
Leveraging Locality and Robustness to Achieve Massively Scalable Gaussian Process Regression

Anonymous Author(s)

Affiliation

Address

email

365 A Theoretical GPnn Results

366 A.1 Preliminary results

367 Let $\rho(\mathbf{x}, \mathbf{x}') = \sigma_f \sqrt{1 - c(\mathbf{x}/l, \mathbf{x}'/l)}$ be the kernel-induced distance function over \mathbb{R}^d ([17]). We
368 define $\mathbf{x}_{(j,n)}(\mathbf{x}^*)$ as the j^{th} nearest-neighbour random variable to a test point \mathbf{x}^* under ρ , which we
369 abbreviate to $\mathbf{x}_{(j)}$ when the context is clear, and $\mathbf{x}_{(j)}(\mathbf{x}^*) \in N_m(\mathbf{x}^*)$ as the realised j^{th} nearest-
370 neighbour of the test point \mathbf{x}^* from a training set X . From this we define $\epsilon_i = \rho^2(\mathbf{x}_{(i)}, \mathbf{x}^*)$ and
371 $\epsilon_{ij} = \rho^2(\mathbf{x}_{(i)}, \mathbf{x}_{(j)})$.

372 **Definition 8** (Support). Let $P_{\mathbf{x}}$ be the probability measure of \mathbf{x} and $S_{\mathbf{x},\epsilon}^\rho$ the closed ball of radius
373 $\epsilon > 0$ under the metric ρ centred at \mathbf{x} . Then we define $\text{support}(P_{\mathbf{x}}) = \{\mathbf{x} : P_{\mathbf{x}}(S_{\mathbf{x},\epsilon}^\rho) > 0 \forall \epsilon > 0\}$.

374 **Definition 9** (Weakly-faithful). We define a pair of metrics $\rho(\cdot, \cdot), \hat{\rho}(\cdot, \cdot)$ to be weakly-faithful w.r.t.
375 each other if the following condition holds: The m^{th} nearest-neighbour under $\hat{\rho}$ converges to the test
376 point as $n \rightarrow \infty$ if and only if the m^{th} nearest-neighbour under ρ converges to the test point in the
377 limit.

378 Assumptions

379 **(A1)** $\mathbf{x} \stackrel{\text{iid}}{\sim} P_{\mathbf{x}}$ and $\mathbf{x}^* \in \text{support}(P_{\mathbf{x}})$ under the generative metric defined by $c(\cdot, \cdot)$.

380 **(A2)** $c(\cdot, \cdot), \hat{c}(\cdot, \cdot)$ are stationary kernels whose induced distance functions are *weakly faithful*
381 metrics (Definition 9).

382 **(A3)** $y_i = f(\mathbf{x}_i) + \xi_i$ with $\xi_i \stackrel{\text{iid}}{\sim} P_\xi$, $f(\mathbf{x}) \sim \mathcal{WSRF}(\sigma_f^2 c(\cdot/l, \cdot/l))$ and $y_i | f(\mathbf{x}_i) \sim P_\xi$ and
383 $\mathbb{E}[\xi] = 0, \mathbb{E}[\xi^2] = \sigma_\xi^2$.

384 **Note:** Assumption (A2) is not overly restrictive and encompasses commonly used kernels such as all
385 those mentioned in this paper.

386 **Lemma 10.** $\epsilon_i \rightarrow 0$ and $\epsilon_{ij} \rightarrow 0$ as $n \rightarrow \infty$ a.e. with respect to the measure over $\mathbf{x} \in \mathbb{R}^d, P_{\mathbf{x}}$, for
387 $i, j \leq m, \frac{m}{n} \rightarrow 0$ and under (A1-2).

388 *Proof.* Lemma 6.1 of [9] states that $\|\mathbf{x}_{(m,n)}(\mathbf{x}) - \mathbf{x}\| \xrightarrow{n \rightarrow \infty} 0$ with probability one (with respect
389 to $P_{\mathbf{x}}$). Their proof can be generalised immediately to state that $\rho(\mathbf{x}_{(m,n)}(\mathbf{x}), \mathbf{x}) \xrightarrow{n \rightarrow \infty} 0$ by using
390 our definition of support, 8, that directly invokes the metric ρ . Hence $\epsilon_i \rightarrow 0$ for all $i \leq m$ (since
391 \mathbf{x}^* is in $\text{support}(P_{\mathbf{x}})$). Since ρ is a metric it satisfies the triangle inequality; hence $\rho(\mathbf{x}_{(i)}, \mathbf{x}_{(j)}) \leq$
392 $\rho(\mathbf{x}_{(i)}, \mathbf{x}^*) + \rho(\mathbf{x}_{(j)}, \mathbf{x}^*) \xrightarrow{n \rightarrow \infty} 0$ for all $i, j \leq m$. \square

