
Optimal Rates for the Random Fourier Feature Method*

Zoltán Szabó (Gatsby Unit, University College London)

Abstract: Kernel methods represent one of the most powerful tools in machine learning to tackle problems expressed in terms of function values and derivatives. While these methods show good versatility, they are computationally intensive and have poor scalability to large data as they require operations on Gram matrices. In order to mitigate this serious computational limitation, recently randomized methods have been proposed in the literature, which allow the application of fast linear algorithms. Random Fourier features (RFF) are among the most popular and widely applied constructions: they provide an easily computable, low-dimensional feature representation for shift-invariant kernels. Despite the popularity of RFFs, very little is understood theoretically about their approximation quality. In this talk, I am going to present the main ideas and results of a detailed finite-sample theoretical analysis about the approximation quality of RFFs by (i) establishing optimal (in terms of the RFF dimension, and growing set size) performance guarantees in uniform norm, and (ii) providing guarantees in L^r ($1 \leq r < \infty$) norms. I will also propose an RFF approximation to derivatives of a kernel with a theoretical study on its approximation quality. [Joint work with Bharath K. Sriperumbudur; equal contribution]

Bio: Zoltán Szabó is a Research Associate at the Gatsby Unit, University College London (2013 - present). He holds a double PhD in Computer Science and Applied Mathematics from the Eötvös Loránd University (2009-2012; Budapest, Hungary). His primary research interests are information theory (<https://bitbucket.org/szzoli/ite/>), statistical machine learning, empirical processes and kernel methods with applications in remote sensing (sustainability), distribution regression, structured sparsity, independent subspace analysis and its extensions, collaborative filtering.

*Carnegie Mellon University: Statistical ML Reading Group; 1 December 2015; abstract.