1124-60-287 **Kostis Gourgoulias*** (gourgoul@math.umass.edu). *Quantification of the bias in approximate inference via information-theoretical bounds.*

A core problem in statistics and machine learning is the construction of a probability distribution, Q, that approximates another, P. For example, consider the case of approximate inference, with P representing the posterior probability of the parameters. We want to use P to carry out inference but, especially in the case of high-dimensional parameter spaces, normalizing and sampling from the posterior is inefficient. Another option is to construct a mean-field approximation to P, selected by minimizing the Kullback-Leibler divergence over a class of distributions of our choosing. Using samples generated from the optimal Q, we can then estimate quantities of interest, such as the expected value of a parameter, its variance, etc. But how much trust should we place on the accuracy of a given Q?

The goal of our project, done jointly with Markos A. Katsoulakis, Luc Rey-Bellet, and Jie Wang, is to quantify the quality of such approximations, both theoretically and practically. Given an observable f, we provide confidence bounds for the bias, $E_P[f] - E_Q[f]$, via *sharp* information-theoretic inequalities. We expect that such bounds will be applicable in a variety of situations outside of approximate inference. (Received September 12, 2016)