393 **Lemma 11.** For an m -GPnn under the assumptions (A1-3),

$$\lim_{n \rightarrow \infty} \mathbf{k}_N^{*T} K_N^{-1} \mathbf{k}_N^* = \sigma_f^2 - \sigma_\xi^2 m^{-1} + \mathcal{O}(m^{-2}).$$

394 *Proof.* From Lemma 10 we have that $\lim_{n \rightarrow \infty} k(\mathbf{x}_{(j)}(\mathbf{x}^*), \mathbf{x}^*) = \lim_{n \rightarrow \infty} (\sigma_f^2 - \epsilon_i) = \sigma_f^2$ and
 395 $\lim_{n \rightarrow \infty} k(\mathbf{x}_{(i)}(\mathbf{x}^*), \mathbf{x}_{(j)}(\mathbf{x}^*)) = \lim_{n \rightarrow \infty} (\sigma_f^2 - \epsilon_{ij}) = \sigma_f^2$. As a result, $\mathbf{k}_N^* \rightarrow \sigma_f^2 \mathbf{1}$ and

$$K^\infty := \lim_{n \rightarrow \infty} K_N = \sigma_\xi^2 I + \sigma_f^2 \mathbf{1}\mathbf{1}^T. \quad (8)$$

396 Now using Sherman-Morrison and the continuity of matrix inverse and matrix-matrix products:

$$(A + \mathbf{b}\mathbf{c}^T)^{-1} = A^{-1} - \frac{A^{-1}\mathbf{b}\mathbf{c}^T A^{-1}}{1 + \mathbf{c}^T A^{-1}\mathbf{b}} \quad (9)$$

$$(K^\infty)^{-1} = (\sigma_\xi^2 I + \sigma_f^2 \mathbf{1}\mathbf{1}^T)^{-1} = \frac{1}{\sigma_\xi^2} \left(I - \sigma_f^2 \frac{\mathbf{1}\mathbf{1}^T}{\sigma_\xi^2 + \sigma_f^2 \mathbf{1}^T \mathbf{1}} \right) \quad (10)$$

$$\begin{aligned} \mathbf{1}^T (K^\infty)^{-1} \mathbf{1} &= \frac{m}{\sigma_\xi^2} \left(1 - \frac{m\sigma_f^2}{\sigma_\xi^2 + m\sigma_f^2} \right) \\ &= \frac{m}{\sigma_\xi^2} \left(1 - m\sigma_f^2 \frac{1}{m\sigma_f^2} \left(1 - \frac{\sigma_\xi^2}{m\sigma_f^2} + \frac{\sigma_\xi^4}{m^2\sigma_f^4} - \mathcal{O}(m^{-3}) \right) \right) \\ &= \frac{1}{\sigma_f^2} - \frac{\sigma_\xi^2}{m\sigma_f^4} + \mathcal{O}(m^{-2}). \end{aligned} \quad (11)$$

397 Thus,

$$\lim_{n \rightarrow \infty} \mathbf{k}_N^{*T} K_N^{-1} \mathbf{k}_N^* = \sigma_f^4 \mathbf{1}^T (K^\infty)^{-1} \mathbf{1} = \sigma_f^2 - \sigma_\xi^2 m^{-1} + \mathcal{O}(m^{-2}). \quad (12)$$

398

□

399 **Lemma 12** (\mathcal{WSRF} expectations). *Under (A3), $\mathbb{E}_{\mathbf{y}, y^*} \{\mathbf{y}y^*\} = \mathbf{k}^*$ and $\mathbb{E}_{\mathbf{y}} \{\mathbf{y}\mathbf{y}^T\} = K$.*

400 *Proof.* By assumption on the covariance properties of y and the independence and zero-mean of
 401 the additive noise, $\mathbb{E}_{\mathbf{y}} \{y_i y_j\} = k(\mathbf{x}_i, \mathbf{x}_j)$. Extending this to the joint distribution over \mathbf{y}, y^* is
 402 straightforward and gives the results stated. □

403 Lemma 12 is subsequently assumed to be in use throughout A.2.

404 A.2 Limit proofs

405 In the following statements only misspecification of type (d) and/or (e) (subsection 4.2) is considered
 406 to be at work.

