

Background Lecture:

Symmetry Primer

Jake Welde

Incoming Assistant Professor,
Sibley School of Mechanical and Aerospace Engineering,
Cornell University

Robotics: Science and Systems
Equivariant Systems Workshop

Los Angeles, California, USA

June 25, 2025

Poster Authors: Upload Your Slides

Add three slides to your section of the Google Slides presentation.



In This Talk

This talk is meant to give a “crash course” on the mathematical foundations of symmetry-informed methods in robotics. We'll briefly touch on:

1. **Rotations and Transformations:** Lie groups you already know and love!
2. **Smooth Manifolds:** the continuous spaces in which we work
3. **Lie Group Fundamentals:** thinking abstractly about continuous transformations
4. **Group Actions:** how Lie groups interact with other spaces
5. **Equivariance and Invariance:** the two most common “flavors” of symmetry
6. **Robotics Sampler:** what equivariance “looks like” in a few example domains

Please, ask questions during the talk!

Rotations and Homogeneous Transforms

Whether you work in control, estimation, perception, or something else, at some point you've no doubt worked with homogeneous transforms:

$$H_b^a = \begin{bmatrix} R_b^a & p_b^a \\ 0 & 1 \end{bmatrix} \in \text{SE}(3),$$

where $p_b^a \in \mathbb{R}^3$ is the position of the origin of frame a in frame b and $R_b^a \in \text{SO}(3)$ is the orientation of frame a in frame b , and

$$\text{SO}(3) = \{R \in \mathbb{R}^{3 \times 3} : RR^T = R^T R = I_{3 \times 3}, \det R = 1\}.$$

In this primer, I'll assume basic familiarity with these ideas, and aim to describe more abstract tools that provide a framework for working with symmetry in robotic systems.

Smooth Manifolds

A *smooth manifold* M is a space that locally “looks like” \mathbb{R}^n , but may have a very different global structure. In this slide, we assume $M \subseteq \mathbb{R}^N$ (simplifies definitions).

Example. The sphere in 3D space is the manifold $\mathbb{S}^2 = \{x \in \mathbb{R}^3 : x^T x = 1\}$.

The *tangent space* at each point $m \in M$ is given by

$$T_m M = \{\dot{\gamma}(0) : \gamma : (-\varepsilon, \varepsilon) \rightarrow M \text{ s.t. } \gamma(0) = m\}.$$

For any map $f : M \rightarrow N$, we can define its *differential* (or *tangent map*, or *derivative*) by

$$df(v_m) := \left. \frac{d}{dt} (f \circ \gamma(t)) \right|_{t=0} = \lim_{\Delta t \rightarrow 0} \frac{f(\gamma(\Delta t)) - f(\gamma(0))}{\Delta t}$$

for *any* smooth curve $\gamma : (-\varepsilon, \varepsilon) \rightarrow M$ such that $\gamma(0) = m$ and $\dot{\gamma}(0) = v_m$.

Groups

A *group* is a set G along with an operation $\cdot : G \times G \rightarrow G$ with three properties:

1. **Associativity:** $(a \cdot b) \cdot c = a \cdot (b \cdot c)$ for all $a, b, c \in G$.
2. **Identity:** There exists some $e \in G$ such that $e \cdot a = a$ for all $a \in G$.
3. **Inverse:** For each $a \in G$, there is some $a^{-1} \in G$ such that $a \cdot a^{-1} = a^{-1} \cdot a = e$.

Example. The permutation group \mathcal{S}_n :

- ▶ matrices where every row and column is a “one-hot” vector $(0, \dots, 1, \dots, 0)$,
- ▶ the group operation is matrix multiplication,
- ▶ and the inverse of any $P \in \mathcal{S}_n$ is P^T , e.g.,

$$\begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Lie Groups

A *Lie group* is a group that is also a smooth manifold and for which the group operation $(g_1, g_2) \mapsto g_1 \cdot g_2$ and inverse map $g \mapsto g^{-1}$ are smooth.

Example. The 2D rotation group under matrix multiplication:

$$\mathrm{SO}(2) = \left\{ \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} : \theta \in [0, 2\pi) \right\}, \quad R(\theta_1)R(\theta_2) = R(\theta_1 + \theta_2), \quad R(\theta)^{-1} = R(-\theta).$$

More Examples. $\mathrm{SO}(n)$, \mathbb{R}^n (under addition), $\mathbb{R}_{>0}$ (under multiplication), ...

