

Probabilistic Programming, Bayesian Nonparametrics, and Inference Compilation

FRANK WOOD¹

¹*University of Oxford, UK*

Abstract

Probabilistic programming uses programming language techniques to make it easy to denote and perform inference in the kinds of probabilistic models that inform decision-making, accelerate scientific discovery, and underlie modern attacks on the problem of artificial intelligence.

Higher order probabilistic programming languages were inspired by Bayesian nonparametric modeling. Evaluators for such languages make prototyping and performing inference in a wide variety of Bayesian nonparametric models easy and feasible.

Deep learning uses programming language techniques to automate supervised learning of program parameter values by gradient-based optimization.

What happens if we marry all these things?

This talk will review higher-order probabilistic programming and show how many Bayesian nonparametric models can be straightforwardly represented. It will also introduce inference compilation and address how linking deep learning and probabilistic programming is leading to powerful new AI techniques while also opening up significant new research questions.