407 **Lemma 13** (MSE limit). *Under the assumptions (A1-3), for fixed $m < \infty$, the predictive GPnn given
 408 in subsection 4.1 converges pointwise in the sense of MSE wrt $P_{\mathbf{x}}$ -a.e. as*

$$\lim_{n \rightarrow \infty} f_n^{\text{MSE}}(\boldsymbol{\theta}) = \sigma_\xi^2 (1 + m^{-1}) - \mathcal{O}(m^{-2}).$$

409 *Proof.* This follows from Lemma 11 by expanding the definition of MSE:

$$\begin{aligned} \lim_{n \rightarrow \infty} f_n^{\text{MSE}}(\boldsymbol{\theta}) &= \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{y}, y^*} \left\{ |y^* - \mu_N^*|^2 \right\} \\ &= \lim_{n \rightarrow \infty} \left[\mathbb{E}_{y^*} \{y^{*2}\} + \mathbb{E}_{\mathbf{y}} \{\mu_N^{*2}\} - 2 \mathbb{E}_{\mathbf{y}, y^*} \{\mathbf{k}_N^{*T} K_N^{-1} \mathbf{y}_N y^*\} \right] \\ &= \sigma_f^2 + \sigma_\xi^2 - \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{y}} \{\mu_N^{*2}\} \\ &= \sigma_\xi^2 (1 + m^{-1}) - \mathcal{O}(m^{-2}). \end{aligned}$$

410 Since $\mathbb{E}_{\mathbf{y}} \{\mu_N^{*2}\} = \mathbb{E}_{\mathbf{y}} \{\mathbf{k}_N^{*T} K_N^{-1} \mathbf{y}_N \mathbf{y}_N^T K_N^{-1} \mathbf{k}_N^*\} = \mathbf{k}_N^{*T} K_N^{-1} \mathbf{k}_N^*$, and by assumption
 411 $\mathbb{E}_{\mathbf{y}, y^*} \{\mathbf{y}_N y^*\} = \mathbf{k}_N^*$, even under a \mathcal{WSRF} generative model (Lemma 12). □

Corollary 14 (NLL limit).

$$\lim_{n \rightarrow \infty} f_n^{\text{NLL}}(\boldsymbol{\theta}) = \frac{1}{2} \log(\sigma_\xi^2 (1 + m^{-1})) + \frac{1}{2} + \frac{1}{2} \log 2\pi - \mathcal{O}(m^{-2}).$$

412 *Proof.* The proof follows straightforwardly from Lemma 11 and because $\sigma_N^{*2} = \sigma_f^2 + \sigma_\xi^2 -$
 413 $\mathbf{k}_N^{*T} K_N^{-1} \mathbf{k}_N^*$.

$$\begin{aligned}
 2 \mathbb{E}_{\mathbf{y}, y^*} \{l_N^*\} &= \mathbb{E}_{\mathbf{y}, y^*} \left\{ \log \sigma_N^{*2} + \frac{(y^* - \mu_N^*)^2}{\sigma_N^{*2}} + \log 2\pi \right\} \\
 &= \log \sigma_N^{*2} + 1 + \log 2\pi \\
 \lim_{n \rightarrow \infty} 2 \mathbb{E}_{\mathbf{y}, y^*} \{l_N^*\} &= \log(\sigma_f^2 + \sigma_\xi^2 - (\sigma_f^2 - \sigma_\xi^2 m^{-1} + \mathcal{O}(m^{-2}))) + 1 + \log 2\pi \\
 &= \log(\sigma_\xi^2(1 + m^{-1}) - \mathcal{O}(m^{-2})) + 1 + \log 2\pi \\
 &= \log \sigma_\xi^2 + m^{-1} + 1 + \log 2\pi - \mathcal{O}(m^{-2}).
 \end{aligned}$$

414

□

415 A.2.1 Full misspecification

416 For the remainder of A.2 we assume that the full range of possible misspecifications ((a)-(e)) outlined
 417 in subsection 4.2 are in action. We refer to this case as “fully-misspecified” and introduce the notation
 418 $\hat{\mu}_N^*, \hat{\sigma}_N^{*2}$ to be understood to mean the predictive mean and variance under these misspecifications.