Question. Is every manifold a Lie group? **Answer.** Definitely not, e.g., \mathbb{S}^2 .

All examples here are (at least isomorphic to) *matrix Lie groups*, meaning $G \subseteq \mathrm{GL}(n, \mathbb{R})$ (the invertible $n \times n$ real matrices) and the group operation is matrix multiplication.

Group Actions

A *left action* of a group G on a smooth manifold M is a smooth map $\phi : G \times M \rightarrow M$ with the following two properties:

1. **Identity:** $\phi(e, m) = m$ for all $m \in M$, where $e \in G$ is the identity.
2. **Compatibility:** $\phi(g_1, \phi(g_2, m)) = \phi(g_1 g_2, m)$ for all $g_1, g_2 \in G$ and $m \in M$.
Note: for a right action, this becomes $\phi(g_1, \phi(g_2, m)) = \phi(g_2 g_1, m)$.

It's often convenient to write $\phi_g : M \rightarrow M$, where $\phi_g(m) = \phi(g, m)$.

Example. Let $\phi : \text{SO}(3) \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be given by $\phi(R, x) = Rx$ (matrix multiplication).

Group actions are the structure we use to describe the symmetry of data defined on M .

- ▶ M could be the state or input space of a control system.
- ▶ M could be the input or output space of a regression problem.
- ▶ M could be the state or action space of a Markov decision process.

Left, Right, and Inner Automorphisms of a Lie Group

Example Any Lie group G comes equipped with a few actions out of the box:

1. The *left automorphism* is a left action of G on itself, given by $L_g(h) = g \cdot h$.
2. The *right automorphism* is a right action of G on itself, given by $R_g(h) = h \cdot g$.
3. The *inner automorphism* is a left action of G on itself, given by $I_g(h) = L_g \circ R_{g^{-1}}(h) = g \cdot h \cdot g^{-1}$.

In an *Abelian* group, the operation is commutative (i.e., $a \cdot b = b \cdot a$ for all $a, b \in G$).

In such a group, $L_g = R_g$, and thus $I_g = \text{id}_G$.

Examples. Abelian: $\text{SO}(2)$, \mathbb{R}^n , $\mathbb{R}_{\geq 0}$. Non-Abelian: $\text{SO}(3)$, $\text{SE}(2)$.

The Lie Algebra

The *Lie algebra* of any Lie group G is $\mathfrak{g} = T_e G$, the tangent space at the identity.

Crucially, every tangent space of G is given by $T_g G = d L_g(\mathfrak{g}) = d R_g(\mathfrak{g})$.

Thus, given any curve $g : (-\varepsilon, \varepsilon) \rightarrow G$, we may express its velocity via either:

1. the *body velocity*, $\xi^b(t) = d L_{g(t)^{-1}}(\dot{g}(t))$ (“in the body frame”), or
2. the *spacial velocity*, $\xi^s(t) = d R_{g(t)^{-1}}(\dot{g}(t))$ (“in the world frame”).

The *adjoint representation* is the action $\text{Ad} : G \times \mathfrak{g} \rightarrow \mathfrak{g}$ given by

$$\text{Ad}_g(\xi) = d I_g(\xi) = d L_g \circ d R_{g^{-1}}(\xi),$$

which converts between these velocities, *i.e.*, $\xi_s = \text{Ad}_g(\xi_b)$.

Properties of Group Actions

We can study properties of a group action $\phi : G \times M \rightarrow M$:

- ▶ ϕ is *free* if for all non-identity $g \in G$, the map ϕ_g has no fixed points.
- ▶ ϕ is *transitive* if, for every $m_1, m_2 \in M$, there exists $g \in G$ such that $\phi_g(m_1) = m_2$.

Example.

$$\begin{aligned}\phi : \text{SO}(3) \times \mathbb{R}^3 &\rightarrow \mathbb{R}^3, \\ \phi(R, x) &= Rx.\end{aligned}$$

✗ not free, since
 $\phi(R, 0) = R0 = 0$.

Example.

$$\begin{aligned}\phi : \text{SO}(3) \times \mathbb{S}^2 &\rightarrow \mathbb{S}^2, \\ \phi(R, x) &= Rx.\end{aligned}$$

✗ not free, since
 $\phi(R, [0, 0, 1]) = [0, 0, 1]$.

