Posterior Consistency for Gaussian Process Surrogate Models with Generalized Observations

Rujian Chen, John W. Fisher III

Massachusetts Institute of Technology

Summary of the main result

Gaussian processes are often used as approximations for complex computational models. We analyze the consistency of parameter inference in large system models with GP surrogate components, under the setting of generalized observations. We provide guarantees and conditions which apply to many practical model setups.

Parameter inference with GP surrogate models

- GP surrogate models approximate intractable *black box functions* with GPs trained on simulation data (Fig. a). The surrogate model may be used to infer unknown parameters of interest (Fig. b):

Parameters: $\theta \sim p(\theta)$

Exact model: Black box system function: f_*

Measurements: $y = f_*(\theta) + \xi$

Posterior: $p_*(\theta \mid y) \propto p(\theta) p(y \mid \theta, f_*)$ (1)

Approx model: GP prior: $f \sim \mathcal{GP}(0, k)$

Simulation model: $D = \{(\bar\theta_j, \bar y_j)\}_{j=1}^{|D|}, \ \bar y_j = f(\bar\theta_j) + \epsilon_j$

Surrogate function model: $f \mid D$

Measurements: $y = f(\theta) + \xi$

Posterior: $p(\theta \mid y, D) \propto \mathbb{E}_f\, p(\theta) p(D \mid f) p(y \mid \theta, f)$ (2)

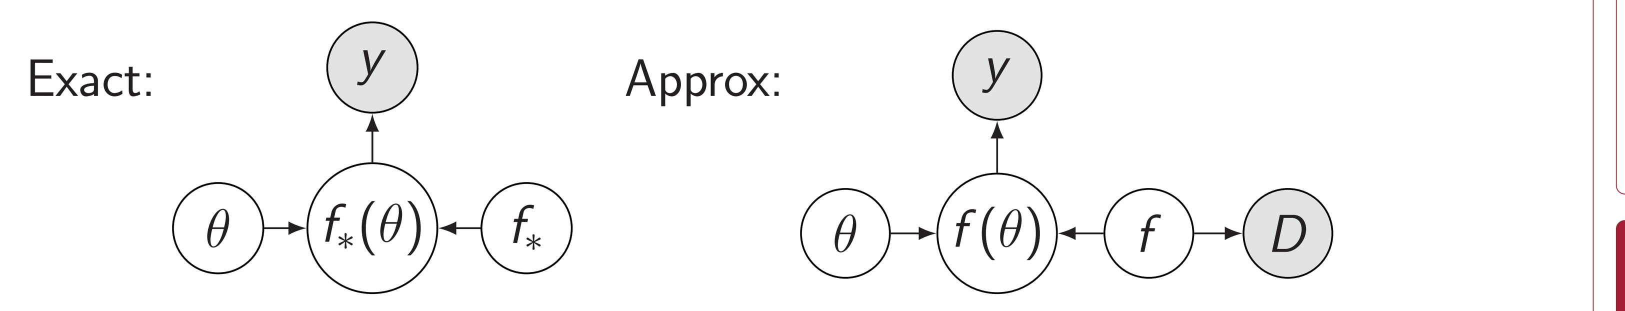

- We provide conditions under which the approximate posterior (Eq. 2) converges to the exact posterior (Eq. 1):

$$p(\theta \mid y, D_n) \to p_*(\theta \mid y) \ \text{ as } n \to \infty.$$

Generalized observations

Generalized simulator and real observations, $\bar y$ and y, are values of linear functionals λ_j, L_θ acting on f:

$$D = \{(\lambda_j, \bar y_j)\}_{j=1}^{|D|}, \quad \bar y_j = \lambda_j(f) + \epsilon, \quad y = L_\theta(f) + \xi$$