419 **Lemma 15** (Fully misspecified MSE limit). *For a fully misspecified model, asymptotically*

$$\lim_{n \rightarrow \infty} f_n^{\text{MSE}}(\hat{\theta}) = \sigma_\xi^2(1 + m^{-1}) - \mathcal{O}(m^{-2}).$$

420 *provided the misspecified kernel distance metric is weakly faithful in the sense that the m^{th} nearest-*
 421 *neighbour converges under both the true and misspecified metrics (Definition 9).*

Proof.

$$\begin{aligned}
 \mathbb{E}_{\mathbf{y}} \left\{ \mathbb{E}_{\mathbf{y}^*} [(y^* - \hat{\mu}_N^*)^2 | \mathbf{y}] \right\} &= \mathbb{E}_{\mathbf{y}} \left\{ \mathbb{E}_{\mathbf{y}^*} [y^{*2} - 2y^* \hat{\mu}_N^* + (\hat{\mu}_N^*)^2 | \mathbf{y}] \right\} \\
 &= \mathbb{E}_{\mathbf{y}} \left\{ \sigma_N^{*2} + \mu_N^{*2} - 2\mu_N^* \hat{\mu}_N^* + (\hat{\mu}_N^*)^2 \right\} \\
 &= \underbrace{\sigma_N^{*2}}_{(a)} + \underbrace{\mathbf{k}_N^{*T} K_N^{-1} \mathbf{k}_N^*}_{(b)} - 2 \underbrace{\mathbf{k}_N^{*T} \hat{K}_N^{-1} \hat{\mathbf{k}}_N^*}_{(c)} + \underbrace{\hat{\mathbf{k}}_N^{*T} \hat{K}_N^{-1} K_N \hat{K}_N^{-1} \hat{\mathbf{k}}_N^*}_{(d)}.
 \end{aligned}$$

422 We can use standard results to state that (a) + (b) = $\sigma_f^2 + \sigma_\xi^2$. Then we define $\hat{\gamma} = \frac{\hat{\sigma}_f^2}{\hat{\sigma}_\xi^2 + m\hat{\sigma}_f^2}$ and
 423 expand it in terms of m^{-1} :

$$1 - m\hat{\gamma} = \frac{\hat{\sigma}_\xi^2}{m\hat{\sigma}_f^2} - \frac{\hat{\sigma}_\xi^4}{m^2\hat{\sigma}_f^4} + \mathcal{O}(m^{-3}).$$

424 In a manner similar to Lemma 11 we use this result to compute:

$$\begin{aligned}
 \lim_{n \rightarrow \infty} (c) &= \sigma_f^2 \mathbf{1}^T \hat{\sigma}_\xi^{-2} (I - \hat{\gamma} \mathbf{1} \mathbf{1}^T) \mathbf{1} \hat{\sigma}_f^2 \\
 &= \frac{\sigma_f^2 \hat{\sigma}_f^2}{\hat{\sigma}_\xi^2} m(1 - m\hat{\gamma}) \\
 &= \frac{\sigma_f^2 \hat{\sigma}_f^2}{\hat{\sigma}_\xi^2} \left(\frac{\hat{\sigma}_\xi^2}{\hat{\sigma}_f^2} - \frac{\hat{\sigma}_\xi^4}{m\hat{\sigma}_f^4} \right) + \mathcal{O}(m^{-2}) \\
 &= \sigma_f^2 - \frac{\sigma_f^2 \hat{\sigma}_\xi^2}{m\hat{\sigma}_f^2} + \mathcal{O}(m^{-2})
 \end{aligned}$$

425 and

$$\begin{aligned}
\lim_{n \rightarrow \infty} (d) &= \frac{\hat{\sigma}_f^4}{\hat{\sigma}_\xi^4} \mathbf{1}^T (I - \hat{\gamma} \mathbf{1} \mathbf{1}^T) (\sigma_\xi^2 I + \sigma_f^2 \mathbf{1} \mathbf{1}^T) (I - \hat{\gamma} \mathbf{1} \mathbf{1}^T) \mathbf{1} \\
&= \frac{\hat{\sigma}_f^4}{\hat{\sigma}_\xi^4} \mathbf{1}^T [\sigma_\xi^2 I + \sigma_f^2 \mathbf{1} \mathbf{1}^T - 2\sigma_\xi^2 \hat{\gamma} \mathbf{1} \mathbf{1}^T + \hat{\gamma}^2 \sigma_\xi^2 m \mathbf{1} \mathbf{1}^T - 2\sigma_f^2 \hat{\gamma} m \mathbf{1} \mathbf{1}^T + \sigma_f^2 \hat{\gamma}^2 m^2 \mathbf{1} \mathbf{1}^T] \mathbf{1} \\
&= \frac{\hat{\sigma}_f^4}{\hat{\sigma}_\xi^4} m (\sigma_\xi^2 + m \sigma_f^2) [1 - 2m\hat{\gamma} + m^2 \hat{\gamma}^2] \\
&= \frac{\hat{\sigma}_f^4}{\hat{\sigma}_\xi^4} m (\sigma_\xi^2 + m \sigma_f^2) (1 - m\hat{\gamma})^2 \\
&= \frac{\hat{\sigma}_f^4}{\hat{\sigma}_\xi^4} m (\sigma_\xi^2 + m \sigma_f^2) \left(\frac{\hat{\sigma}_\xi^4}{m^2 \hat{\sigma}_f^4} - 2 \frac{\hat{\sigma}_\xi^6}{m^3 \hat{\sigma}_f^6} + \mathcal{O}(m^{-4}) \right) \\
&= \sigma_f^2 + \frac{\sigma_\xi^2}{m} - 2 \frac{\sigma_f^2 \hat{\sigma}_\xi^2}{\hat{\sigma}_f^2 m} - \mathcal{O}(m^{-2}),
\end{aligned}$$

