

Manifolds and Lie Groups

Visual SLAM Basics, Session A1: Manifolds

The geometry underneath every state we estimate

Ryan Tse

Roadmap

A1: manifolds (about 55 min)

1. Smooth structure: charts
2. Smooth maps
3. Tangent vectors
4. Regular level sets
5. Vector fields and flows
6. Retractions


Three running robots: **two-link arm** **unicycle** **camera**

Goal: by the end, " $\mathbf{T} \leftarrow \mathbf{T} \exp(\hat{\delta})$ " from Part 1 is a sentence you can defend.

A2: Lie groups (about 45 min)

7. Lie groups
8. Actions and equivariance
9. Left-invariant fields, Lie algebra
10. Lie bracket
11. The exponential map

Everything we estimate lives on a curved space




Joint angles: torus. Orientation: $SO(3)$. Camera pose: $SE(3)$. Directions and bearings: sphere.

Why “just use vectors” fails



Yaw 179° vs -179° : far in $\mathbb{R}$, close in truth.



$\|q + \Delta q\| \neq 1$: left the circle


An additive update leaves the set of rotations.

We need: **coordinates** (to compute), **tangent spaces** (to linearize), and a way to **step** (to stay on the space).


A1: Smooth manifolds

Charts, maps, tangent vectors, vector fields, retractions

Locally, it looks like $\mathbb{R}^n$



curved globally




flat up close: a chart $\varphi : U \rightarrow \mathbb{R}^2$

A **chart** is a bijection $\varphi : U \rightarrow \mathbb{R}^n$ from an open patch $U \subset \mathcal{M}$ onto an open set of $\mathbb{R}^n$.

- Gives **coordinates** on the patch
- Every point is covered by some chart
- No single chart works globally (in general)

Example: the circle needs two charts




Angle charts: $\varphi_1 = \theta \in (0, 2\pi)$ and $\varphi_2 = \theta \in (-\pi, \pi)$.

On the overlap, the **transition map** is $\varphi_2 \circ \varphi_1^{-1}(\theta) = \theta$ or $\theta - 2\pi$: smooth.

This is the yaw wrap-around from before: not a bug, a chart boundary.

Atlas and transition maps






A **smooth atlas**: charts covering $\mathcal{M}$ whose transition maps are smooth (C^∞) wherever they are defined.

Smooth manifold = space + (maximal) smooth atlas.

Fine print we skip: Hausdorff, second countable.


Robot example: the two-link arm lives on a torus




Configuration $(q_1, q_2) \in S^1 \times S^1 = T^2$.

Cut the torus open: a square with opposite edges glued. The charts are the angle charts from the circle, twice.

Euler angles are one chart of $SO(3)$, and it breaks



pitch = 25°: three distinct axes



pitch = 90°: gimbal lock

Roll–pitch–yaw is a **chart**:

$$(\phi, \theta, \psi) \mapsto \mathbf{R}_z \mathbf{R}_y \mathbf{R}_x.$$


At pitch = $\pm 90^\circ$ two axes align and the angles-to-rotation map loses rank (gimbal lock). These rotations lie outside this chart's domain.

The rotation group is fine there. Only this coordinate system fails.


Invertible vs not: when do coordinates work?

A chart must be a **two-way street**: coordinates $\leftrightarrow$ points, smooth in both directions. Locally that means the Jacobian of the map has **full rank**.

invertible: nearby points $\leftrightarrow$ nearby angles



rank lost: not invertible



- **Gimbal lock**: many angle pairs give one rotation, and some nearby rotations are unreachable.
- **Arm at $q_2 = 0, \pi$** : the hand cannot move radially (the ellipse collapses to a segment).

Same disease, two robots. Singularities belong to the coordinates or the map, not to the space.


Smooth maps: smooth in coordinates

$$\begin{array}{ccc} \mathcal{M} & \xrightarrow{F} & \mathcal{N} \\ \downarrow \varphi & & \downarrow \psi \\ \mathbb{R}^m & \xrightarrow{\hat{F} = \psi \circ F \circ \varphi^{-1}} & \mathbb{R}^n \end{array}$$

F is **smooth** if its coordinate representation $\hat{F}$ is smooth for the charts around each point.

