Lecture 1: Minimal Submanifolds and Calibrated Geometry

An introduction to minimal submanifolds, calibrations, calibrated submanifolds, and their link with holonomy

Prelude: Where are we? Where are we Going?

Sometimes we forget, especially when fully immersed in distinct modules at university, that mathematics is not a collection of disconnected subjects but is more of a web of intersecting ideas. Ideas that begin in one corner of maths often reappear in a new way in another corner. I would like to begin these notes by tracing one such path: the web of geometry, highlighting calibrated geometry as a meeting place. It stands at the intersection of differential geometry, Riemannian geometry, symplectic and Kähler geometry, and even algebraic geometry.

Let us begin by sketching a rough sequence of the geometry that one may have thus met:

- **Topological manifolds:** We meet arbitrary-dimensional generalisations of curves and surfaces. The governing property is that these spaces locally resemble Euclidean space, though globally they may differ. Formally, these are topological spaces that are second countable, Hausdorff, and locally Euclidean. Key topics include:
  - Connectedness and compactness
  - Cell complexes
  - Homotopy and the fundamental group
  - The Seifert-Van Kampen theorem
  - Covering maps
  - Homology
- **Smooth manifolds:** A refinement of the topological notion of manifolds, where we define the notion of a smooth atlas, to allow us to speak of derivatives, vector fields, and flows. Here, calculus finds a new home in many dimensions. Key topics:
  - Tangent vectors and covectors
  - Submanifolds
  - Vector and tensor fields
  - Flows
  - Lie derivatives
  - Lie groups
  - Differential forms
  - Orientations
  - de Rham cohomology
  - Foliations
  - Group actions
- **Riemannian geometry:** A branch of differential geometry where we study smooth manifolds with a Riemannian metric (that is, a symmetric, positive-definite metric that varies smoothly from point to point). The advantage of this is that we then have notions of the length of curves, angles, surface area, and volume. Key topics:
  - Riemannian metrics and manifolds
  - Geodesics and exponential map
  - Riemannian curvature tensor
  - Submanifold theory
  - First and second variation formula for curves
  - Jacobi fields
  - Conjugate points and minimising geodesics

  - Completeness of Riemannian manifolds
  - Bonnet Myers and Cartan Hadamard Theorem
  - Space forms
  - Comparison theorem in Riemannian geometry

- **Symplectic geometry:** The correct and geometric formalism to describe dynamics in the Hamitonian framework. Here, we study smooth manifolds which are equipped with a closed, non-degenerate 2-form, called the symplectic 2-form. This allows for the measurement of sizes of two-dimensional objects in the space. This is analogous to the role played by the metric tensor in Riemannian geometry. Key topics:
  - Symplectic forms and Darboux's theorem
  - Canonical symplectic form on cotangent bundles
  - Hamiltonian vector fields and flows
  - Poisson brackets and Poisson manifolds
  - Symplectic group and linear symplectic geometry
  - Lagrangian submanifolds
  - Moment maps and symplectic reduction
- **Algebraic geometry:** The study of spaces defined by polynomial equations. At first glance, this is a geometry of zeros and loci, but it quickly reveals deep connections with number theory, topology, and complex geometry. Over the complex numbers, varieties can be viewed as complex manifolds, linking algebraic geometry to both symplectic and Kähler geometry. Modern algebraic geometry uses the language of schemes, which allows for a uniform treatment of varieties over arbitrary fields and brings arithmetic into the picture. Key topics:
  - Affine and projective varieties
  - Morphisms of varieties and maps given by polynomials
  - Coordinate rings and the Nullstellensatz
  - Schemes and sheaves
  - Line bundles and divisors
  - Cohomology of sheaves
  - Intersection theory
  - Moduli spaces (e.g. curves, vector bundles)
  - Connection with complex and Kähler geometry

Whilst it can feel like each of these topics is somewhat solitary, really they are much connected. The following diagram aims to demonstrate this:

Calibrated geometry enters precisely here, sitting at the intersection, as demonstrated in the diagram. To even define a calibration, we need the language of smooth manifolds and differential forms. To understand the minimality they guarantee, we rely on ideas from Riemannian geometry, where notions of length, area, and volume are made precise. To see their most natural examples, we turn to symplectic, complex and Kähler geometries, where Lagrangian and holomorphic submanifolds reveal themselves as calibrated. And in the broader picture, algebraic geometry provides the moduli-theoretic and topological frameworks where these objects can be studied in families.