426 where we have used the expansion of $1 - m\hat{\gamma}$ given earlier. Putting these results together gives

$$\begin{aligned}
\lim_{n \rightarrow \infty} f_n^{\text{MSE}}(\hat{\boldsymbol{\theta}}) &= \lim_{n \rightarrow \infty} [(a) + (b) - 2(c) + (d)] \\
&= \sigma_f^2 + \sigma_\xi^2 - 2 \left(\sigma_f^2 - \frac{\sigma_f^2 \hat{\sigma}_\xi^2}{m \hat{\sigma}_f^2} \right) + \sigma_f^2 + \frac{\sigma_\xi^2}{m} - 2 \frac{\sigma_f^2 \hat{\sigma}_\xi^2}{\hat{\sigma}_f^2 m} - \mathcal{O}(m^{-2}) \\
&= \sigma_\xi^2 (1 + m^{-1}) - \mathcal{O}(m^{-2}).
\end{aligned}$$

427

□

Lemma 16 (Calibration limit under full misspecification).

$$\lim_{n \rightarrow \infty} f_n^{\text{CAL}}(\hat{\boldsymbol{\theta}}) = \frac{\sigma_\xi^2}{\hat{\sigma}_\xi^2} + \mathcal{O}(m^{-2}).$$

428 *Proof.* We use continuity to write

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{y}, \mathbf{y}^*} \left\{ \frac{(y^* - \hat{\mu}_N^*)^2}{\hat{\sigma}_N^{*2}} \right\} = \left(\lim_{n \rightarrow \infty} \frac{1}{\hat{\sigma}_N^{*2}} \right) \left(\lim_{n \rightarrow \infty} f_n^{\text{MSE}}(\hat{\boldsymbol{\theta}}) \right).$$

429 By direct application of Lemma 11 $\hat{\sigma}_N^{*2} \xrightarrow{n \rightarrow \infty} \hat{\sigma}_\xi^2 (1 + m^{-1}) - \mathcal{O}(m^{-2})$ and thus

$$\lim_{n \rightarrow \infty} f_n^{\text{CAL}}(\hat{\boldsymbol{\theta}}) = \frac{\sigma_\xi^2}{\hat{\sigma}_\xi^2} + \mathcal{O}(m^{-2}).$$

430

□

Corollary 17 (NLL limit under full misspecification).

$$\lim_{n \rightarrow \infty} f_n^{\text{NLL}}(\hat{\boldsymbol{\theta}}) = \frac{1}{2} \log(\hat{\sigma}_\xi^2 (1 + m^{-1})) + \frac{1}{2} \frac{\sigma_\xi^2}{\hat{\sigma}_\xi^2} + \frac{1}{2} \log 2\pi - \mathcal{O}(m^{-2}).$$

431 *Proof.* We start with

$$2f_n^{\text{NLL}}(\hat{\boldsymbol{\theta}}) = \mathbb{E}_{\mathbf{y}, \mathbf{y}^*} \left\{ \log \hat{\sigma}_N^{*2} + \frac{(y^* - \hat{\mu}_N^*)^2}{\hat{\sigma}_N^{*2}} + \log 2\pi \right\}.$$

432 For the second term we use Lemma 16 so that we have

$$\lim_{n \rightarrow \infty} 2f_n^{\text{NLL}}(\hat{\boldsymbol{\theta}}) = \log \hat{\sigma}_\xi^2 + m^{-1} + \frac{\sigma_\xi^2}{\hat{\sigma}_\xi^2} + \log 2\pi - \mathcal{O}(m^{-2}).$$

433

□

434 *Proof of Theorem 1.* We construct the proof using all of the intermediate results given above. In
435 particular item (i) follows from Lemma 15, item (ii) from Lemma 16 and item (iii) from Corollary 17.

