

Supplementary Material: Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP)

Andrew Gordon Wilson
Carnegie Mellon University
andrewgw@cs.cmu.edu

Hannes Nickisch
Philips Research Hamburg
hannes@nickisch.org

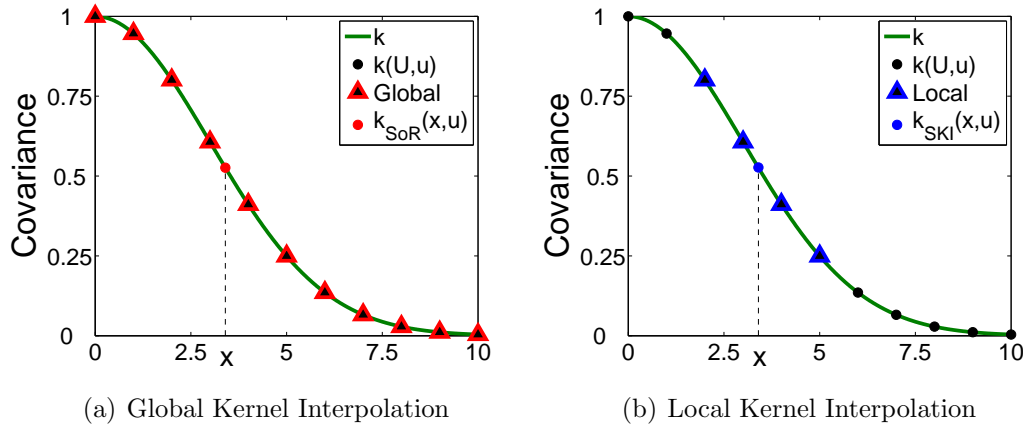


Figure 1: Global vs. local kernel interpolation. Triangle markers denote the inducing points used for interpolating $k(x, u)$ from $k(U, u)$. Here $u = 0$, $U = \{0, 1, \dots, 10\}$, and $x = 3.4$. a) All conventional inducing point methods, such as SoR or FITC, perform global GP regression on $K_{U,u}$ (a vector of covariances between all inducing points U and the point u), at test point $x_* = x$, to form an approximate \tilde{k} , e.g., $k_{\text{SoR}}(x, u) = K_{x,U} K_{U,U}^{-1} K_{U,u}$, for any desired x and u . b) SKI can perform local kernel interpolation on $K_{U,u}$ to form the approximation $k_{\text{SKI}}(x, u) = \mathbf{w}_x^\top K_{U,u}$.