Example.

$$\begin{aligned}\phi : \text{SO}(2) \times \text{SO}(3) &\rightarrow \text{SO}(3), \\ \phi(\theta, R) &= R_z(\theta)R.\end{aligned}$$

✓ free, since at most one
column of R can be
(anti)parallel to $[0, 0, 1]$.

✗ not transitive, since
 $\|Rx\| = \|x\|$, so consider
any $\|x_1\| \neq \|x_2\|$.

✓ transitive, since we can
pick an axis of rotation
based on $x_1 \times x_2$

✗ not transitive, since three
axes of rotation needed to
achieve arbitrary orientation

Orbits of Group Actions

The *orbit* of any point $m \in M$ is given by $\phi_G(m) := \{\phi(g, m) : g \in G\}$.

Orbits form a partition of the manifold M .

- ▶ If ϕ is transitive, there is a single orbit, given by M itself!
- ▶ Define an equivalence relation on M where $m_1 \sim m_2$ whenever $m_1 \in \phi_G(m_2)$.
- ▶ This induces a *quotient space* M/G , each element of which is an orbit of ϕ .
- ▶ The quotient space M/G may be quite messy, e.g., orbits may vary in dimension!

Fact. If ϕ is free and G is compact, M/G inherits the structure of a smooth manifold. *Actually, this happens whenever ϕ is free and “proper”, but we will gloss over the details.*

Example. $\mathrm{SO}(3)/\mathrm{SO}(2) \cong \mathbb{S}^2$ for the group action $\phi(\theta, R) = R_z(\theta)R$.

Groups Made Out of Other Groups

Consider two Lie groups G and H . We can construct a new Lie group in two ways:

- ▶ The *direct product* $G \times H$ is the Lie group whose operation is performed componentwise, i.e., it is given by $(g_1, h_1) \cdot (g_2, h_2) = (g_1 g_2, h_1 h_2)$.
- ▶ Given an action $\phi : G \times H \rightarrow H$, the *semidirect product* $G \ltimes H$ is the Lie group whose operation is given by $(g_1, h_1) \cdot (g_2, h_2) = (g_1 g_2, h_1 \phi(g_1, h_2))$.

Question. What's the difference between $\mathbb{R}^3 \times \text{SO}(3)$ and $\text{SE}(3)$?

Answer. As a *manifold*, nothing! As a *Lie group*, $\text{SE}(3) = \text{SO}(3) \ltimes \mathbb{R}^3$, the *semidirect product* (vs. the direct product) using the action $\phi : (R, x) \mapsto Rx$. We verify:

$$\begin{bmatrix} R_1 & x_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} R_2 & x_2 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} R_1 R_2 & R_1 x_2 + x_1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} R_1 R_2 & x_1 + \phi(R_1, x_2) \\ 0 & 1 \end{bmatrix},$$

since the group operation for \mathbb{R}^3 is addition.

Two Common Flavors of Symmetry

Consider a map $f : M \rightarrow N$ as well as actions $\phi : G \times M \rightarrow M$ and $\psi : G \times N \rightarrow N$.

We consider two basic flavors of symmetry:

1. f is *invariant* with respect to ϕ if $f \circ \phi_g = f$ for all $g \in G$.
2. f is *equivariant* with respect to ϕ and ψ if $f \circ \phi_g = \psi_g \circ f$ for all $g \in G$.

We illustrate these ideas with the following commutative diagrams:

invariance:

$$\begin{array}{ccc} M & & \\ \phi_g \downarrow & \searrow f & \\ M & \xrightarrow{f} & N \end{array}$$

equivariance:


$$\begin{array}{ccc} M & \xrightarrow{f} & N \\ \phi_g \downarrow & & \downarrow \psi_g \\ M & \xrightarrow{f} & N \end{array}$$

Thus, invariance is just equivariance with $\psi_g = \text{id}_N$ for all $g \in G$.

Equivariance Sampler: Perception

Example. Consider the design of an edge detection filter for square grayscale images, modeled as a map $f : \mathbb{R}^{k \times k} \rightarrow \{0, 1\}^{k \times k}$, where edges are 1's and all other pixels are 0's.

Define an action $\phi : \mathcal{D}_4 \times \mathbb{R}^{k \times k} \rightarrow \mathbb{R}^{k \times k}$, where \mathcal{D}_n is the *dihedral group* consisting of all rotations and reflections that preserve the vertices of a regular n -sided polygon.



Convolutional layers have translational equivariance (shifting input shifts output).

What about rotations, reflections, etc.?