436

□

437 B Parameter Calibration (Proof of Lemma 4)

438 *Proof of Lemma 4.* (a) Replacing parameters $\hat{\theta} = (\hat{l}, \hat{\sigma}_\xi^2, \hat{\sigma}_f^2)$ with $\hat{\theta}' = (\hat{l}, \alpha\hat{\sigma}_\xi^2, \alpha\hat{\sigma}_f^2)$ changes all of
439 the σ_i^{*2} values to $\alpha\sigma_i^{*2}$ and therefore changes the calibration value on C from $\alpha = \frac{1}{c} \sum_{i=1}^c \frac{(y_i^* - \mu_i^*)^2}{\sigma_i^{*2}}$
440 to $\alpha/\alpha = 1$. (b) The NLL on C arising from parameters $(\hat{l}, \alpha\hat{\sigma}_\xi^2, \alpha\hat{\sigma}_f^2)$ is $\frac{1}{2c} \sum_{i=1}^c \{\log(\alpha\hat{\sigma}_\xi^2) +$
441 $(y_i^* - \mu_i^*)^2 / (\alpha\sigma_i^{*2}) + \log 2\pi\}$ which, on taking first and second derivatives w.r.t. α , is found to
442 be uniquely minimised by $\alpha = \frac{1}{c} \sum_{i=1}^c \frac{(y_i^* - \mu_i^*)^2}{\sigma_i^{*2}}$. (c) It is easily shown that replacing parameters
443 $(\hat{\sigma}_\xi^2, \hat{\sigma}_f^2)$ by $(k\hat{\sigma}_\xi^2, k\hat{\sigma}_f^2)$ (for any $k > 0$) in the formula for μ^* (Equation 3 and Equation 6) does not
444 alter μ^* . Hence the value of $\text{MSE} = \frac{1}{n^*} \sum_{i=1}^{n^*} (y_i^* - \mu_i^*)^2$ on any size- n^* test set is unchanged when
445 parameters $\hat{\theta}'$ are used in place of $\hat{\theta}$. \square

446 C Real world datasets

447 We consider a variety of datasets from the standard UCI machine learning repository¹. These datasets
448 are commonly used in the GP literature (see [8] for instance) and are, in principle, easily available
449 online. In practice, we encountered some difficulties: the dataset documentation is often limited;
450 the dataset names commonly used in other published papers do not always match the UCI database
451 naming and important details about data pre-processing, which features to use etc, are often omitted.
452 There are numerous attempts on GitHub and elsewhere at cataloguing these datasets along with any
453 pre-processing, however we had limited success using them, with many appearing unmaintained. Our
454 focus in this work is on testing our methods on a variety of real world datasets and in a way that is,
455 as far as possible, consistent with other papers. We therefore rejected datasets about which there is
456 ambiguity over the correct features to use, or even which column to regress on or for which outlier
457 rejection is required but undocumented elsewhere.

