we found especially useful. There are various ways of defining convexity. In fact `mathlib`'s definition changed several times during the course of our work and on several occasions we had to make adjustments in our work to cater for this. Such adjustments became easier as the convexity API stabilised and we regard this as encouraging empirical evidence in favour of using `mathlib` as a dependency.

While `mathlib` already contained key concepts such as the convex hull of a set, of course not every result we needed about convexity theory already existed. Missing results were added directly to `mathlib`. Examples include Carathéodory's lemma, a characterisation of when a convex set has non-empty interior, and numerous minor technical lemmas. These are all now available to any future `mathlib` consumer.

2.4 Barycentric coordinates and their smoothness

An affine basis for a d-dimensional real vector space F is a set of $d+1$ points $p_0, p_1, \ldots, p_d$ in F such that any q in F can be written as $q = \sum w_i p_i$ for unique scalars w_i such that $\sum w_i = 1$. The scalars w_i are known as the barycentric coordinates of q with respect to the basis p_i.

Barycentric coordinates are the natural coordinates when working with affine-invariant concepts such as convexity. Several of our key constructions required the use of barycentric coordinates and so we added a definition and corresponding API to `mathlib`. Amongst the basic facts we needed to add were:

- any real vector space[2] has an affine basis,
- a point lies in the convex hull of an affine basis iff all of its barycentric coordinates are non-negative,
- a point lies in the *interior* of the convex hull of an affine basis iff all of its barycentric coordinates are strictly positive.

A slightly trickier result we needed concerned the smoothness of barycentric coordinates. Let $A \subseteq F^{d+1}$ be the set of affine bases of F. The result we needed was that the map $F \times A \to \mathbb{R}^{d+1}$, $(q, p_0, \ldots, p_d) \mapsto (w_0, \ldots, w_d)$ is smooth. Note that this is joint smoothness on $F \times A$ so both the point in F and the basis in A are varying simultaneously. This is an example of a geometrically obvious fact that wouldn't receive any explanation in an informal context.

Although one could prove this result using the language of the exterior algebra, we opted for a more elementary proof using determinants. At the time `mathlib`'s theory of determinants was more complete than its theory of the exterior algebra (see also [WS22]). The key ingredients were:

- a barycentric change-of-basis formula,
- the smoothness of the determinant function.

[2] Actually we proved this for affine space over any division ring.

2.5 Partitions of unity

Partitions of unity are an important tool in differential topology. They often allow one to patch a family of local solutions to some problem into the desired global solution. They are used so often that it is challenging to render a complete informal presentation without becoming repetitive. This kind of repetition is especially bad from a formalisation point of view so we extracted an "induction principle" associated to partitions of unity.

`mathlib` already contained a sizeable library of results about partitions of unity, including versions for normal topological spaces and versions for smooth manifolds. Notably, there is no version specialised to normed vector spaces and so one must regard the vector spaces as manifolds in order to invoke `mathlib`'s API for partitions of unity. Technically this comes at the cost of making partitions of unity for vector spaces depend on a certain amount of `mathlib`'s differential geometry library. However from the user's point of view, this is completely hidden as can be seen in the lemmas below.

In the first lemma we want to construct maps $f : E \to F$ such that $\forall x, P(x, f(x))$ for some predicate P on $E \times F$. The assumptions are a convexity assumption on P and the existence of local solutions. Partitions of unity appear in the proof but not in the statement.

Lemma 2.1. *Let E and F be real normed vector spaces. Assume that E is finite dimensional. Let P be a predicate on $E \times F$ such that for all x in E, $\{y \mid P(x, y)\}$ is convex. Let n be a natural number or $+\infty$. Assume that every x has a neighbourhood U on which there exists a $\mathcal{C}^n$ function f such that $\forall x \in U, P(x, f(x))$. Then there is a global $\mathcal{C}^n$ function f such that $\forall x, P(x, f(x))$.*

As a simple application of this lemma, one can prove that, given any function $f : E \to \mathbb{R}$, if f is locally bounded below by a positive constant then f is globally bounded below by a smooth positive function. Indeed the convexity assumption is satisfied because for each x the interval $(0, f(x)]$ is convex and the local boundedness assumption provides (constant) local solutions.

More generally, this lemma allows one to deduce useful approximation results that would be painful to obtain using obvious alternatives such as Stone-Weirstrass or convolution. In particular, we will use the following.

Corollary 2.2. *Given n, E and F as in the lemma, let C be a closed set in E, let $\varepsilon : E \to \mathbb{R}$ be a continuous positive function and let $f : E \to F$ be a continuous function. If f is of class $\mathcal{C}^n$ near C then there exists a function $f' : E \to F$ which is everywhere of class $\mathcal{C}^n$, coincides with f on C and such that $\|f'(x) - f(x)\| \le \varepsilon(x)$ everywhere.*

The abstraction captured in Theorem 2.1 also exposes a mathematical awkwardness that exists for smooth maps, both formally and informally. The problem is that smooth manifolds and maps do not form a cartesian closed category; equivalently, one cannot curry smooth functions.

The point is that, even for finite-dimensional vector spaces E, F (where there are no topological issues) there is no natural smooth structure on the space of smooth maps $E \to F$. For our work, we needed a variant of Theorem 2.1 for a product space $E = E_1 \times E_2$. One would like to be able to deduce this from Theorem 2.1 by taking E and F to be E_1 and $E_2 \to F$ respectively but there is no norm on $E_2 \to F$ and so one cannot talk about smooth maps taking values there.

Even the informal mathematics community has not converged on a solution to this problem[3]. We thus used partitions of unity to prove a second lemma, similar to Theorem 2.1 except for products $E_1 \times E_2$. It is notable that formalism forced us to confront this usually harmless fact about the category of smooth manifolds.

3 Constructing loop families

3.1 Loop families

The corrugation procedure outlined in Section 1 requires having families of loops taking values in some prescribed set, with some prescribed base point, and with prescribed average value. All those loops should also come with a global homotopy from the family of constant loops at the prescribed base points. Everything is gathered in the following proposition where Ω prescribes values, the map β prescribes base points, and g prescribes the average values. The parameter $t \in \mathbb{R}$ is the homotopy parameter while $s \in \mathbb{S}^1$ is the loop parameter. We fix finite-dimensional normed vector spaces E and F over $\mathbb{R}$.

