Learning dynamics on manifolds with neural geodesic flows

Supervisors: Ben Moseley, Luca Magri

Keywords: dynamical systems, geodesic flows, Riemannian geometry, manifold learning, neural differential equations, scientific machine learning

Background

Many physical, biological, and engineering systems evolve according to **geometric laws**. Planets follow elliptical orbits shaped by gravity, fluids swirl along curved streamlines, and biological processes often unfold on low-dimensional manifolds embedded in high-dimensional spaces. A central challenge in machine learning is to **discover and model** such dynamics faithfully and efficiently. Most existing deep learning approaches approximate trajectories in flat Euclidean spaces, often ignoring the geometric structure of the data. While expressive, these models can suffer from poor interpretability, weak generalisation, and lack of physical consistency.

To address this, our lab is developing neural geodesic flows (NGFs) [1] – a new class of machine learning models that assume dynamical systems **evolve along the geodesics** of a learned latent **Riemannian manifold**. Observations (e.g. positions of bodies in orbit, or fluid velocities) are mapped to coordinates on a latent manifold, evolved over time on the manifold according to the geodesic flow equations, and mapped back to the data space (Figure 1). Importantly, NGFs learn both the manifold metric and the mappings to and from the manifold using deep neural networks. Our core assumption – that dynamics can be described as geodesic flow on a latent manifold – may seem like a strong assumption, but multiple systems can be mathematically interpreted in this way, including classical mechanics and spacetime evolution in general relativity, with deep links to the principle of **least action**. This makes NGFs a powerful framework for physics discovery. By construction, NGFs naturally conserve geodesic energy, and their learned metric can be used to interpret the dynamical system.

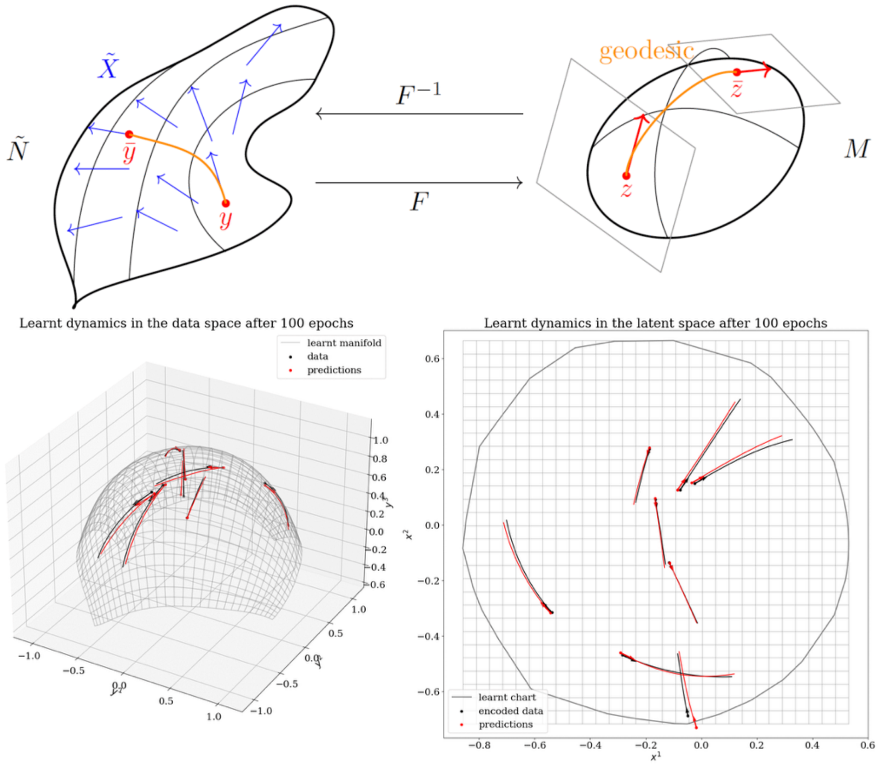

Figure 1: Top - Schematic of a neural geodesic flow (NGF). Observations of a physical system (e.g. positions of bodies in orbit, or fluid velocities), y, are mapped to coordinates, z, on a latent manifold, evolved over time according to the geodesic flow equations, and mapped back to the data space. Bottom - Example NGF modelling the dynamics of particles flowing over the surface of a sphere.