Thus, bear in mind whilst studying the topic, that calibrated geometry is not a topic one can introduce in isolation. It is the natural meeting place of these strands, and such synthesis only makes sense once the background is sketched. Having refreshed one's mind on the areas of geometry, what follows will be the study of a particular and beautiful path across it.

1.1 An Introduction to Minimal Submanifolds

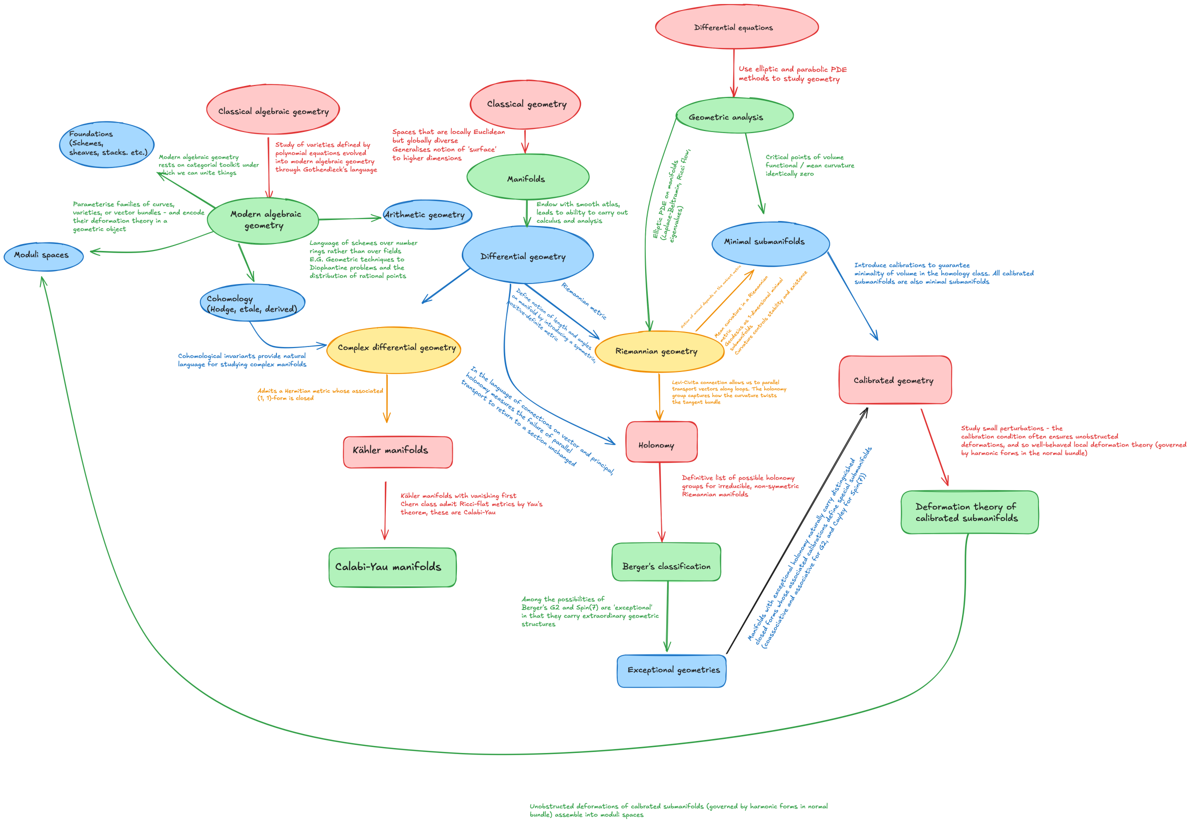

Figure 1: The geometric connections

1.1.1 Minimal Surfaces in Euclidean Space

Our first question is an ancient one:

'Among all surfaces in $\mathbb{R}^3$ with a fixed boundary curve, which one has least area?'

This is known as Plateau's problem.

Historical Note

Despite its name, the roots of this problem trace back to Lagrange in the 18th century, who wrote down the Euler-Lagrange equations for the area functional. In the 19th century, Plateau gave physical intuition for the concept, studying soap films suspended across wire frames.

Minimal surfaces are natural solutions to this problem. They are characterised by the property that every sufficiently small piece of the surface minimises area among nearby competitors with the same boundary. Analytically, this means that they are critical points of the area functional.

$$\mathcal{A}(\Sigma) = \int_\Sigma dA, \tag{1.1}$$

so the first variation of $\mathcal{A}$ vanishes under compactly supported deformations.