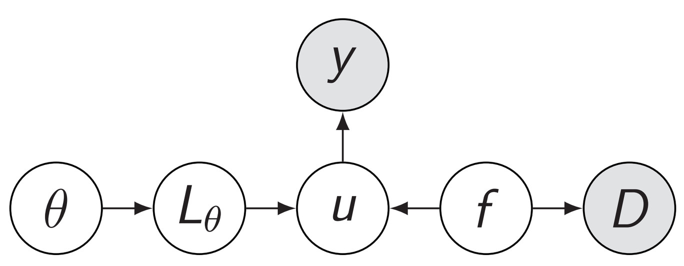

- Generalized observations include direct observations (the evaluation functional), derivatives, integrals, linear combinations, etc (Fig. c).
- They arise from various sensor types and simulation methods.
- $\boldsymbol{\Theta} \equiv \text{Dom}(\theta)$ can be different from $\mathcal{X} \equiv \text{Dom}(f)$ to represent model structure.
- λ_j's can be different types from L_θ's.

While generalized observations allow for more diverse sensor types, additional technical conditions are necessary to guarantee posterior consistency.

Posterior Consistency

- We show consistency of the approximate posterior under the generalized setting:

Posterior expectation: $\int g(\theta) p(\theta \mid y, D_n)\, d\theta \xrightarrow{p} \int g(\theta) p_*(\theta \mid y)\, d\theta$

KL-divergence: $\text{KL}[p(\cdot \mid y, D_n) \,\|\, p_*(\cdot \mid y)] \xrightarrow{p} 0$

Entropy: $\text{H}[p(\cdot \mid y, D_n)] \xrightarrow{p} \text{H}[p_*(\cdot \mid y)]$

- Assumptions (sketch):
  1. Regularity conditions: $f_* \in H_k$ and $L_\theta, \lambda_{nj} \in H_k^*$, where H_k is the RKHS of k and H_k^* its dual space.
  2. Kernel k satisfies smoothness property $\partial/\partial x_i \in H_k^*$.
  3. Fraction of evaluation functionals in D_n does not diminish anywhere in Dom(f).
  4. Independent Gaussian simulation noise ϵ_{nj} with upper bounded variance.

Illustrations

- Sufficient condition on f_* and $L_{\theta m}$ for consistency:

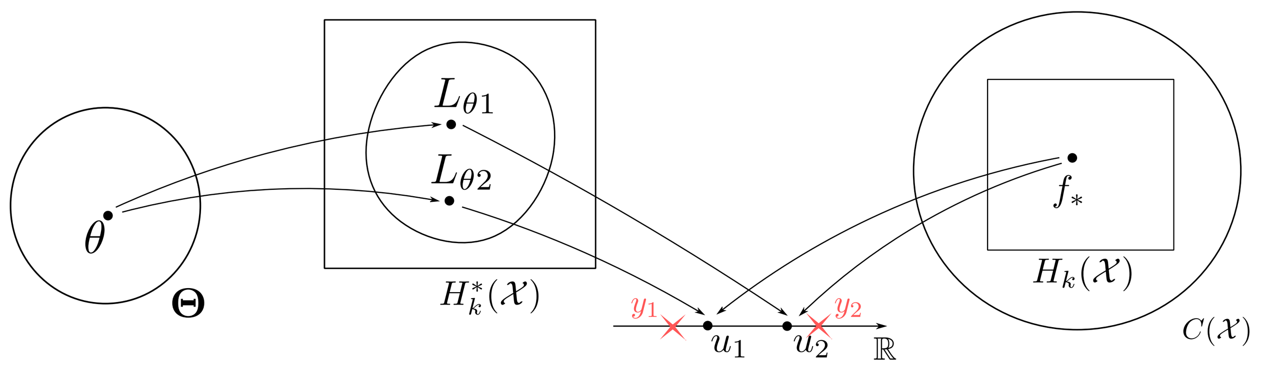

- Example: Posterior inference with generalized observations:

