

Abstract: Gaussian process (GP) regression is a Bayesian nonparametric method for regression and interpolation, offering a principled way of quantifying the uncertainties of predicted function values. For the quantified uncertainties to be well-calibrated, however, the covariance kernel of the GP prior has to be carefully selected. In this work, we theoretically compare two methods for choosing the kernel in GP regression: cross-validation and maximum likelihood estimation. Focusing on the scale-parameter estimation of a Brownian motion kernel in the noiseless setting, we prove that cross-validation can yield asymptotically well-calibrated credible intervals for a broader class of ground-truth functions than maximum likelihood estimation, suggesting an advantage of the former over the latter.