Proposition 3.1. *Let K a compact set in E. Let Ω be an open set in $E \times F$.*

Let β and g be smooth maps from E to F. Write $\Omega_x := \{y \in F \mid (x, y) \in \Omega\}$, assume that $\beta(x) \in \Omega_x$ for all x, and that $g(x) = \beta(x)$ for x near K. Lastly suppose that for every x, $g(x)$ is in $\mathrm{Conv}(\mathrm{Conn}_{\beta(x)}(\Omega_x))$, the convex hull of the connected component of Ω_x containing $\beta(x)$, then there exists a smooth family of loops

$$\gamma \colon E \times \mathbb{R} \times \mathbb{S}^1 \to F, (x, t, s) \mapsto \gamma_x^t(s)$$

such that, for all $x \in E$, all $t \in \mathbb{R}$ and all $s \in \mathbb{S}^1$,

[3]For example though the category of Frölicher spaces is Cartesian closed, it is hardly used and is not locally Cartesian closed.

- $\gamma_x^t(s) \in \Omega_x$,
- $\gamma_x^t(s) = \beta(x)$ *when* $t = 0$ *or* $s = 0$ *or* x *is near* K,
- *the average of* γ_x^1 *is* $g(x)$.

Part of the challenge in formalising this lemma is that there is a lot of freedom in constructing γ, which is problematic when trying to do it consistently when x varies. This issue will be addressed in Section 3.2.

From a formalisation point of view, two questions must be addressed from the beginning. The first one is how to handle maps defined on $\mathbb{S}^1$ and the second one is juggling between curryfied and uncurryfied functions.

On paper the notation $\mathbb{S}^1$ is already ambiguous. It could refer to the unit circle in $\mathbb{R}^2$ (or in $\mathbb{C}$) or to a quotient of $\mathbb{R}$, by the subgroup $\mathbb{Z}$ or $2\pi\mathbb{Z}$ depending on taste. Note that we need to be able to talk about continuous and smooth maps defined on $\mathbb{S}^1$. We ultimately settled on the following:

```
structure loop (X : Type*) :=
(to_fun : ℝ → X)
(per' : ∀ t, to_fun (t + 1) = to_fun t)
```

which packages a function `to_fun` with a periodicity assumption (with period 1 so that our implementation is closest to $\mathbb{R}/\mathbb{Z}$). We then record a coercion from `loop X` to `ℝ → X` so that, given `t : ℝ` and `γ : loop X` one can write `γ t`. Note in particular that there is no type $\mathbb{S}^1$ in this story, only a type which plays the role of functions on $\mathbb{S}^1$ but is not a function type from a foundational point of view. One last note about `loop`: the reason why the `per'` field of `loop` is named with a prime is that it is later restated as `loop.per` in terms of this coercion to function.

With this definition of loops, it is very easy to state that a loop is continuous or smooth. But this is not enough. We also need families of loops, parametrised by topological spaces. In particular we also need loops parametrised by a normed space E, or by $E \times \mathbb{R}$ as in the statement of Theorem 3.1. This creates some tension since we would like to think of such a family of loops as a function on $E \times \mathbb{R} \times \mathbb{S}^1$ but our loops are not true function types and so we must do some extra work in order to obtain the usual conveniences of partial evaluation and currying when working with families of loops. We thus introduced a type class:

```
class has_uncurry (α β γ : Type*) :=
(uncurry : α → β → γ)
```

which records a way to turn an element of `α` into a function from `β` to `γ`, with notation `↿` for `has_uncurry.uncurry`. The most generic use is to uncurry recursively. For instance a function `f : α → β → γ → δ` will be fully uncurried to `↿f : α × β × γ → δ`. Using this typeclass, for every pair of types `α` and `X`, we then register the key instance `has_uncurry (α → loop X) (α × ℝ) X` which allows us to convert any function `φ : α → loop X` to a function `α × ℝ → X`. This way we can state the smoothness conclusion from Theorem 3.1 simply as `𝒞 ∞ ↿γ`.

This setup is not completely bullet-proof: sometimes the elaborator gets confused and needs some help, despite the fact that, contrary to the slightly simplified code displayed above, the actual code declares `β` and `γ` as output parameters for type class instance search. However we are globally satisfied by this encoding.

3.2 Surrounding families

We now discuss the proof of a version of Theorem 3.1, subject to two simplifications. Firstly, we work with *continuous* families of loops. We will smooth these families at the end, taking advantage of the fact that Ω and the surrounding condition are open. Secondly, we work with families of loops that don't have a prescribed average, but which we can reparametrise to have the prescribed average. We will do the reparametrisation in Section 3.3. To ensure that this reparametrisation exists, we need to require that the loop γ_x *surrounds* $g(x)$.

Definition 3.2. A loop γ *surrounds* a point v if there is an affine basis in the image of g such that v has positive barycentric coordinates with respect to this basis.

From the discussion in Section 2.4 it seems that we could have given the definition equivalently as $v \in \mathrm{Conv}(\mathrm{im}\,\gamma)^o$. This is indeed the definition used in the standard references [Gro86, Spr98]. However, this is not clearly equivalent. Notice that $v \in \mathrm{Conv}(A)^o$ does not always imply that there is an affine basis in A such that v has positive barycentric coordinates with respect to this basis. As a counterexample, consider A to be the vertices of a square in the plane, and let v be the center of the square. We define a loop surrounding a point as above, because this is exactly the condition we need.

The first main task in proving the special case of Theorem 3.1 is to construct suitable families of loops γ_x surrounding $g(x)$, by assembling local families of loops. We therefore introduce the following definition.

Definition 3.3. A continuous family of loops $\gamma : E \times [0,1] \times \mathbb{S}^1 \to F$, $(x, t, s) \mapsto \gamma_x^t(s)$ surrounds a map $g : E \to F$ with base $\beta : E \to F$ on $U \subseteq E$ in $\Omega \subseteq E \times F$ if, for every x in U, every $t \in [0, 1]$ and every $s \in \mathbb{S}^1$,

- $\gamma_x^t(s) = \beta(x)$ if $t = 0$ or $s = 0$,

- γ_x^1 surrounds $g(x)$,
- $(x, \gamma_x^t(s)) \in \Omega$.

The space of such families will be denoted by $\mathcal{L}(g, \beta, U, \Omega)$.