458 Referring to the datasets used in [8], we were able to locate the following:

- 459 • Song ([https://archive.ics.uci.edu/ml/machine-learning-databases/00203/](https://archive.ics.uci.edu/ml/machine-learning-databases/00203/YearPredictionMSD.txt.zip)
460 [YearPredictionMSD.txt.zip](https://archive.ics.uci.edu/ml/machine-learning-databases/00203/YearPredictionMSD.txt.zip))
- 461 • Bike ([https://archive.ics.uci.edu/ml/machine-learning-databases/00275/](https://archive.ics.uci.edu/ml/machine-learning-databases/00275/Bike-Sharing-Dataset.zip)
462 [Bike-Sharing-Dataset.zip](https://archive.ics.uci.edu/ml/machine-learning-databases/00275/Bike-Sharing-Dataset.zip))
- 463 • Poletele ([https://archive.ics.uci.edu/ml/machine-learning-databases/](https://archive.ics.uci.edu/ml/machine-learning-databases/parkinsons/telemonitoring/parkinsons_updrs.data)
464 [parkinsons/telemonitoring/parkinsons_updrs.data](https://archive.ics.uci.edu/ml/machine-learning-databases/parkinsons/telemonitoring/parkinsons_updrs.data))
- 465 • Keggdirected ([https://archive.ics.uci.edu/ml/machine-learning-databases/](https://archive.ics.uci.edu/ml/machine-learning-databases/00220/Relation%20Network%20(Directed).data)
466 [00220/Relation%20Network%20\(Directed\).data](https://archive.ics.uci.edu/ml/machine-learning-databases/00220/Relation%20Network%20(Directed).data))
- 467 • Keggundirected ([https://archive.ics.uci.edu/ml/](https://archive.ics.uci.edu/ml/machine-learning-databases/00221/Reaction%20Network%20(Undirected).data)
468 [machine-learning-databases/00221/Reaction%20Network%20\(Undirected\)](https://archive.ics.uci.edu/ml/machine-learning-databases/00221/Reaction%20Network%20(Undirected).data)
469 [.data](https://archive.ics.uci.edu/ml/machine-learning-databases/00221/Reaction%20Network%20(Undirected).data))
- 470 • CTSlice ([https://archive.ics.uci.edu/ml/machine-learning-databases/](https://archive.ics.uci.edu/ml/machine-learning-databases/00206/slice_localization_data.zip)
471 [00206/slice_localization_data.zip](https://archive.ics.uci.edu/ml/machine-learning-databases/00206/slice_localization_data.zip))
- 472 • Road3d ([https://archive.ics.uci.edu/ml/machine-learning-databases/](https://archive.ics.uci.edu/ml/machine-learning-databases/00246/3D_spatial_network.txt)
473 [00246/3D_spatial_network.txt](https://archive.ics.uci.edu/ml/machine-learning-databases/00246/3D_spatial_network.txt))
- 474 • Protein ([https://archive.ics.uci.edu/ml/machine-learning-databases/](https://archive.ics.uci.edu/ml/machine-learning-databases/00265/CASP.csv)
475 [00265/CASP.csv](https://archive.ics.uci.edu/ml/machine-learning-databases/00265/CASP.csv))
- 476 • Buzz ([https://archive.ics.uci.edu/ml/machine-learning-databases/00248/](https://archive.ics.uci.edu/ml/machine-learning-databases/00248/regression.tar.gz)
477 [regression.tar.gz](https://archive.ics.uci.edu/ml/machine-learning-databases/00248/regression.tar.gz))
- 478 • HouseElectric ([https://archive-beta.ics.uci.edu/dataset/235/individual+](https://archive-beta.ics.uci.edu/dataset/235/individual+household+electric+power+consumption)
479 [household+electric+power+consumption](https://archive-beta.ics.uci.edu/dataset/235/individual+household+electric+power+consumption))

480 We were unable to find any documentation on the Kegg datasets to indicate which of the columns
481 should be used as the independent variable (the regressor) and neither is this mentioned in any

¹<https://archive-beta.ics.uci.edu>, accessed April 2023.

482 literature of which we are aware. Initial runs of standard exact GP training and prediction produced
 483 RMSEs much higher than reported in [8]. Combining these two observations, we chose to exclude
 484 both Kegg datasets. Likewise we faced problems with Buzz. An analysis of the y values revealed a
 485 small proportion of extremely large outliers that we found could unduly distort performance results
 486 (e.g. depending on whether these outliers appeared in the test set for some of the random splits).
 487 With the lack of documentation we were unable to identify an outlier rejection scheme that we were
 488 confident would be consistent with results quoted in other papers. For this reason we have excluded
 489 Buzz.

490 The choice of (x, y) value that we applied for each of the used datasets is as follows:

- 491 • Song. The first column is y , all remaining columns are x .
- 492 • Bike. We use `hour.csv`. The y value is `cnt`. `dteday` (the date) is transformed to just be
 493 the integer representation of the day. `instant` is just an index so is dropped. `registered`
 494 and `casual` are dropped as `registered + casual = cnt`.
- 495 • Poletele. The y value is `total_UPDRS`. The columns `subject#` and `test_time` are not
 496 relevant to the problem so are dropped.
- 497 • CTSlice. y value is the final column. The first column is dropped as it is just an index. We
 498 additionally drop six columns which are constant over the majority of the dataset, namely
 499 columns 59, 69, 179, 189, 279 and 351.
- 500 • Road3d. y value is the final column. The first column is dropped as it is just an index.
- 501 • Protein. This dataset was processed as per [https://github.com/hughsalimbeni/](https://github.com/hughsalimbeni/bayesian_benchmarks)
 502 [bayesian_benchmarks](https://github.com/hughsalimbeni/bayesian_benchmarks), whereafter we used our own random (seeded) train/test split.
- 503 • HouseElectric. y value is the column labelled “Global active power”, rescaled by 1000/60
 504 and with “Sub metering 1,2,3” columns subtracted. We convert the date column into day-of-
 505 year/365 and the time column into time of day in minutes. Further, we remove any rows
 506 with null entries.

