Scalable Gaussian Process Inference with Finite-data Mean and Variance Guarantees

Mikolaj Kasprzak (Université du Luxembourg)

Time **Friday, Sept 20, 2019 at 10:00**
Place **Campus Belval, MNO 5A (fifth floor)**

A number of Bayesian methods, including scalable Gaussian Process Inference, while appealing from the point of view of applications, lack a finite-data approximation theory and tools for evaluating their accuracy. In this talk, covering parts of the joint work with Jonathan Huggins, Trevor Campbell and Tamara Broderick (published in AISTATS 2019), I will present our novel approach to bounding posterior mean and uncertainty estimates of scalable inference algorithms. In the context of GP regression, the approach is based on a new objective called the preconditioned Fisher (pF) divergence. This objective bounds the 2-Wasserstein distance from above, which in turn provides tight bounds on the pointwise difference of the mean and variance functions. At the same time, as opposed to the Wasserstein distance, it is tractable computationally. Some of its appealing properties may be proved using probabilistic techniques involving infinite-dimensional SDEs, which I will introduce and explain. I will also mention other areas in which I suspect similar methods to be useful.