

DEPARTMENT OF STATISTICS

**Arthur Gretton**

Gatsby Computational Neuroscience Unit
Centre for Computational Statistics and Machine Learning
University College London

Kernel Tests of Goodness-of-Fit using Stein's Method

May 20, 2020
11:45am – 12:45pm

Zoom Meeting: Meeting ID: 896 0733 1426
Password: 055424

Abstract: I will describe nonparametric, kernel-based tests to assess the relative goodness of fit of models with intractable unnormalized densities. We will begin with the case of models for which the marginal densities are known in closed form, up to normalisation. In this case, we compare expectations of infinite dictionaries of features under the model and data distributions, where these expectations agree when the model and data match. The features are chosen to have zero expectation under the model, which can be achieved for unnormalised densities using the Stein trick. We will obtain improved test performance by computing a small set of Stein features, chosen so as to maximise the power of the resulting test. These features yield an interpretable picture of model and data mismatch and are computable in linear time. Finally, I will describe a test of relative goodness of fit for multiple models, where it is desired to find which model fits best, with the understanding that “all models are wrong.” This final test applies even in the case where the models contain latent variables and closed-form marginal distributions of the observed variables cannot be computed. In the case of models with low-dimensional latent structure and high-dimensional observations, our test significantly outperforms the relative maximum mean discrepancy test, which cannot exploit the latent structure.

Bio: Arthur Gretton is a Professor with the Gatsby Computational Neuroscience Unit of University College London, His recent research interests in machine learning include the design and training of generative models, both implicit (e.g. GANs) and explicit (high/infinite dimensional exponential family models), nonparametric hypothesis testing, and kernel methods. Arthur was program chair for AISTATS in 2016 (with Christian Robert), tutorials chair for ICML 2018 (with Ruslan Salakhutdinov), workshops chair for ICML 2019 (with Honglak Lee), program chair for the Dali workshop in 2019 (with Krikamol Muandet and Shakir Mohammed), co-organiser of the Machine Learning Summer School 2019 in London (with Marc Deisenroth), and is a member of Royal Statistical Society Research Section Committee.

