Learnable optimisation algorithms for inverse problems

Supervisor: Ben Moseley

Keywords: inverse problems, learnable optimisation, meta learning, medical imaging, seismic imaging, scientific machine learning

Background

Being able to **infer underlying properties** of a physical system given some **observations** of the system is an essential scientific task. For example, magnetic resonance imaging (MRI) involves inferring the composition of the body from magnetic field measurements, seismic imaging involves inferring the Earth's structure from vibrations on its surface, and modelling the spread of viruses involves estimating transmission rates from numbers of confirmed cases over time.

Such problems are known as inverse problems, and they occur across many scientific domains. They are fundamentally **search** problems: find the underlying properties which produced the observations. A popular approach for solving them is to use **optimisation** algorithms, such as gradient descent. These work by defining a forward model which maps some guess of the underlying properties to predicted observations, and minimising the mismatch between the model predictions and the observations.

However, many real-world inverse problems **do not have enough information** in the observations to uniquely determine the underlying properties. For example, MRI images can suffer from low spatial resolution if low magnetic field strengths are used and determining virus transmission rates can be challenging with low reporting rates. For an inversion algorithm to perform well in this setting, it is essential to assume some **prior** knowledge of the underlying properties to **guide** the search process. Traditional methods, such as Bayesian inference, can be effective, but they are often extremely computationally expensive, and it can be difficult to hand-design prior knowledge [1].

Instead, it can be very effective to **learn** priors. Much recent work has focused on this, for example learning generative models of underlying properties [2], learning regularising terms in the optimisation objective function [3], and learning to directly estimate underlying parameters from observations [4]. Learning priors is more effective than designing them by hand, and they can be adapted to a specific inverse problem. However, it remains an open research question how best to 1) learn priors and 2) incorporate them into the search process for a given inverse problem.

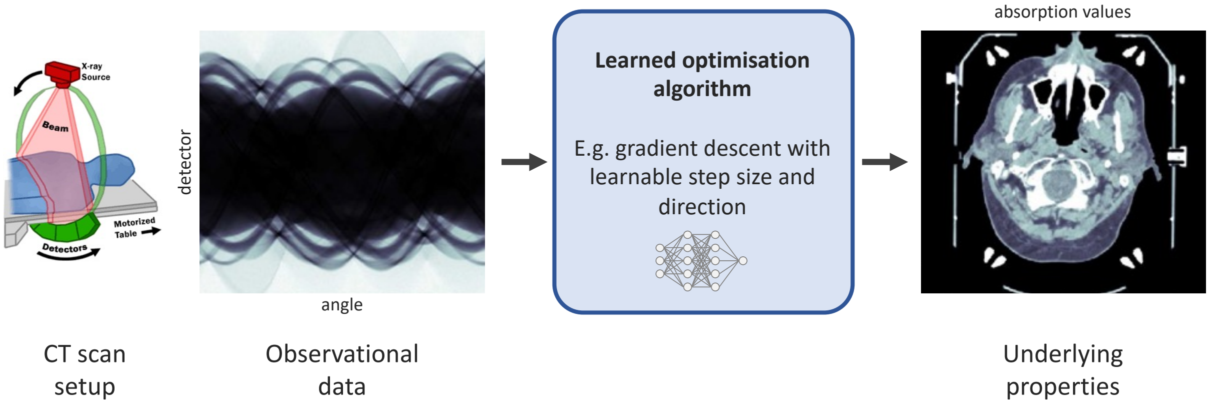

Figure 1: In this project we will learn optimisation algorithms for solving inverse problems. Given an inverse problem (computed tomography example shown), we will define an optimisation algorithm with learnable components and train it to solve many example inverse problems accurately (image credits: FDA, [5]).

Project

In this project we will investigate this question across multiple scientific domains. We will consider a range of inverse problems (including computed tomography imaging, MRI, and seismic imaging), and design and benchmark a range of **learnable optimisation algorithms** for solving them. These approaches will vary in how they learn priors and incorporate them into the search algorithm. Two specific approaches we will consider are learned gradient descent (where the step size and direction are learned, similar to [5,6]) and using diffusion models to generate guesses of underlying properties [2]. Key research questions are: what are the best ways to learn priors and incorporate them into the optimisation algorithm? How should the optimisation algorithm change with decreasing observations? Overall, the goal is to understand the effectiveness of learned optimisers for solving inverse problems.

Impact

Being able to solve inverse problems more accurately will have a transformative impact across science. It would allow researchers to extract more precise and reliable information from limited or noisy data, enabling breakthroughs in fields like medical imaging, geophysics, climate science, and materials engineering. Improved inverse problem-solving techniques would enhance our ability to reconstruct hidden structures, predict system behaviours, and design optimized solutions, accelerating innovation and discovery.

Supervision

Dr. **Ben Moseley** is a Lecturer 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.

We are open to identifying a further co-supervisor or project advisor as necessary. We actively encourage collaboration with industry and other research groups.

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 Earth science, space science, and other domains.

We value **interdisciplinarity** – 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 in scientific machine learning and inverse problems. Given the interdisciplinary nature of this project and group, we do not expect candidates with deep experience in all areas; instead, we are looking for someone who has a strong foundation, a willingness to work across disciplines, a passion for continuous learning, and the ability to approach complex problems with creativity and curiosity.

Essential qualifications / experience:

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

Desirable qualifications / experience:

- Experience in medical imaging, seismic imaging, or other inverse problem field
- Understanding of scientific machine learning, in particular learned optimisation
- Familiarity with different deep learning architectures
- 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: https://www.imperial.ac.uk/earth-science/prosp-students/phd-opportunities/funding/

Apply

If you are interested, please start by sending us a CV and detailed description (~200-400 words) of your relevant experience and specific research interests.

For more details on the Imperial PhD application process see here: https://www.imperial.ac.uk/earth-science/prosp-students/phd-opportunities/

Contact

Ben Moseley will be joining Imperial in November 2024. Before this date, please visit his personal website for contact details: https://benmoseley.blog/

References

[1] Arridge et al., (2019). Solving inverse problems using data-driven models. Acta Numerica.
[2] Song et al., (2022). Solving Inverse Problems in Medical Imaging with Score-Based Generative Models. ICLR.
[3] Lunz et al., (2018). Adversarial regularizers in inverse problems. Advances in Neural Information Processing Systems.
[4] Zhu et al., (2018). Image reconstruction by domain-transform manifold learning. Nature.
[5] Adler et al., (2017). Solving ill-posed inverse problems using iterative deep neural networks, Inverse Problems.
[6] Andrychowicz et al., (2016). Learning to learn by gradient descent by gradient descent. NeurIPS.