In this section we assume the hypotheses of Theorem 3.1, i.e. β and g are smooth maps, $\beta(x) \in \Omega_x$ for all x, and $g(x) \in \mathrm{Conv}(\mathrm{Conn}_{\beta(x)}(\Omega_x))$.

Using Carathéodory's lemma, we can construct a surrounding loop at a single point x and thus obtain an element of $\mathcal{L}(g, \beta, \{x\}, \Omega)$. Since $g(x) \in \mathrm{Conv}(\mathrm{Conn}_{\beta(x)}(\Omega_x))$ and $\mathrm{Conn}_{\beta(x)}(\Omega_x)$ is open, we can obtain an affine basis $B \subseteq \mathrm{Conn}_{\beta(x)}(\Omega_x)$ such that $g(x) \in \mathrm{Conv}(B)^o$. Since $\mathrm{Conn}_{\beta(x)}(\Omega_x)$ is path-connected, we can then find a path in $\mathrm{Conn}_{\beta(x)}(\Omega_x)$ starting at $\beta(x)$ through all points in B. To make it a null-homotopic loop in Ω_x based at $\beta(x)$, we traverse the same path backward. This homotopy provides an element of $\mathcal{L}(g, \beta, \{x\}, \Omega)$.

Moreover, we can even construct families of surrounding loops locally around a point x_0. We take our element $\gamma \in \mathcal{L}(g, \beta, \{x_0\}, \Omega)$ and set

$$\gamma_x^t(s) = \gamma^t(s) + \beta(x) - \beta(x_0).$$

Since Ω is open and barycentric coordinates are smooth, this will give a surrounding family in $\mathcal{L}(g, \beta, U, \Omega)$ for some neighbourhood U of x_0.

The difficulty in constructing global families of surrounding loops is that there are plenty of surrounding loops and we need to choose them consistently. The key feature of the above definition is that the parameter t not only allows us to carry out the corrugation process in the next section, but also brings a "satisfied or refund" guarantee, as explained in the next lemma.

Lemma 3.4. *For any set $U \subseteq E$, $\mathcal{L}(g, \beta, U, \Omega)$ is contractible: for every γ_0 and γ_1 in $\mathcal{L}(g, \beta, U, \Omega)$, there is a continuous map $\delta : [0,1] \times E \times [0,1] \times \mathbb{S}^1 \to F, (\tau, x, t, s) \mapsto \delta_{\tau,x}^t(s)$ which interpolates between γ_0 and γ_1 in $\mathcal{L}(g, \beta, U, \Omega)$.*

The tricky part of this lemma is that we need to make sure that δ always surrounds g. The informal proof is again a nice picture: the idea is to build a path of loops that starts with γ_0 then γ_0 concatenated with a longer and longer initial segment of γ_1 until one reaches the full concatenation of γ_0 and γ_1 at $\tau = 1/2$, and start replacing γ_0 by a shorter and shorter initial segment. For each τ this contains a full copy of either γ_0 or γ_1 hence surrounds g. We now describe how we implemented this picture.

Let $\rho : \mathbb{R} \to \mathbb{R}$ be a piecewise-affine function with $\rho(t) = 1$ for $t \leq \frac{1}{2}$ and $\rho(t) = 0$ for $t \geq 1$. We can define the homotopy δ as follows:

- $\delta_{\tau,x}^t$ moves along the loop $\gamma_{0,x}^{\rho(\tau)t}$ on $[0, 1-\tau]$ (if $\tau < 1$)

- $\delta^t_{\tau,x}$ moves along the loop $\gamma^{\rho(1-\tau)t}_{1,x}$ on $[1-\tau, 1]$ (if $\tau > 0$)

Note that the image of $\delta^1_{\tau,x}$ contains the image of $\gamma^1_{0,x}$ for $\tau \leq \frac{1}{2}$, and the image of $\gamma^1_{1,x}$ for $\tau \geq \frac{1}{2}$. Hence it always surrounds $g(x)$.

The argument that δ is continuous is surprisingly tricky. First of all, δ is defined piecewise, so we have to check that the different cases agree on the frontier. Furthermore, note that if $\tau \to 1$ then $\delta^1_{\tau,x}$ will move along loop $\gamma^{\rho(\tau)t}_{0,x}$ at a speed that tends to $+\infty$, so we need to show that $\gamma^{\rho(\tau)t'}_{0,x'}$ tends uniformly to the constant loop as $(x', \tau, t') \to (x, 1, t)$, which follows from the fact that γ is continuous.

Using this lemma, we can transition between two solutions. Therefore, if we have a solution γ_i near K_i for a compact sets K_i ($i \in \{0, 1\}$), we can find a solution near $K_0 \cup K_1$ that coincides with γ_0 near K_0.

Finally, we can apply this recursively to obtain the following result.

Lemma 3.5. *In the setup of Theorem 3.1, assume we have a continuous family γ of loops defined near K which is based at β, surrounds g and such that each γ^t_x takes values in Ω_x. Then there such a family which is defined on all of E and agrees with γ near K.*

The proof requires finding a countable locally-finite family of compact sets covering E and extending the solution recursively.

3.3 The reparametrisation lemma

The reparametrisation lemma concerns the behaviour of the average value of a smooth loop $\gamma : \mathbb{S}^1 \to F$ when the loop is reparametrised by precomposing it with a diffeomorphism $\phi : \mathbb{S}^1 \to \mathbb{S}^1$.

Given a loop $\gamma : \mathbb{S}^1 \to F$, for some finite-dimensional real vector space F, one may integrate to obtain its average $\overline{\gamma} = \int_0^1 \gamma$. Although this average depends on the loop's parametrisation[4], it satisfies a constraint that depends only on the image of the loop: the average is contained in the closure of the convex hull of the image of γ. Indeed the integral defining the average value is a limit of average values over a finite sample of values and those finite averages belong to $\mathrm{Conv}(\mathrm{im}\,\gamma)$.

The reparametrisation lemma says that conversely, given any point g *surrounded* by γ, there exists a reparametrisation ϕ such that $\gamma \circ \phi$ has average value g.

[4]Intuitively, the parametrisation is the speed at which $\gamma(s)$ moves when s moves with unit speed in $\mathbb{S}^1$.

The reparametrisation lemma thus allows one to reduce the problem of constructing a loop whose average is a given point, to the problem of constructing a loop subject to a condition that depends only on its image.