Common Approach. Design a neural architecture that explicitly enforces equivariance constraints by building networks out of *group convolutions*.

Equivariance Sampler: Reinforcement Learning

Example. Consider a Markov decision process $M = (S, A, R, \tau)$, where $R : S \times A \rightarrow \mathbb{R}$ and $\tau : S \times A \rightarrow \Delta(S)$, along with group actions $\phi : G \times S \rightarrow S$ and $\psi : G \times A \rightarrow A$.

Fact. If an MDP $M = (S, A, R, \tau)$ has:

1. equivariant transitions (i.e., $\tau(\phi_g(s), \psi_g(a)) = \tau(s, a) \circ \phi_g$) and
2. an invariant reward (i.e., $R(\phi_g(s), \psi_g(a)) = R(s, a)$),

then the optimal policy is equivariant and the value function is invariant.

Note: this is a sufficient (but not necessary) condition!

Common Approach. Constrain the value and policy networks to be equivariant using specialized architectures, then apply standard reinforcement learning (RL) algorithms.

Common Approach. Construct an *MDP homomorphism* from the original MDP “upstairs” to one of lower dimension “downstairs” (often with $\tilde{S} = S/G$). Then, learn a policy “downstairs”, and “lift” it back “upstairs” with guaranteed value equivalence.

Equivariance Sampler: Estimation and Control

Question. Is there a global error between actual and reference states $m, m_d \in M$?

On a Lie group, yes, e.g., $e = g_d^{-1}g$. But there is no inverse on a general manifold!

Let M be a *homogeneous space*, i.e., we have a transitive group action $\phi : G \times M \rightarrow M$.

Examples. Some that show up often in robotics include:

1. Any Lie group $M = G$ with $\phi_g = L_g$ or $\phi_g = R_g$.
2. $M = \mathbb{S}^n$ under the usual action of $G = \text{SO}(3)$ given by $\phi_R(x) = Rx$.

Pick an arbitrary fixed *origin* $0_M \in M$. For any reference $m_d : \mathbb{R} \rightarrow M$, a *lift* is a curve $g_d : \mathbb{R} \rightarrow G$ satisfying $\phi(g_d, 0_M) = m_d$. Then, define the error $e = \phi(g_d^{-1}, m)$.

Fact. $e = 0_M$ if and only if $m_d = m$.

Such a globally-defined error can be used to design intrinsic tracking controllers and state observers with global or almost global convergence and better linearization properties.

Equivariance Sampler: Dynamics

For a robotic system with configuration $q \in Q$ actuated by configuration-dependent forces, the dynamics take the form of the so-called “manipulator equations”

$$\mathbf{M}(q)\dot{v} + \mathbf{C}(q, v)v + \mathbf{g}(q) = \mathbf{B}(q)u,$$

where \mathbf{M} , \mathbf{C} , \mathbf{g} , and \mathbf{B} depend in general on all parts of the configuration q .

However, it's well-known that for free-floating multibody robots with configuration $q = (H, \theta) \in Q = \text{SE}(3) \times \mathbb{R}^n$ and actuated by “body-fixed” forces (e.g., thrusters, actuated joints, etc.), we may express the dynamics in the simplified form

$$\mathbf{M}(\theta)\dot{v} + \mathbf{C}(\theta, v)v + \mathbf{g}(\theta) = \mathbf{B}(\theta)u,$$

where \mathbf{M} , \mathbf{C} , and \mathbf{B} depend in particular on θ (but not on H).

More formally, this is a direct result of the fact that the mass matrix $\mathbf{M}(q)$ (and the input forces) are invariant to the action $\phi : \text{SE}(3) \times Q \rightarrow Q$ which acts by left translation on the $\text{SE}(3)$ component, so all dependence on H can be “reduced” away.

Summary

- ▶ In this background lecture, we've reviewed some basic tools with which we can describe, analyze, and ultimately exploit symmetries in robotic systems.
 - ▶ Lie groups are the basic structure of continuous symmetries.
 - ▶ Group actions allow us to describe the symmetry of data defined on other manifolds.
 - ▶ With a free group action, often invariance along the orbits is desirable.
 - ▶ With a transitive group action, we are able to obtain global error states.
- ▶ These ideas show up across numerous “verticals” in robotics.
- ▶ This overview was nowhere near complete, but I hope it provides a basis to make the rest of today's workshop more accessible to newcomers

Q U E S T I O N S ?