HW for 2019-06-04

(due: 2019-06-11)

1: Preserving positive definiteness Let $k(x, y)$ be a positive definite kernel function on some set Ω. Suppose $g : \Omega' \to \Omega$ is a one-one map and $h : \Omega' \to \mathbb{R}$ is nonzero for every $u \in \Omega'$. Argue that

$$\hat{k}(u, v) \equiv k(g(u), g(v))h(u)h(v)$$

is a positive definite kernel function on Ω'. Why are the hypotheses that g is one-one and that h is nonzero needed?

2: Sample smoothness A standard method for sampling from a multivariate normal distribution $\mathcal{N}(\mu, K)$ is to compute a Cholesky factorization $K = R^T R$ and then sample

$$Y = R^T Z + \mu$$

where $Z\mathcal{N}(0, I)$, i.e. Z is a vector whose entries are independent standard normal random variables. Use this technique to plot draws from a mean zero GP on $[-1, 1]$ using the exponential kernel and the squared exponential kernel with length scales $\ell = 0.1, 0.5, 1, 2$. Comment on the apparent smoothness of the samples.