The idea of the proof is simple: since g is contained in the the convex hull of the image of γ, there exist $s_0, s_1, \ldots, s_d$ and barycentric coordinates $w_0, w_1, \ldots, w_d$ such that:

$$g = \sum w_i \gamma(s_i).$$

If there were no smoothness requirement on ϕ one could define it to be a step function which spends time w_i at each s_i. However because there is a smoothness condition, one must approximate by rounding off the corners of the would-be step function. Using such an approximation it is easy to see that the average of $\gamma \circ \phi$ can be made arbitrarily close to g. In order to find ϕ such that the average is exactly g we use the additional freedom that we may also vary the w_i. Because being surrounded is an open condition, a simple continuity argument shows that this additional freedom is sufficient.

Because the s_i are constant, it is easy to construct the inverse of ϕ, which is what we did. It is constructed as the integral of a sum of approximations to the Dirac delta functions, which we call delta mollifiers.

In fact the reparametrisation lemma holds for *families* of loops and this was the version that we needed. More precisely we proved the following:

Lemma 3.6. *Let E, F be a finite-dimensional normed real vector spaces, γ a smooth family of loops:*

$$\begin{aligned} \gamma : E \times \mathbb{S}^1 &\to F, \\ (x, s) &\mapsto \gamma_x(s), \end{aligned}$$

and $g : E \to F$ a smooth function such that γ_x surrounds $g(x)$ for all x. Then there exists a smooth family of diffeomorphisms ϕ_x of $\mathbb{S}^1$ such that $\overline{\gamma_x \circ \phi_x} = g(x)$ and $\phi_x(0) = 0$ for all x.

The argument outlined above for a single loop works locally in the neighbourhood of any x in E and one uses a partition of unity to globalise all the local solutions into the required family.

Actually in our formalisation, the statement of this lemma is distributed across a definition, a lemma about its smoothness, and a lemma about its average values.

3.4 Proof of the loop construction proposition

Using these ingredients, we can now prove Theorem 3.1.

Proof sketch of Theorem 3.1. Let γ^* be any family of loops in $\mathcal{L}(0, 0, \{0\}, F)$. In a neighborhood U^* of K where $g = \beta$, we can set $\gamma'_x = g(x) + \epsilon\gamma^*$. Here $\epsilon > 0$ must be small enough to ensure that γ'_x lands in Ω_x, which is possible since Ω is open and K is compact. From Theorem 3.5 we obtain a continuous family of surrounding loops γ'_x for all x. We can now approximate γ' with a smooth family of loops γ^S. Next, we reparametrise γ^S_x as discussed in Section 3.3 to obtain a smooth family of loops γ^R_x with average $g(x)$. Finally, we use a smooth cut-off function to transition between $g(x)$ near K to γ^R_x outside U^* to obtain our final family γ that equals g near K. □

There are a couple of nuances to this argument. First, we have to ensure that all our constructions remain in Ω. To do this, we must strengthen the condition on U^*. We can require that there is a $\delta > 0$ such that for all $x \in U^*$ the ball with center $(x, \beta(x))$ and radius 2δ lies on Ω and that the distance between γ'_x and $\beta(x)$ is at most δ. When smoothing, we then require that γ^S lies at most δ from γ'. This ensures that γ (which lies on the segment from g to γ^R) lies in Ω.

A second nuance is that we need to make sure that the smoothed family γ^S still surrounds $g(x)$. We ensure this by requiring that γ^S is close enough to γ' and invoking a lemma that states that all loops close enough to a given loop still surround a given point.

A third nuance is the question of how we obtain a smooth function near a continuous one. Our first plan was to use a convolution with a smooth bump function. We need to require that $(\gamma^S)^t_x(s) = \beta(x)$ on $C = \{(t, s) \mid t = 0 \vee s = 0\}$. We planned to continuously reparametrise γ'_x so that it becomes constant near C and then use the fact that the convolution of a function that is constant near x_0 with a bump function with small enough support doesn't change the value at x_0. However, the problem is that we need to smooth γ' in all arguments (x, t, s), and $(\gamma')^t_x(s)$ varies as x varies, even near C, since in that region it equals $\beta(x)$.

We did not find a way to solve this problem with convolutions, since convolutions do not give you enough control over the resulting function in this case. Instead, we used an argument based on partitions of unity. We use the same argument to ensure that $(\gamma')^t_x(s) = \beta(x)$ near C (which is smooth!) and then we apply Theorem 2.2.

After taking these nuances into account, we finally obtain a proof of Theorem 3.1.

4 Convex integration

4.1 A theorem giving parametricity for free

In this section we explain how to reduce Theorem 1.1 to the case where the parameter space P is trivial.

Denote by Ψ the map from $J^1(E \times P, F)$ to $J^1(E, F)$ sending (x, p, y, ψ) to $(x, y, \psi \circ \iota_{x,P})$ where $\iota_{x,P} : E \to E \times P$ sends v to $(v, 0)$.

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

Lemma 4.1. *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 $\mathcal{R} \subseteq J^1(E, F)$ if and only if $\bar{F}$ is a formal solution of $\mathcal{R}^P := \Psi^{-1}(\mathcal{R})$.*
- *If $\mathcal{R}$ is ample then, for any parameter space P, $\mathcal{R}^P$ is also ample.*

As far as we know, the last item is new. Indeed in an informal account it does not cost much to write that handling parameters only requires complicating notation or proving variations of known lemmas, so the incentive to prove the above lemma is low. Using it, we obtain the parametricity in Theorem 1.1 for free.

Lemma 4.2. *If $\mathcal{R}^P$ satisfies the h-principle (i.e. the conclusion of Theorem 1.1 for all appropriate C, K and $\mathcal{F}_0$) with a trivial parameter space P, then $\mathcal{R}$ satisfies the h-principle with parameter space P.*

4.2 Corrugations

In this section we comment on our formalisation of Theillière's corrugation operation introduced in [The22]. In fact for our purposes we need to generalise the results of [The22] slightly.

