Probabilistic Programming

Hongseok Yang

University of Oxford, UK
hongseok00@gmail.com

—— **Abstract** ——

Probabilistic programming refers to the idea of using standard programming constructs for specifying probabilistic models from machine learning and statistics, and employing generic inference algorithms for answering various queries on these models, such as posterior inference and estimation of model evidence. Although this idea itself is not new and was, in fact, explored by several programming-language and statistics researchers in the early 2000, it is only in the last few years that probabilistic programming has gained a large amount of attention among researchers in machine learning and programming languages, and that expressive and efficient probabilistic programming systems (such as Anglican, Church, Figaro, Infer.net, PSI, PyMC, Stan, and Venture) started to appear and have been taken up by a nontrivial number of users.

The primary goal of my talk is to introduce probabilistic programming to the CONCUR/QUEST/FORMATS audience. At the end of my talk, I want the audience to understand basic results and techniques in probabilistic programming and to feel that these results and techniques are relevant or at least related to what she or he studies, although they typically come from foreign research areas, such as machine learning and statistics. My talk will contain both technical materials and lessons that I learnt from my machine-learning colleagues in Oxford, who are developing a highly-expressive higher-order probabilistic programming language, called Anglican. It will also include my work on the denotational semantics of higher-order probabilistic programming languages and their inference algorithms, which are jointly pursued with colleagues in Cambridge, Edinburgh, Oxford and Tubingen.