Equivalently, if $\Sigma \subset \mathbb{R}^3$ is a smooth surface with unit normal N, then Σ is minimal if and only if its mean curvature vector H vanishes identically:

$$H \equiv 0. \tag{1.2}$$

Geometrically, this means that there is a perfect balance of curvature, meaning that no point on the surface bends more than than another. Physically, this is the reason that soap films naturally realise minimal surfaces: surface tension drives them to configurations where curvature cancels out and area is locally minimised.

1.1.2 Minimal Submanifolds in a Riemannian Manifold

The notion of minimal surfaces in Euclidean space extends naturally to submanifolds of arbitrary Riemannian manifolds. To set the stage, let us recall the definition.

Definition 1.1: Minimal submanifolds

Let (M, g) be an n-dimensional Riemannian manifold, and let $\Sigma^k \subset M$ be an immersed k-dimensional submanifold. We say that Σ is minimal if it is a critical point of the volume functional

$$\mathrm{Vol}(\Sigma) = \int_\Sigma d\mu_\Sigma, \tag{1.3}$$

where $d\mu_\Sigma$ is the induced k-dimensional volume form.
That is, for every smooth variation $\Phi_t : \Sigma \to M$ with $\Phi_0 = \mathbb{1}_\Sigma$, we have:

$$\frac{d}{dt}\mathrm{Vol}(\Phi_t(\Sigma))|_{t=0} = 0. \tag{1.4}$$

Equivalently, Σ is minimal if and only if its mean curvature vector H vanishes identically:

$$H \equiv 0. \tag{1.5}$$

The condition that $H \equiv 0$ is the higher-dimensional analogue of the vanishing mean curvature for surfaces in $\mathbb{R}^3$. In every normal direction, the extrinsic curvatures cancel out, so that the submanifold 'balances' itself with the ambient geometry.

One remarks that at no point do we require that Σ minimises the volume. The definition only requires stationarity of the volume functional, not that Σ globally minimises volume among all competitors with the same boundary. For example, a saddle surface in $\mathbb{R}^3$ is minimal, but does not minimise area globally. In many contexts, the stronger property of being volume-minimising is what one truly seeks, so it would be useful to have a stronger notion...

1.2 Calibrated Geometry

We begin with a question:

Can we identify submanifolds which minimise volume within their entire homology class?

In practice, to prove such global minimality is notoriously difficult. The minimality condition $H = 0$ translates into a non-linear elliptic PDE, solvable with only delicate analytic estimates and variational methods. Calibrated geometry offers a different path: it provides a purely differential-geometric criterion ensuring that certain submanifolds are not only stationary, but are absolutely volume-minimising in their homology class.

How, then, should one picture this? A calibration is like a geometric measuring tool. Apply it to any oriented k-dimensional surface in the manifold, and it will always overestimate the volume. Yet for a distinguished family of submanifolds, the calibration is perfectly accurate: it measures their volume exactly. Those special submanifolds are called calibrated, and they are automatically volume-minimising in their homology class.

The advantage is this: calibrated submanifolds are automatically volume-minimising in their homology class, no PDEs required. The existence of a calibration rigidly locks in their minimality.

Interlude: Differential Forms as Measuring Devices

Before introducing calibrated geometry in its mathematical formality, it is useful to recall how differential forms act as measuring tools on a Riemannian manifold. There is nothing new here for a reader who is familiar with differential geometry, but it will allow us to introduce the notation and perspective that shall be used throughout.

Let (M, g) be an n-dimensional Riemannian manifold.

- A k-form $\varphi \in \Omega^k(M)$ is a smooth alternating k-linear functional on tangent vectors. At each point $p \in M$m the form φ_p 'eats' k tangent vectors and returns a number $\varphi_p(v_1, ..., v_k) \in \mathbb{R}$.
- Geometrically, if $v_1, ..., v_k$ span a k-dimensional parallelepiped in T_pM, then $\varphi_p(v_1, ..., v_k)$ is (up to a factor) the signed k-volume of that parallelepiped.

The Riemannian metric g equips each tangent space with a natural notion of k-dimensional volume. If ξ is a decomposable unit k-vector (that is, an oriented k-plane with orthonormal basis $e_1, ..., e_k$), then the induced volume form $d\mu_\xi$ satisfies

$$d\mu_\xi(e_1, ..., e_k) = 1. \tag{1.6}$$

Given this, it makes sense to compare a k-form φ with the true geometric volume. The relevant measure is the comass norm:

$$||\varphi_p||^* = \sup\{\varphi_p(\xi) : \xi \in \Lambda^k(T_pM), \xi \text{ unit, simple, oriented}\}. \tag{1.7}$$

Intuitively, $|\varphi_p *|$ tells us how large φ can appear when tested against a unit k-plane. If $\varphi_p|* \leq 1$, then φ never reports a number greater than the actual volume.