$$L_{\theta 1}(f) = f(\theta), \quad L_{\theta 2}(f) = \frac{df}{dx}(\theta)$$

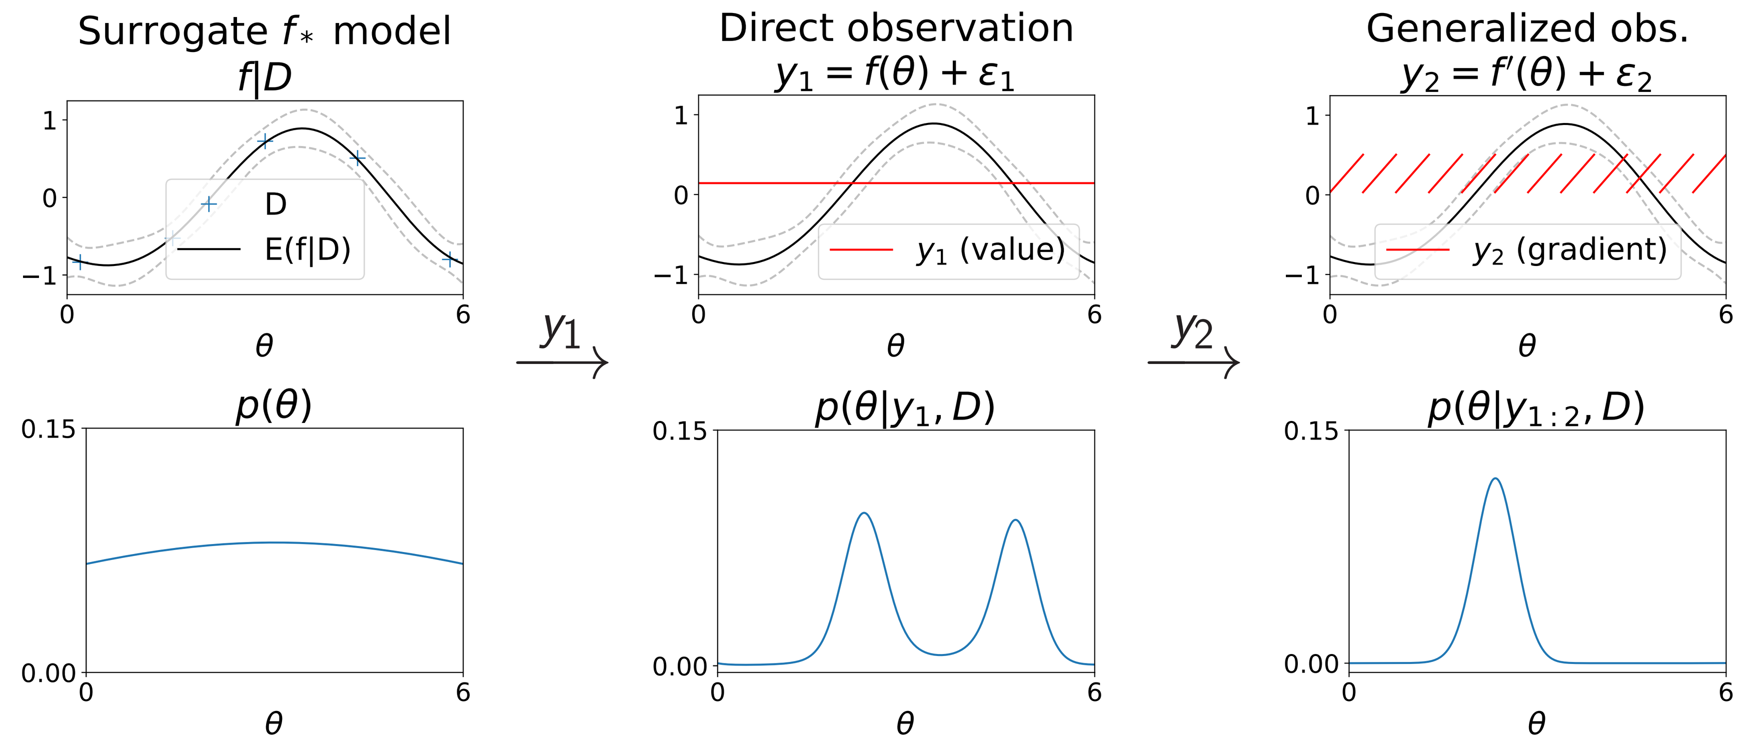

(a) (b) (c)

What happens if some assumptions are violated?