Fix a dual pair $p = (\pi, v)$ on E. Recall that this means $\pi \in E^*$, $v \in E$, and $\pi(v) = 1$. Given a family of loops γ_x in F parametrised by x in E, and a real number N the corrugation map $\mathcal{T}_p^N \gamma \colon E \to F$ is defined by:

$$\mathcal{T}_p^N \gamma(x) = \frac{1}{N} \int_0^{N\pi(x)} (\gamma_x(s) - \overline{\gamma}_x)\, ds.$$

We also define the remainder term: $R_p^N \gamma := \mathcal{T}_p^N \partial_x \gamma$ where $\partial_x \gamma \colon E \times \mathbb{S}^1 \to L(E, F)$ is the partial derivative of γ in the direction of E.

Proposition 4.3. *Let $\gamma\colon [0,1] \times E \times \mathbb{S}^1 \to F$ be a smooth family of loops in F parametrised by $[0,1] \times E$. Let K be a compact set in E and let ε be a positive real number. Then:*

- *$(t, x) \mapsto \mathcal{T}_p^N \gamma_t(x)$ is smooth,*
- *for every large N and every x in K, $\|\mathcal{T}_p^N \gamma_t(x)\| \leq \varepsilon$,*
- *for every large N and every x in K, $\|R_p^N \gamma_t(x)\| \leq \varepsilon$,*
- *for every t,*

$$D\mathcal{T}_p^N \gamma_t(x) = \pi \otimes \left(\gamma_{t,x}(N\pi(x)) - \overline{\gamma}_{t,x}\right) + R_p^N \gamma_t(x).$$

The first point in the above proposition wouldn't be stated in an informal context, let alone be proven. If pressed to provide a hint of proof, we would say this map is smooth as a composition of smooth maps. We already discussed how to state and prove smoothness lemma in Section 2.2. The next ingredient of course is a strong library of calculus and integration. At the beginning of this project, `mathlib` already contained such libraries, including the fundamental theorem of calculus (see [vD21] for general explanation about integration and measure theory in `mathlib`). However it had nothing about parametric integrals. We thus needed to add the following lemma.

Lemma 4.4. *Let F be a real Banach space and H be a finite-dimensional real normed space. Let n be a natural number or $+\infty$, and let a be a real number. Assume $s\colon H \to \mathbb{R}$ and $\Phi\colon H \times \mathbb{R} \to F$ are of class $\mathcal{C}^n$. Then the function defined by $x \mapsto \int_a^{s(x)} \Phi(x,t)\,dt$ is of class $\mathcal{C}^n$ and, assuming $n > 0$, its derivative is*

$$x \mapsto \int_a^{s(x)} \frac{\partial \Phi}{\partial x}(x,t)\,dt + \Phi(x, s(x)) \otimes ds(x).$$

Note this lemma includes a version of the fundamental theorem of calculus when $H = \mathbb{R}$, Φ does not depend on x, and $s = \mathrm{id}$. So this theorem is obviously an ingredient in the above lemma. The other ingredient is the dominated convergence theorem which allows one to swap the order integration and differentiation when s is constant. One might argue that invoking dominated convergence is excessive in our situation where F is finite dimensional, we assume continuous differentiability everywhere, and the integration domain is compact. However we proved these lemmas in the broader context of building `mathlib` which strives to be a general-purpose mathematics library. We thus first prove much more general lemmas. The following is a sample statement.

Lemma 4.5. *Let F be a real Banach space and H be a real normed space. Let a_0, b_0, a and ε be real numbers such that $a \in (a_0, b_0)$ and $\varepsilon > 0$. Let x_0 be a point in H. Let $s\colon H \to \mathbb{R}$, $\Phi : H \times \mathbb{R} \to F$, $\Phi' : \mathbb{R} \to L(H, F)$, and $b\colon \mathbb{R} \to \mathbb{R}$ be functions. Suppose that the following properties hold:*

- *$s(x_0) \in (a_0, b_0)$ and s is differentiable at x_0 with derivative $s' \in L(H, \mathbb{R})$;*
- *$\Phi(x, \cdot)$ is almost everywhere strongly measurable on (a_0, b_0) for all $x \in B(x_0, \varepsilon)$;*
- *$\Phi(x_0, \cdot)$ is integrable on (a_0, b_0) and continuous at $s(x_0)$;*
- *Φ' is almost everywhere strongly measurable on $(a, s(x_0))$;*
- *for almost every t in (a_0, b_0), $\Phi(\cdot, t)$ is $b(t)$-Lipschitz on the ball $B(x_0, \varepsilon)$;*
- *b is integrable on (a_0, b_0), non-negative and continuous at x_0;*
- *for almost every t in $(a, s(x_0))$, $\Phi(\cdot, t)$ is differentiable at x_0 with derivative $\Phi'(t)$.*

Then Φ' is integrable on $(a, s(x_0))$ and the function defined by $x \mapsto \int_a^{s(x))} \Phi(x, t)\, dt$ is differentiable at x_0 with derivative

$$\int_a^{s(x_0)} \Phi'(t)\, dt + \Phi(x_0, s(x_0)) \otimes s'.$$

We have not been able to find the above statement in any textbook. It has rather minimalistic assumptions that are quite subtle. For instance the positive radius ε is meant to ensure some uniformity in t when requiring that $\Phi(\cdot, t)$ is $b(t)$-Lipschitz near x_0. Also note that differentiability of $\Phi(\cdot, t)$ is assumed only at x_0 and we do not require any bound on Φ', this is deduced from the local Lipschitz assumption. We have many variations on this lemma including versions where the measure used isn't the Lebesgue measure. They would be rather difficult to write without a proof assistant (we certainly wrote several wrong variations in the process). Those lemmas are not really meant to be used in such a generality, but they are meant as common foundations for various lemmas with stronger assumptions.

A further remark about proving these kinds of lemmas is that informal accounts are typically very sloppy about handling the fact that $s(x)$ could cross a in the context of Theorem 4.4. The formalised version requires care here.

One last ingredient in the proof of Theorem 4.3 is that any continuous loop is bounded, and this holds uniformly with respect to any parameter moving in a compact set. Our representation of loops as periodic functions instead of as functions on the compact space $\mathbb{S}^1$ means we must do some extra work to invoke this fact. We thus introduced $\mathbb{R}/\mathbb{Z}$, denoted $\mathbb{S}_1$ in the code, and some glue to go back and forth between 1-periodic functions and functions on $\mathbb{S}_1$. However this glue is tightly encapsulated: we only use it to prove that, given any separated topological space X and any compact set K in X, every continuous function f from $X \times \mathbb{R}$ to a normed space such that each $f(x, \cdot)$ is 1-periodic is bounded on $K \times \mathbb{R}$. Yet again we benefitted from `mathlib`'s strong topology library, invoking results about quotient maps and separated quotient spaces.