This is precisely the idea behind calibrations: special closed forms whose comass is at most one, and which exactly match the volume on certain distinguished submanifolds.

1.2.1 Defining a Calibration

Let us present the formal definition of a calibration, which, as described informally, is a type of differential form that *'never overestimates volume'*.

Definition 1.2: Calibration

Let (M, g) be a Riemannian manifold. A differential k-form $\eta \in \Omega^k(M)$ is called a calibration if:

1. η is closed, $d\eta = 0$, and
2. for every point $p \in M$ and every oriented orthonormal k-tuple $(e_1, ..., e_k) \in T_pM$, we have

$$\eta_p(e_1, ..., e_k) \leq 1. \tag{1.8}$$

Equivalently, $|\eta_p|_* \leq 1$ for all p, where $|\cdot|_*$ denotes the comass norm.

Example 1.1

Consider $\mathbb{R}^n$ with its standard Euclidean metric. Any constant-coefficient k-form $\varphi \in \Lambda^k(\mathbb{R}^n)^*$ is closed. By rescaling φ so that its comass norm is one, we obtain a calibration. Moreover, equality holds on at least one k-plane (namely, the one spanned by a wedge of basis vectors appearing in φ).

Of course, we wish to use the definition of a calibration to single out submanifolds where this 'measuring device' exactly matches the true Riemannian volume.

Definition 1.3: Calibrated submanifold

Let η be a calibration k-form on (M, g). An oriented k-dimensional submanifold $N \subset M$ is said to be calibrated by η if the restriction of η to N agrees with the Riemannian volume form of N:

$$\eta|_N = \mathrm{Vol}_N \tag{1.9}$$

Example 1.2

Any oriented k-plane in $\mathbb{R}^n$ can be calibrated. In coordinates, if $P = \{x \in \mathbb{R}^n : x_{k+1} = \cdots x_n = 0\}$ with its natural orientation, then the form

$$\eta = dx_1 \wedge \ldots \wedge dx_k \tag{1.10}$$

is a calibration, and P is calibrated by η.

Aside

In the next section, we shall provide some interesting examples of calibrations (e.g. Kähler, special Lagrangian, associative, coassociative, Cayley), but frankly it doesn't make sense to introduce these without first defining holonomy. We shall see that calibrations appear naturally in manifolds with reduced holonomy, giving us many examples.

Note that the calibration condition is purely algebraic on the tangent spaces of N. Thus, the problem of finding calibrated submanifolds translates into solving a first-order non-linear system for the immersion of N into M, in contrast with the second-order PDE of mean curvature for general minimal submanifolds.

1.2.2 Properties of Calibrated Manifolds

The first proposition formalises an idea that we have previously alluded to, which is that calibrated submanifolds are automatically minimal submanifolds. For simplicity, we shall consider the compact case, but the notion carries to non-compact submanifolds as well.

Proposition 1.1

Let (M, g) be a Riemannian manifold, φ be a calibration on M, and N a compact φ-submanifold in M. Then N is volume-minimising in its homology class.

Proof. Let $\dim(N) = k$, $[N] \in H_k(M, \mathbb{R})$, and $[\varphi] \in H^k(M, \mathbb{R})$ be the homology and cohomology class of N and φ. Then since $\varphi|_{T_xN} = \mathrm{Vol}_{T_xN}$ for each $x \in N$ by the fact it is a calibrated submanifold, one has that:

$$[\varphi] \cdot [N] = \int_{x \in N} \varphi|_{T_xN} = \int_{x \in N} \mathrm{Vol}_{T_xN} = \mathrm{Vol}(N). \tag{1.11}$$

Let N' be any other compact k-submanifold of M, with $[N'] = [N]$ in $H_k(M, \mathbb{R})$. Then since $\varphi|_{T_xN'} \leq \mathrm{Vol}_{T_xN'}$ as φ is a calibration, one has:

$$[\varphi] \cdot [N] = [\varphi] \cdot [N'] = \int_{x \in N'} \varphi|_{T_xN'} \leq \int_{x \in N'} \mathrm{Vol}_{T_xN'} = \mathrm{Vol}(N'). \tag{1.12}$$

One can then conclude that $\mathrm{Vol}(N) \leq \mathrm{Vol}(N')$. Thus, N is volume-minimising in its homology class. □