Project

Our preliminary work [1] has shown that NGFs are able to model the dynamics of simple systems, such as particles flowing over a sphere, and the two-body problem. This project will **take the significant next step** towards extensively understanding their predictive power, interpretability, and real-world applicability. We will 1) examine whether NGFs can model more complex dynamical systems (e.g. the N-body problem, the double pendulum, and fluid flows), 2) benchmark their predictive performance, efficiency, and interpretability against alternative methods such as neural ordinary differential equations, and 3) solve manifold learning and dimensionality reduction tasks on other real-world datasets, such as human activity recognition (HAR) and medical patient trajectories (PhysioNet).

Key research questions are: What types of physical systems can be faithfully represented as geodesic flows? How do NGFs compare to existing predictive methods in terms of accuracy and efficiency? Can they provide interpretable latent manifolds that reveal intrinsic dynamics? Can they be used more generally for manifold learning tasks?

For more details: please see our project page and GitHub repository.

Impact

Neural geodesic flows offer a principled way to embed geometric structure into machine learning models, allowing us to discover, interpret, and accurately model complex dynamics. Beyond physics, NGFs could transform data-driven machine learning by uncovering intrinsic low-dimensional manifolds. This project will therefore provide new tools for scientific discovery and manifold learning.

Supervisory team

Dr. **Ben Moseley** is an Assistant Professor in AI (Schmidt AI in Science Fellow) at the Department of Earth Science and Engineering and a Fellow at the Imperial I-X Centre. He heads the Scalable Scientific Machine Learning Lab and is an expert in scientific machine learning, physics-informed neural networks, neural differential equations, hybrid modelling, learned inverse algorithms, high-performance computing, geophysics, and planetary data science.

Prof. **Luca Magri** is a Professor of Scientific Machine Learning at the Department of Aeronautics at Imperial. He is an expert in scientific machine learning, real-time modelling and control, optimization, dynamical systems, chaos, data assimilation, reduced-order modelling, fluid mechanics, and quantum machine learning.

Research group

The student will be part of the **Scalable Scientific Machine Learning Lab** headed by Dr. Ben Moseley. The lab accelerates scientific research by designing scientific machine learning algorithms and applying them to impactful problems across science. See our lab website for more information.

We are a highly **cross-disciplinarity** team – we train our members across machine learning, applied mathematics, high-performance computing, and in domain-specific areas including geophysics, climate science, and planetary science. We **collaborate** with other groups at the Department of Earth Science and Engineering, I-X (Imperial's AI initiative), other Imperial departments, and with external universities and industry partners. Lab members are encouraged to present and publish at high-impact conferences and journals.

Student profile

We are looking for someone who is motivated to complete a PhD at the intersection of scientific machine learning, dynamical systems, and manifold learning, across scientific domains. This is an interdisciplinary project, and we do not expect candidates to arrive with expertise in all areas; instead, we are looking for someone with a strong technical foundation, enthusiasm for interdisciplinary research, and an ability to approach complex problems with creativity and curiosity.

Essential qualifications / experience:

- Good Master's degree in a relevant field (e.g. applied mathematics, physics, computer science, engineering, or related areas). Motivated candidates with an excellent bachelor's degree and a relevant research portfolio are encouraged to apply
- Completed courses in machine learning and/or applied mathematics
- Coding proficiency in e.g. Python/ C++/ Julia/ Fortran

Desirable qualifications / experience:

- Understanding of dynamical systems
- Some understanding of Riemannian geometry
- Experience with neural differential equations and deep learning
- Proficiency with Python machine learning frameworks (PyTorch, JAX (with Equinox))
- Experience in scientific, HPC, GPU, and/or parallel computing
- Relevant publications and/or industry experience are a plus

Funding

This project is not currently funded through a research grant and is eligible for College and/or Departmental scholarship funding. For more details on scholarship funding and deadlines see here.

Apply

If you are interested, please start by sending us:

- CV (including education, and any research experience).
- Brief motivation letter (200 – 400 words) where you should highlight how your experience enables you to pursue the project (can be in the email body).
- Any additional materials that support your application (optional).

For more details on our lab's PhD application process see here.
For more details on the Imperial PhD application process see here.

Contact

Dr. Ben Moseley: b.moseley@imperial.ac.uk

References

[1] Julian Bürge. Neural Geodesic Flows. Master's Thesis, ETH Zurich (2025).