Smooth maps are the only maps we may differentiate. All of SLAM's cost functions are of this kind.

Robot example: forward kinematics



$F(q)$ is the **hand (end-effector) position** in the plane, given joint angles $q = (q_1, q_2)$ and link lengths l_1, l_2 .

$$F : T^2 \rightarrow \mathbb{R}^2,$$

$$F(q) = \begin{pmatrix} l_1 \cos q_1 + l_2 \cos(q_1 + q_2) \\ l_1 \sin q_1 + l_2 \sin(q_1 + q_2) \end{pmatrix}$$


Torus in, plane out: not injective, not onto, but smooth.

The same pattern in SLAM

- **Camera projection:** $\pi(\mathbf{K}, \mathbf{T}_{CW}, \mathbf{P}_w) : \text{SE}(3) \times \mathbb{R}^3 \rightarrow \mathbb{R}^2$
- **Reprojection error:** $\mathbf{e}(\mathbf{T}, \mathbf{P}) = \mathbf{u} - \pi(\mathbf{K}, \mathbf{T}, \mathbf{P})$
- **Cost:** $J(\mathbf{T}_1, \dots, \mathbf{T}_N, \mathbf{P}_1, \dots) = \sum \|\mathbf{e}_{ij}\|_{\Sigma}^2$, a smooth map from a product of manifolds to $\mathbb{R}$

To minimize J we need its derivative. A derivative at a point of a manifold: a map between **tangent spaces**.

Tangent vectors: velocities of curves



Take a smooth curve $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M}$ with $\gamma(0) = p$. Its velocity

$$v = \dot{\gamma}(0)$$

is a **tangent vector** at p .


$T_p \mathcal{M}$ = all such velocities: an n -dimensional **vector space** attached to p .

Different curves can have the same velocity; that is fine.

In a chart, v is just the vector $(\varphi \circ \gamma)'(0) \in \mathbb{R}^n$.

Robot example: the unicycle's velocity

pose $(x, y, \theta) \in SE(2)$



Pose $(x, y, \theta) \in SE(2)$. A drive path is a curve $\gamma(t)$.


Its tangent vector is $(\dot{x}, \dot{y}, \dot{\theta})$.

The set of all velocities at one pose is a 3-dimensional space. The robot can only realize the 2D subspace where $(\dot{x}, \dot{y})$ points along its heading.

Differential: the Jacobian, coordinate-free



For smooth $F : \mathcal{M} \rightarrow \mathcal{N}$ and $p \in \mathcal{M}$, the **differential**

$$dF_p : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{N}, \quad dF_p(\dot{\gamma}(0)) = \left. \frac{d}{dt} F(\gamma(t)) \right|_{t=0}$$



In coordinates, dF_p is the **Jacobian matrix** of $\hat{F}$. The chain rule is the usual one:
 $d(G \circ F)_p = dG_{F(p)} \circ dF_p$.

Robot example: the arm Jacobian



$$dF_q = \begin{pmatrix} -l_1 s_1 - l_2 s_{12} & -l_2 s_{12} \\ l_1 c_1 + l_2 c_{12} & l_2 c_{12} \end{pmatrix}$$


where $s_1 = \sin q_1$, $c_1 = \cos q_1$,
 $s_{12} = \sin(q_1 + q_2)$, $c_{12} = \cos(q_1 + q_2)$.

Columns: how the hand moves per unit $\dot{q}_1$,
 $\dot{q}_2$. $\det = l_1 l_2 \sin q_2$.

The unit circle of joint rates maps to an **ellipse** of hand velocities. At $q_2 = 0$ or π it collapses to a segment: a **singularity**, where dF_q loses rank.

This is the same object that Gauss–Newton will linearize in session B.

The tangent bundle: all velocities at once



TS^1 : a line of velocities $\dot{\theta}$ at each θ

$T\mathcal{M} = \bigsqcup_{p \in \mathcal{M}} T_p\mathcal{M}$, itself a smooth manifold of dimension $2n$.