4.3 The inductive argument

The proof of Theorem 1.1 repeatedly uses the corrugation operation to improve the given formal solution in more and more directions. Stating this precisely requires a refinement of the notion of being holonomic. Given a linear subspace $E' \subseteq E$, we say that $(f, \varphi) : E \to F \times L(E, F)$ is E'-holonomic at x if $Df(x)$ and $\varphi(x)$ coincide on E'.

Lemma 4.6. *Let $\mathcal{F}$ be a formal solution of some open and ample $\mathcal{R} \subseteq J^1(E, F)$. Let $K_1 \subseteq E$ be a compact subset, and let K_0 be a compact subset of the interior of K_1. Let C be a closed subset of E. Let p be a dual pair on E and let E' be a linear subspace of E contained in $\ker \pi$. Let ε be a positive real number.*

Assume that $\mathcal{F}$ is E'–holonomic near K_0, and holonomic near C. Then there is a homotopy $\mathcal{F}_t$ such that:

1. *$\mathcal{F}_0 = \mathcal{F}$,*
2. *$\mathcal{F}_t$ is a formal solution of $\mathcal{R}$ for all t ,*
3. *$\mathcal{F}_t(x) = \mathcal{F}(x)$ for all t when x is near C or outside K_1 ,*
4. *$\| \mathrm{pr}_F \mathcal{F}_t(x) - \mathrm{pr}_F \mathcal{F}(x)\| \leq \varepsilon$ for all t and all x ,*
5. *$\mathcal{F}_1$ is $E' \oplus \mathbb{R}v$–holonomic near K_0.*

Proof sketch. We denote the components of $\mathcal{F}$ by f and φ. Since $\mathcal{R}$ is ample, Theorem 3.1 applied to $g : x \mapsto Df(x)v$, $\beta : x \mapsto \varphi(x)v$, $\Omega_x = \mathcal{R}(\mathcal{F}(x), p)$, and $K = C \cap K_1$ gives us a smooth family of loops $\gamma : E \times [0, 1] \times \mathbb{S}^1 \to F$ such that, for all x:

- $\forall t\, s,\ \gamma_x^t(s) \in \mathcal{R}(\mathcal{F}(x), \pi, v)$,

- $\forall s,\ \gamma_x^0(s) = \varphi(x)v$,
- $\bar{\gamma}_x^1 = Df(x)v$,
- if x is near C, $\forall t\, s,\ \gamma_x^t(s) = \varphi(x)v$.

Let $\rho : E \to \mathbb{R}$ be a smooth cut-off function which equals one on a neighbourhood of K_0 and whose support is contained in K_1.

Let N be a positive real and set $\mathcal{F}_t(x) = \big(f_t(x), \varphi_t(x)\big)$ where:

$$f_t(x) = f(x) + t\rho(x)\mathcal{T}_p^N \gamma_t,$$

and:

$$\varphi_t(x) = \Upsilon_p\left(\varphi(x),\, \gamma_x^{t\rho(x)}(N\pi(x))\right) + R_p^N \gamma^1.$$

One then checks this homotopy is suitable using Theorem 4.3 □

Given Theorem 4.2, the preceding lemma allows us to prove Theorem 1.1 by induction on a basis of directions in E. Specifically we choose a basis $e : \{1, \ldots, n\} \to E$, take the dual basis e^*, and apply Theorem 4.6 n times using the sequence of dual pairs $p_i = (e_i^*, e_i)$ and subspaces $E_i = \mathrm{Span}(e_1, \ldots, e_i)$.

One formalisation issue is that the whole construction carries around a lot of data. On paper it is easy to state one lemma listing all this data once and proving many properties. For us it was more convenient to give each property its own lemma. Carrying around data, assumptions and constructions thus required some planning. We mitigated this issue by using two ad-hoc structures which partly bundle the data.

The `landscape` structure records three sets in a vector space, a closed set `C` and two nested compact sets $\texttt{K}_0$ and $\texttt{K}_1$. This is the ambient data for the local h-principle result. We call this partly bundled because it doesn't include the data of the formal solution we want to improve. Instead we have a Prop-valued structure `landscape.accepts` that takes a landscape and a formal solution and asserts some compatibility conditions. There are four conditions, which is already enough motivation to introduce a structure instead of one definition using the logical conjunction operator that would lead to awkward and error prone access to the individual conditions.

The proof of this proposition involves an induction on a flag of subspaces (nested subspaces of increasing dimensions). For the purpose of this induction we use a second structure `step_landscape` that extends `landscape` with two more pieces of data, a subspace and a dual pair, and a compatibility condition, namely the subspace has to be in the hyperplane defined by the dual pair.

In this setup the loop family constructed by Theorem 3.1 is used to construct a function whose arguments are some `(L : step_landscape E)`, a formal solution $\mathcal{F}$ and an assumption `(h : L.accepts R` $\mathcal{F}$`)`. Together with corrugation, it is used to build the homotopy of 1-jet sections appearing in the proof of Theorem 4.6 improving the formal solution $\mathcal{F}$ in that step of the main inductive proof. A rather long series of lemmas prove all the required properties of that homotopy, corresponding to all conclusions of Theorem 4.6.

In the inductive construction itself, all conclusions are stated at once since the induction requires knowing about each of them to proceed to the next step. We could have introduced one more ad-hoc structure to record those conclusions but this isn't needed (at least in that part of the project) since we need to access its components only once.

We finish with a comment on induction in the context of Lean. Since it is based on the calculus of inductive construction, Lean's foundations have built-in support for inductive constructions. This can be used for instance to build the addition on natural numbers. In this project we are talking about a *much* more involved inductive construction where each piece requires one to prove many facts to proceed. In principle it would be possible to use an inductive construction in the foundational sense but that would be extremely cumbersome. Instead we state some existential statement and prove it by induction.

