

The Gaussian Process Density Sampler

Ryan Prescott Adams and David J.C. MacKay
Cavendish Laboratory
University of Cambridge
Cambridge CB3 0HE, U.K.
rpa23@cam.ac.uk, mackay@mrao.cam.ac.uk

The Gaussian process is a useful prior on functions for Bayesian regression and classification. Density estimation with a Gaussian process prior has been difficult, however, due to the requirements that densities be nonnegative and integrate to unity. The statistics community has explored the use of a logistic Gaussian process for density estimation, relying on various methods of approximating the normalization constant (e.g. [1, 3]).

We propose the Gaussian Process Density Sampler (GPDS), a nonparametric, practical and consistent method of constructing a Markov chain on the properties of a posterior distribution on an unknown density, without approximation. The GPDS is composed of four parts. The first part is a GP-based prior on density functions. We develop an exchangeable procedure for generating exact samples in data space from a common density drawn from this prior. Second, we show that this prior allows practical inference of specific values of the unnormalized density, using the recently-developed technique of *exchange sampling* [2]. Third, we extend this MCMC algorithm to draw samples from the predictive distribution on data space that arises when the posterior on density functions is integrated out. This is our primary result. Finally, we demonstrate a sampling procedure for inference of the Gaussian process hyperparameters.

An extended abstract is available at
<http://www.inference.phy.cam.ac.uk/rpa23/papers/snowbird08-extended.pdf>

References

- [1] P.J. Lenk. Towards a practicable Bayesian nonparametric density estimator. *Biometrika*, 78(3):531–543, 1991.
- [2] Iain Murray, Zoubin Ghahramani, and David J.C. MacKay. MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 2006.
- [3] S.T. Tokdar and J.K. Ghosh. Posterior consistency of logistic Gaussian process priors in density estimation. *Journal of Statistical Planning and Inference*, 137:34–42, 2007.

Presenter: RPA
Category: Learning Algorithms
Oral Presentation Preferred