Example: a pendulum's phase space $(\theta, \dot{\theta})$ is TS^1 , a cylinder.


State = configuration plus velocity: the IMU-style state of a robot lives in a tangent bundle.

Remark: other definitions of $T_p\mathcal{M}$

- **Curves** (used here): velocities of curves. Best for pictures.
- **Derivations**: linear maps D on smooth functions with $D(fg) = f Dg + g Df$. Best for proofs (Lee).
- **Coordinates**: vectors in $\mathbb{R}^n$ that transform by the Jacobian of the transition map.

All three give the same vector space. Derivation version in the backup slides.

Regular level sets are manifolds





Let $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^k$ be smooth. If $d\Phi_p$ has full rank k at every $p \in \Phi^{-1}(c)$, then

- $\Phi^{-1}(c)$ is a smooth manifold of dimension $N - k$
- $T_p \Phi^{-1}(c) = \ker d\Phi_p$

The unit sphere: $\Phi(x) = \|x\|^2$, $c = 1$, $T_p = p^\perp$.

Robot example: a constraint on the arm

end effector kept on the wall



Keep the hand on a wall $x = a$:

$$\Phi(q) = F_x(q) - a = 0.$$

The feasible set is a 1-D curve in T^2 (dimension $2 - 1$).

Feasible joint velocities: $\ker d\Phi_q$.

Constrained robotics and constrained estimation are the same math.

Example: $SO(3)$ is a level set


Let $\Phi : \mathbb{R}^{3 \times 3} \rightarrow \text{Sym}_3$ (6-dim), $\Phi(\mathbf{R}) = \mathbf{R}^\top \mathbf{R}$. Then $SO(3) = \Phi^{-1}(\mathbf{I}) \cap \{\det > 0\}$.

- $d\Phi_{\mathbf{R}}(\mathbf{V}) = \mathbf{R}^\top \mathbf{V} + \mathbf{V}^\top \mathbf{R}$, which is onto the symmetric matrices (rank 6, checked numerically)
- **Dimension** $= 9 - 6 = 3$
- **Tangent space:** $T_{\mathbf{R}}SO(3) = \{\mathbf{V} : \mathbf{R}^\top \mathbf{V} \text{ is skew-symmetric}\}$

At $\mathbf{R} = \mathbf{I}$: $T_{\mathbf{I}}SO(3) = \{\text{skew-symmetric matrices}\}$. Hold this thought: it is $\mathfrak{so}(3)$.

By counting (rotation + translation): $\dim SE(3) = 3 + 3 = 6$, $\dim SE(2) = 1 + 2 = 3$. Also, by Whitney, every manifold embeds in some $\mathbb{R}^N$ (backup).

Vector fields: a tangent vector at every point



$$X_p = \omega \times p$$

A **vector field** X assigns $X_p \in T_p\mathcal{M}$ to every p , smoothly.

Example: rotation about an axis on S^2 :
 $X_p = \omega \times p$.


A smooth ODE $\dot{p} = X(p)$ on $\mathcal{M}$ is exactly a vector field. Robot dynamics and kinematics are vector fields.

Robot example: the unicycle model

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{pmatrix} = v \underbrace{\begin{pmatrix} \cos \theta \\ \sin \theta \\ 0 \end{pmatrix}}_{X_1} + \omega \underbrace{\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}}_{X_2}$$

Two vector fields on $SE(2)$; the controls (v, ω) choose a combination.

Only two directions out of three at each pose. Can the robot still reach every pose? Yes, and the reason is the Lie bracket (later).




fixed (v, ω) : a circular arc

Integral curves and flows

An **integral curve** of X through p solves $\dot{\gamma}(t) = X_{\gamma(t)}$, $\gamma(0) = p$.

The **flow** $\theta_t(p) = \gamma_p(t)$ pushes every point along the field for time t .




- $\theta_0 = \text{id}$, $\theta_{s+t} = \theta_s \circ \theta_t$
- Unicycle with constant (v, ω) : the flow moves a pose along a circular arc

This is “integrating the robot’s motion”, and it is how the exponential map will be defined.