5 Sphere eversion

In this section we explain how to derive Smale's theorem from Theorem 1.1. This is less direct than deriving it from the global version of Gromov's theorem (for maps between manifolds), but still rather easy since the source manifold $\mathbb{S}^2$ has a simple model as a subset of a vector space and the target manifold is just a vector space. In this section E is a 3-dimensional real vector space equipped with a scalar product and $\mathbb{S}^2$ is the unit sphere in E. For any point x in $\mathbb{S}^2$, the tangent space $T_x\mathbb{S}^2$ to $\mathbb{S}^2$ at x is the subspace $x^\perp$ of E orthogonal to the line spanned by x. An immersion of $\mathbb{S}^2$ into E is a smooth map f defined near $\mathbb{S}^2$ and such that for every x in $\mathbb{S}^2$, $Df(x)$ is injective on $T_x\mathbb{S}^2$. At face value this may sound slightly stronger than the definition of an abstract immersion from manifold theory. But one can easily prove that any abstract immersion extends to an immersion in the elementary sense. In any case, using a definition of immersion that is too strong would only make Smale's theorem stronger.

Theorem 5.1 (Smale [Sma58])**.** *There is a homotopy of immersions of $\mathbb{S}^2$ into E from the inclusion map to the antipodal map $a\colon q \mapsto -q$.*

Because we want to deduce this from our statement about maps between vector spaces, we need to be slightly careful. We denote by B the ball with radius 9/10 in E. The relation we use is

$$\mathcal{R} = \left\{(x, y, \varphi) \in J^1(E, E) \mid x \notin B \implies \varphi|_{x^\perp} \text{ is injective}\right\}.$$

Proposition 5.2. *Any solution of $\mathcal{R}$ is an immersion of $\mathbb{S}^2$ into E. The relation $\mathcal{R}$ is open and ample.*

As far as we know, this proposition is new. This makes sense because it is not needed to deduce Smale's theorem from the global version of Gromov's theorem. We will explain some ideas of the proof since it also provides examples of geometrically-obvious facts whose proofs need some thought.

In order to prove that $\mathcal{R}$ is open, the main task is to fix $x_0 \notin B$ and $\varphi_0 \in L(E, E,)$ which is injective on $x_0^\perp$ and prove that, for every x close to x_0 and φ close to φ_0, φ is injective on $x^\perp$. This is a typical situation where geometric intuition makes it feel like there is nothing to prove.

One difficulty is that the subspace $x^\perp$ moves with x. We reduce to a fixed subspace by considering the restriction to $x_0^\perp$ of the orthogonal projection onto $x^\perp$. One can check this is an isomorphism as long as x is not perpendicular to x_0. More precisely, we consider $f\colon J^1(E, E) \to \mathbb{R} \times L(x_0^\perp, E)$ which sends (x, y, φ) to $(\langle x_0, x\rangle, \varphi \circ \mathrm{pr}_{x^\perp} \circ j_0)$ where j_0 is the inclusion of $x_0^\perp$ into E. The set U of injective linear maps is open in $L(x_0^\perp, E)$ and the map f is continuous hence the preimage of $\{0\}^c \times U$ is open. This is good enough for us because injectivity of $\varphi \circ \mathrm{pr}_{x^\perp} \circ j_0$ implies injectivity of φ on the image of $\mathrm{pr}_{x^\perp} \circ j_0$ which is $x^\perp$ whenever $\langle x_0, x\rangle \neq 0$.

The next thing to prove is ampleness of $\mathcal{R}$. The key observation is that if one fixes vector spaces F and F', a dual pair (π, v) on F, and an injective linear map $\varphi\colon F \to F'$ then the updated map $\Upsilon_p(\varphi, w)$ is injective if and only if w is not in $\varphi(\ker \pi)$. It then only remains to prove one last intuitively obvious result: the complement of a line in a three dimensional space is ample (we actually prove a more general result). We have not been able to find any informal source that provides any explanation of this.

When using the global version of Gromov's theorem to prove Smale's theorem, the preceding key observation is enough to deduce the ampleness of the relevant relation from the ampleness of the complement of a subspace with codimension at least two. In our case we still have some work to do. It suffices to prove that for every $\sigma = (x, y, \varphi) \in \mathcal{R}$ and every dual pair $p = (\pi, v)$ on E, the slice $\mathcal{R}(\sigma, p)$ is ample. If x is in B then $\mathcal{R}(\sigma, p)$ is all of E which is obviously ample. So we assume x is not in B. Since σ is in $\mathcal{R}$, φ is injective on $x^\perp$. The slice is the set of w such that $\Upsilon_p(\varphi, w)$ is injective on $x^\perp$. Assume first $\ker \pi = x^\perp$. Then $\Upsilon_p(\varphi, w)$ coincides with φ on $x^\perp$ hence the slice is all of E. Assume now that $\ker \pi \neq x^\perp$. The slice is not very easy to picture in this case but one should remember that, up to affine isomorphism, the slice depends only on $\ker \pi$. More precisely, if we keep π but change v then the slice is simply translated in E. Here we replace v by the projection on $x^\perp$ of the vector dual to π rescaled to keep the property $\pi(v) = 1$. What has been gained is that we now have $v \in x^\perp$ and $x^\perp = (x^\perp \cap \ker \pi) \oplus \mathbb{R}v$. Since φ is injective on $x^\perp$, $\varphi(x^\perp \cap \ker \pi)$ is a line, and $\Upsilon_p(\varphi, w)$ is injective on $x^\perp$ if and only if w is in the complement of this line according to the key observation above. So we are indeed back to the fact that the complement of a line is ample.

Lastly, we need a homotopy of formal solutions of $\mathcal{R}$. Roughly, we want to use $\mathcal{F} : (t, x) \mapsto (1 - 2tx, \mathrm{Rot}_{\pi t, x})$ where $\mathrm{Rot}_{\alpha,x}$ is the rotation with angle α around the axis spanned by x. But some extra care is needed to ensure smoothness near the origin (this artificial difficulty is the price we must pay for extending our domain to all of E).

6 Blueprint infrastructure

Before the formalisation started the first author wrote a detailed blueprint in LaTeX with all the definitions and lemmas that were expected to be required for the proof. This was meant to prepare the formalisation work and allow contributions from people who did not know the area of mathematics. This is a well-known strategy, see [Hal12] for an implementation at a much bigger scale.

The new ingredient was to write and use a plugin for plasTeX, a very extensible TeX compiler written in python. This allows one to render the blueprint document in HTML with hyperlinks to precise locations in Lean files corresponding to each result. The software also produces a dependency graph that shows the progress of the project and assists with coordination. This graph is based on manual dependency declarations so that we can indicate a dependency even in the absence of a LaTeX reference. Each node of the graph is either a definition or a lemma statement and is colour-coded to indicate whether something is stated or proven or ready to be stated or proven. This leanblueprint plugin is now used by at least half a dozen formalisation projects. Adapting it to work with other proof assistants would be very easy.

