SAND2022-12368C

Sandia Academic Alliance - Fall 2022 - University of Illinois LDRD Mini-Conference

Transfer Learning of Gaussian Processes to Capture Unmodeled Physics

Erin Swansen (PhD Candidate in MechSE), Mohammad Khalil (Sandia, Quant. Modeling & Analysis), Naira Hovakimyan (UIUC, Dept of MechSE)

Introduction / Motivation

Gaussian process regression is a machine learning strategy of increasing popularity because of its ability to consolidate prior knowledge and observed data in a Bayesian manner. Calibrated Gaussian process models are able to capture unmodeled phenomena in physics-based predictive modeling and simulation. However, much like most machine learning algorithms, Gaussian process regression can become limited in applicability due to its requirement of large training data volume and computational resources. We propose to leverage a novel probabilistic transfer learning strategy whereby knowledge gained through Gaussian process modeling on similar source tasks is transferred to a novel task of interest. This learning scheme will assess the similarity of the source and target tasks, both for the priors and posteriors, and determine the optimal amount of knowledge transfer. We will apply the methodology to a vehicle control problem, whereby Gaussian process regression is employed to learn unmodeled nonlinear effects and transfer learning is utilized to alleviate the data sparsity and computational complexity challenges across vehicles of similar types of unmodeled physics.

Approach

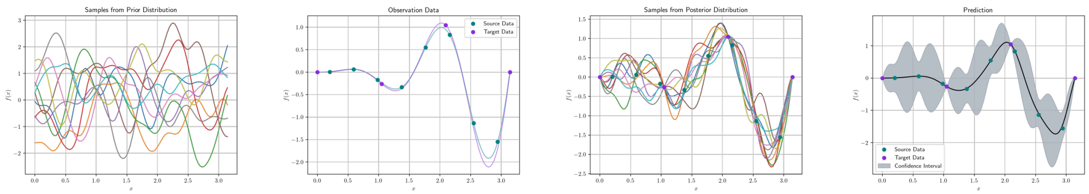

In Gaussian process regression, a prior probability distribution is combined with observed data to generate a posterior distribution. The methodology supports assimilation of "noisy" data, and we use the observation noise as a channel to diffuse knowledge from the source data. We utilize "tempering transformations" as a mechanism to control the amount of information transferred. The transferred information from available source data is then combined with that obtained from target data to make a prediction. Bayes' law yields a posterior distribution over the unknown state of interest, from which we can extract probabilistic predictions in addition to point estimates (e.g. mean process). Usually, the variance is larger near source data points than target data points, indicating less reliance on the source data. To evaluate the quality of the prediction, we evaluate the posterior predictive at the ground truth, loosely representing the likelihood that the truth could be generated from the calibrated model.

Prior

$f(x) \sim GP\big(m(x), k(x, x')\big)$

$m(x) = \mathbb{E}[f(x)]$

$k(x, x') = \mathbb{E}[\big(f(x) - m(x)\big)\big(f(x) - m(x)\big)^T]$

Bayesian Update

$$p(f|\mathcal{D}) = \frac{p(\mathcal{D}|f)p(f)}{p(\mathcal{D})}$$

Current Status / Results

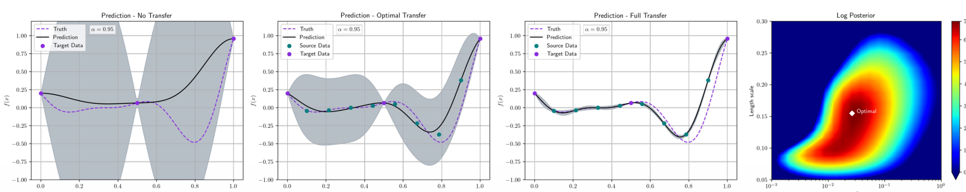

Problem Statement

We have obtained 8 observations from a source function $f(x) = (6x - 2)^2 \sin(12x - 4)$. We would like to use Gaussian process regression and transfer learning to predict a target function $f'(x) = (6x - 2)^2 \sin(12\alpha x - 4)$, but with sparse data of only 3 observations. The tempering transformation is applied via a parameter, β, which scales the source observation noise matrix. β and the length scale hyperparameter of the GP, ℓ, are optimized simultaneously to determine optimal transfer.

Results

The optimal parameters β and ℓ produce a solution for which the true function lies within the confidence interval, despite a biased mean prediction. No transfer results in poor prediction with large variance and bias. Full transfer causes an overly confident prediction with small uncertainty that does not predict the true response with any practical level of confidence.

Challenges

- Computational challenges with non-invertible matrices. When there are two observations at the same point and β is relatively small, the GP kernel matrix is rank deficient. Similarly, noiseless observations lead to a non-invertible covariance matrix for the posterior distribution. This affects log posterior calculations for optimization tasks.
- In practice, we will need to evaluate the quality of the prediction using only sparse observation data, as we will not have knowledge of the true function. Data sparsity prohibits division of the observations into training and testing sets, so other methods need to be explored.

Next Steps / Future Work

- Explore methods for dealing with poorly conditioned GP matrices.
- Compare alternative solutions to log posterior, including marginal log likelihood and leave-one-out cross validation.
- Consider other approaches for hyperparameter optimization. Determine whether optimization should occur before transfer, i.e. optimized for source data, or during transfer.
- Use Gaussian process regression to learn unmodeled vehicle dynamics. Then, employ transfer learning with a vehicle with similar dynamics. Evaluate changes in performance and safety metrics.

ILLINOIS

Sandia National Laboratories