- An example where consistency does not hold if $L_\theta \in H_k^*$ is violated:

$$k = \text{Matern}_{3/2}, \ L_\theta(f) = \frac{\partial^2 f}{\partial x^2}(\theta), \ \text{then } L_\theta \notin H_k^*.$$

  - The following simulation shows $\mathbb{E}_{D_n}[\mathbb{E}_f[L_\theta(f) \mid D_n] - L_\theta(f_*)]^2$ diverges as $n \to \infty$, leading to failure of pointwise posterior consistency (cf. Proposition 3.1).
  - L_θ is the pointwise limit of $L_i \in H_k^*$, where $L_i(f) = \frac{1}{\epsilon_i}\left(\frac{\partial f}{\partial x}(\theta + \epsilon_i) - \frac{\partial f}{\partial x}(\theta)\right)$, as $\epsilon_i \to 0$. However, $\|L_i\|_{H_k^*} \to \infty$. The following simulation shows $\mathbb{E}_{D_n}[\mathbb{E}_f[L_i(f) \mid D_n] - L_i(f_*)]^2 \to 0$ but increasingly slowly.
- If assumptions are satisfied, e.g. $k = \text{RBF}$, then $L_\theta \in H_k^*$. Then the convergence of Prop. 3.1 holds.

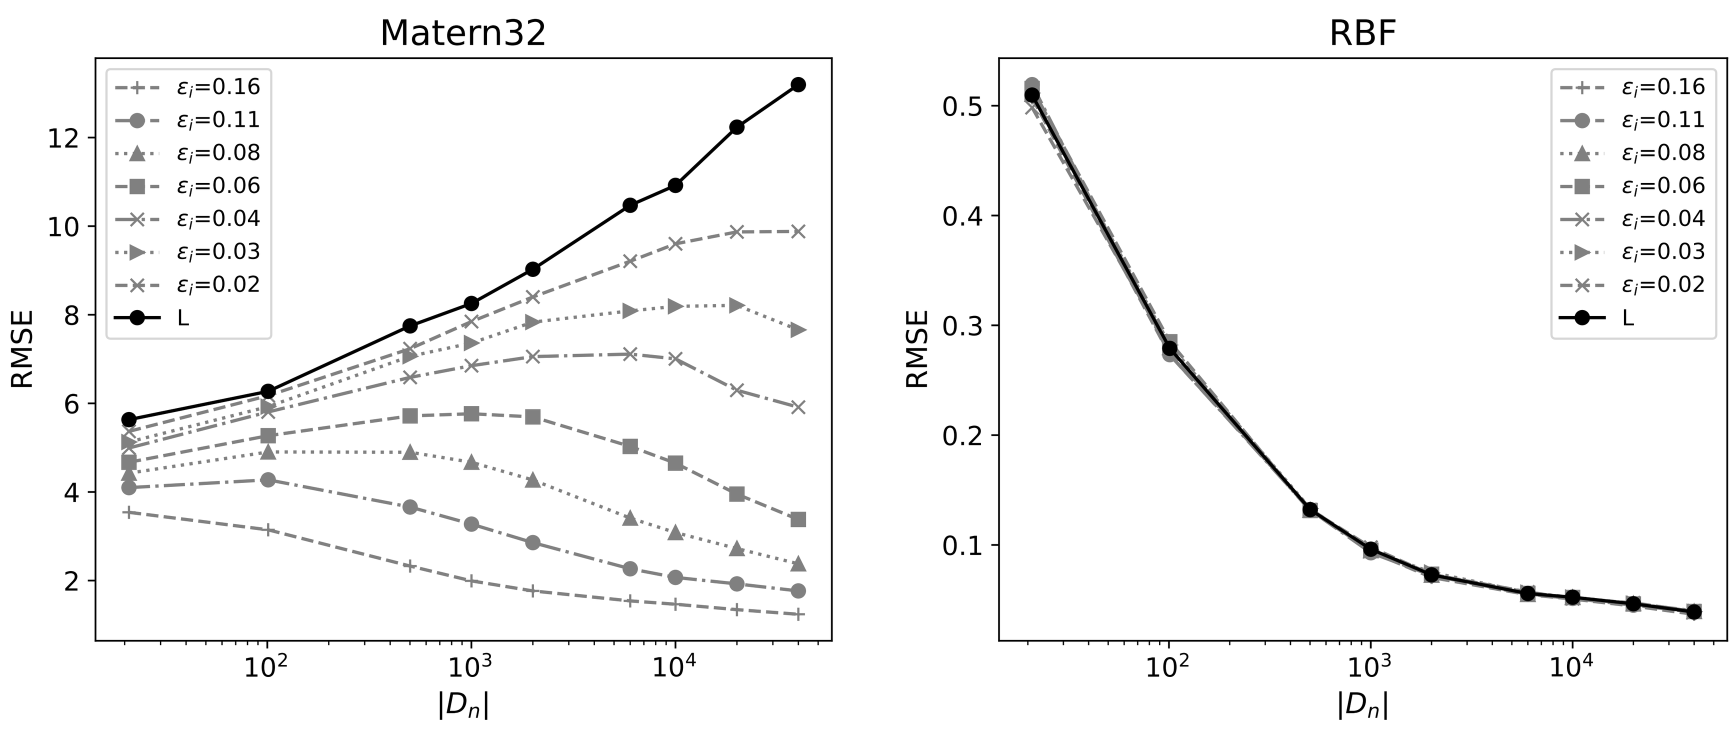

Multiple components with additional structure

Posterior consistency can be extended to a larger class of models with multiple black box components and more complex structure:

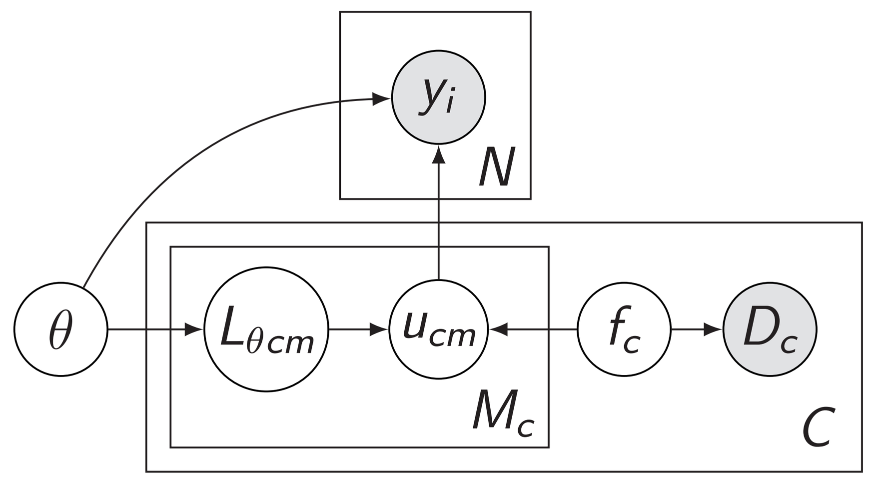

- Each component c is modeled by a GP surrogate f_c.
- Each c generates multiple responses $L_{\theta cm}(f_c)$, $m = 1, \ldots, M_c$.
- y_i measures linearly combined responses across components and/or θ.

Emails: {rjchen,fisher}@csail.mit.edu