Pushforward

A smooth map $F : \mathcal{M} \rightarrow \mathcal{N}$ carries tangent vectors forward, $dF_p : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{N}$. For a diffeomorphism F it also pushes a vector *field* along:

$$(F_*X)_{F(p)} = dF_p(X_p).$$




We will use this to compare “same velocity” at different poses (left translation, in A2).

Capstone: retractions

How to take a step without leaving the manifold

The problem: the step leaves the manifold

Optimizers produce a step $\delta \in T_p\mathcal{M}$. We need a new point on $\mathcal{M}$, not $p + \delta$.



A **retraction** at p is a smooth map

$\mathcal{R}_p : T_p\mathcal{M} \rightarrow \mathcal{M}$ with

- $\mathcal{R}_p(0) = p$
- $d(\mathcal{R}_p)_0 = \text{id}_{T_p\mathcal{M}}$ (first-order agreement with straight-line motion)


Any such map is a legitimate way to “add” δ to p .

Examples of retractions

- **Sphere:** $\mathcal{R}_p(\delta) = \frac{p + \delta}{\|p + \delta\|}$ (step, then normalize)
- **SO(3), projection:** $\mathcal{R}_{\mathbf{R}}(\mathbf{R}\hat{\delta}) = \text{proj}_{\text{SO}(3)}(\mathbf{R} + \mathbf{R}\hat{\delta})$ (SVD)
- **SO(3), exponential:** $\mathcal{R}_{\mathbf{R}}(\hat{\delta}) = \mathbf{R} \exp(\hat{\delta})$ (A2)

All agree to first order, which is all Gauss–Newton needs per step. They differ in cost and in how well they track the true motion.

Test: integrate a gyroscope




Constant $\omega = (0.3, 0.5, 0.2)$, 2000 steps, $dt = 0.01$.

	$\ R^T R - I\ $	$\ R - R_{true}\ $
$R + R\hat{\omega}dt$	1.1×10^{-1}	5.5×10^{-2}
+ SVD	8×10^{-16}	2.2×10^{-4}
$R \exp(\hat{\omega}dt)$	2×10^{-13}	9.5×10^{-14}

Additive: leaves $SO(3)$. Projected: valid but slightly off.

Exp: both.

Closing the loop on Part 1



Gauss–Newton on a manifold:

1. Linearize e at the current estimate, in $T_p\mathcal{M}$
2. Solve for $\delta \in T_p\mathcal{M}$
3. Update $p \leftarrow \mathcal{R}_p(\delta)$

This is the “ $\oplus$ ” we waved at in Part 1.

Session B: how to compute the Jacobian with respect to δ .

References

- J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed., Springer, 2013.
- P.-A. Absil, R. Mahony, R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton, 2008.
- J. Solà, J. Deray, D. Atchuthan, “A micro Lie theory for state estimation in robotics”, arXiv:1812.01537.
- T. Barfoot, *State Estimation for Robotics*, Cambridge, 2017.
- K. Lynch, F. Park, *Modern Robotics*, Cambridge, 2017.
- R. Murray, Z. Li, S. Sastry, *A Mathematical Introduction to Robotic Manipulation*, CRC, 1994.

Backup

Backup: tangent vectors as derivations

A **derivation** at p is a linear map $D : C^\infty(\mathcal{M}) \rightarrow \mathbb{R}$ with

$$D(fg) = f(p) Dg + g(p) Df.$$

Each curve γ gives one: $D_\gamma f = (f \circ \gamma)'(0)$, and every derivation arises this way. In coordinates, $D = \sum_i v^i \partial_i|_p$.

This definition needs no curves and no embedding, which is why proofs prefer it.

Backup: Whitney embedding

Every smooth n -manifold embeds smoothly in $\mathbb{R}^{2n}$ (strong version; the weak version gives $\mathbb{R}^{2n+1}$).

So “manifold = level set or submanifold of some $\mathbb{R}^N$ ” loses nothing, though the intrinsic point of view is often cleaner.

$SO(3)$ sits in $\mathbb{R}^9$ by construction; Whitney only promises it also fits in $\mathbb{R}^6$.