

Inductive Bias in Learning:

How to inject scientific knowledge in a learning algorithm to speed up scientific discovery

- Scientific Knowledge?

Systematic description of phenomena, *simple* rules, based on human understanding, that *generalizes*

- What is a **Discovery**?

Some consensus, field dependent

- How to inject scientific knowledge as bias in ML?

- Tuning, or correcting, given models from data
- Learning predictors from features that are informed by knowledge
- Formalization of Bayesian priors, defining architectures that explicitly limit hypothesis class

- General framework for knowledge and ML interaction?

- Domain expert formulates 20 questions that the ML expert should be able to answer from data
- ML expert formulates 20 questions about the problem, informing “what algorithm” to choose (with the ‘right’ inductive bias for this problem).