During the project we continuously edited the blueprint text to include more explanation that we found to be necessary, or to cater for minor changes in strategy. The end result is that we now have a somewhat bilingual informal/formal account of all our results, in which each side is useful for illuminating the other.

7 Conclusion and future work

We believe that this work demonstrates that arguments in differential topology are not beyond the reach of formalisation. There were indeed many places where informal sources do not provide any explanation, except perhaps a a picture. Providing formal proofs was a rather pleasant process overall.

In order to see what was gained, one can return to the proof sketch we wrote for Theorem 4.6. It is typical of many proofs in differential topology. We rather carefully described a construction, with a quite a lot of input data, and then the reader is expected to agree that this data satisfies the desired conditions. Here we think the main benefit of a formalised version is that the input, assumptions, and desired output are very clearly stated. Reaching that level of clarity is difficult to achieve without a proof assistant. After seeing such a statement, most readers probably still prefer to work out a mental picture rather than seeing details of the proof, but they are in a much better position to do so. And of course having such a precise statement doesn't prevent anyone from also writing a more vague but less intimidating version first.

The next step in this project is to deduce from Theorem 1.1 the global version for maps between smooth manifolds. In informal accounts this seldom occupies more than a couple of paragraphs. Of course this assumes the theory of smooth manifolds is known, including jet spaces. We already have many pieces in `mathlib` but reducing to vector spaces actually requires some care.

In the more distant future, we hope that formalised mathematics will allow people to engage with much of differential topology at many different scales, from rough heuristic pictures to full details according to the reader's wishes.

Acknowledgements

We would like to thank the `mathlib` community for developing a usable library with the requisite material used in this project. There are too many contributors that developed the parts of the library that we used to mention here explicitly, but we want to specifically thank a few of the contributors. We would like to thank Yury Kudryashov and Sébastien Gouëzel for developing the main part of the library for calculus and integration. We also want to thank Yury for the development of partitions of unity. We greatly appreciate Heather Macbeth's work for writing the theory of rotations specifically for this project. We appreciate Anatole Dedecker's contributions to the theory of paths and ample sets. We thank Johan Commelin and Scott Morrison for the first version of Carathéodory's theorem. Special thanks to Gabriel Ebner for help with the `has_uncurry` type-class.

The second author appreciates the support by Fondation Mathématique Jacques Hadamard for the support via a postdoc fellowship.

References

[Com20] The mathlib Community, *The Lean mathematical library*, Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (New York, NY, USA), CPP 2020, Association for Computing Machinery, 2020, p. 367–381, doi:10.1145/3372885.3373824.

[CP88] Thierry Coquand and Christine Paulin, *Inductively defined types*, COLOG-88, International Conference on Computer Logic, Tallinn, USSR, December 1988, Proceedings, 1988, pp. 50–66.

[dMKA+15] Leonardo Mendonça de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer, *The Lean Theorem Prover (System Description)*, Automated Deduction - CADE-25 - 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings, 2015, pp. 378–388.

[Gro86] Mikhael Gromov, *Partial differential relations*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 9, Springer-Verlag, Berlin, 1986, doi:10.1007/978-3-662-02267-2, https://doi.org/10.1007/978-3-662-02267-2. MR 864505

[Hal07] Thomas C. Hales, *The jordan curve theorem, formally and informally*, Am. Math. Mon. **114** (2007), no. 10, 882–894, http://www.jstor.org/stable/27642361.

[Hal12] Thomas C. Hales, *Dense sphere packings*, London Mathematical Society Lecture Note Series, vol. 400, Cambridge University Press, Cambridge, 2012, A blueprint for formal proofs, doi:10.1017/CBO9781139193894, https://doi.org/10.1017/CBO9781139193894. MR 3012355

[HIH13] Johannes Hölzl, Fabian Immler, and Brian Huffman, *Type Classes and Filters for Mathematical Analysis in Isabelle/HOL*, Interactive Theorem Proving - 4th International Conference, ITP 2013, Rennes, France, July 22-26, 2013. Proceedings, 2013, pp. 279–294.

[IT20] Fabian Immler and Yong Kiam Tan, *The poincaré-bendixson theorem in isabelle/hol*, Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020, New Orleans, LA, USA, January 20-21, 2020 (Jasmin Blanchette and Catalin Hritcu, eds.), ACM, 2020, pp. 338–352, doi:10.1145/3372885.3373833, https://doi.org/10.1145/3372885.3373833.

[Lew19] Robert Y. Lewis, *A formal proof of hensel's lemma over the p-adic integers*, Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs (New York, NY, USA), CPP 2019, Association for Computing Machinery, 2019, p. 15–26, doi:10.1145/3293880.3294089, https://doi.org/10.1145/3293880.3294089.

[Sma58] Stephen Smale, *A classification of immersions of the two-sphere*, Trans. Amer. Math. Soc. **90** (1958), 281–290, doi:10.2307/1993205, https://doi.org/10.2307/1993205. MR 104227

[Spr98] David Spring, *Convex integration theory*, Monographs in Mathematics, vol. 92, Birkhäuser Verlag, Basel, 1998, Solutions to the h-principle in geometry and topology, doi:10.1007/978-3-0348-0060-0, https://doi.org/10.1007/978-3-0348-0060-0. MR 1488424

[The22] Mélanie Theillière, *Convex integration theory without integration*, Math. Z. **300** (2022), no. 3, 2737–2770, doi:10.1007/s00209-021-02785-9, https://doi.org/10.1007/s00209-021-02785-9. MR 4381219

[vD21] Floris van Doorn, *Formalized Haar Measure*, 12th International Conference on Interactive Theorem Proving (ITP 2021), Leibniz International Proceedings in Informatics (LIPIcs), vol. 193, 2021, pp. 18:1–18:17, doi:10.4230/LIPIcs.ITP.2021.18.

[WS22] Eric Wieser and Utensil Song, *Formalizing geometric algebra in lean*, Advances in Applied Clifford Algebras **32** (2022), no. 3, doi:10.1007/s00006-021-01164-1, https://doi.org/10.1007/s00006-021-01164-1.