507 We note that although we are using a standard set of real-world datasets, it is not always clear exactly
 508 how others in the field have carried out their own preprocessing, limiting the ability to make direct
 509 comparisons to other results reported in the literature.

510 D Additional Implementation Details

511 D.1 Pre-whitening of Data

512 For all datasets covered in subsection 7.1 the following “whitening” preprocessing step is adopted:
 513 Let \mathbf{y} be the vector of all regressor values in the *training* dataset only, and X the matrix of all
 514 regressands in the *training* dataset only, where each row of X is a feature. Let μ_y, σ_y^2 be the sample
 515 mean and variance of \mathbf{y} respectively in the training dataset, then the whitened y values used in both
 516 the training and test set are simply $\sigma_y^{-1}(y - \mu_y)$. Let μ_X, Σ_X be the sample mean and covariance
 517 matrix of X respectively. Let $\Sigma_X = MM^T$, then the whitened x values in both the training and test
 518 data are $\frac{1}{\sqrt{d}}M^{-1}(\mathbf{x} - \mu_X)$, where d is the feature dimension of X . **Note:** the performance metrics
 519 given in subsection 7.1 are expressed in terms of the whitened y values rather than the y values in
 520 their original form. This appears to be common practice in the literature and has no bearing on the
 521 *comparative* performance of the different methods within this paper.

522 D.2 Test-Set Batching

523 To prevent excessive memory consumption, we perform all predictions for the distributed and
 524 variational methods in batches of 1000 points at a time. Where this is not possible (e.g. for especially
 525 large datasets), we use smaller batches of 500 or 250 points, as appropriate.

526 D.3 Additional Implementation Details for SVGP

527 We use the sparse variational inducing point approach of [11], following the implementation
 528 provided by GPyTorch, which in particular uses a Cholesky decomposition to parameterise the

529 covariance matrix of the variational prior. We broadly follow the SVGP implementation exam-
530 ple provided by [https://docs.gpytorch.ai/en/stable/examples/04_Variational_and_](https://docs.gpytorch.ai/en/stable/examples/04_Variational_and_Approximate_GPs/SVGP_Regression_CUDA.html)
531 [Approximate_GPs/SVGP_Regression_CUDA.html](https://docs.gpytorch.ai/en/stable/examples/04_Variational_and_Approximate_GPs/SVGP_Regression_CUDA.html). In particular, we follow their example in us-
532 ing the Adam optimiser to train our model over 100 epochs with a minibatch size of 1024 and a
533 learning rate of 0.01. We opt to use 1024 inducing points. All experiments under this method are
534 run on a SageMaker ml.p3.2xlarge instance, consisting of a single Tesla V100 GPU with 16GB of
535 memory.

536 **D.4 Additional Implementation Details for Distributed methods**

537 A good introduction to distributed methods for Gaussian process inference is [7]. Here we run
538 the product-of-experts (PoE) [12], generalised product-of-experts (gPoE) [2], Bayesian committee
539 machine (BCM) [24], robust Bayesian committee machine (rBCM) [7] and generalised robust
540 Bayesian committee machine (GrBCM) [13] following the recommendation in [4] to aggregate in
541 f -space. There are three components to any distributed method: the hyperparameter inference, the
542 *partitioner* and the *aggregator*. Hyperparameter estimation is the same for all of the methods: we use
543 the method in section 3.1 of [7], randomly partitioning the entire training set into subsets of size 625
544 (or as close as possible with equal-sized experts given that in general n is not a multiple of 625). A
545 block diagonal approximation (with $n/625$ blocks) is then used to approximate to the full $n \times n$ gram
546 kernel matrix. To recover hyperparameters with this we use Gaussian Process models with a zero
547 prior mean and a scaled square-exponential kernel. Training is conducted using the Adam optimiser
548 with a learning rate of 0.1 over 100 optimiser iterations. Once the hyperparameters are trained, we
549 run our distributed prediction mechanism to evaluate performance against the test-set. The 625-sized
550 partitioned blocks are referred to as “experts” and the shared hyperparameter values are distributed to
551 each expert and held fixed thereafter. In the *aggregator*, or distributed prediction phase, each expert
552 produces an individual predictive distribution and these are then aggregated to a final predictive mean
553 and variance for each of our test points. GRBCM prediction is a little more complex than this as it
554 makes use of an additional “communications” expert as explained in [13], aggregating in f -space as
555 recommended in [4]. We provide timing statistics for training these models.

556 We use our own GPyTorch-based implementation of distributed GP approximations. All exact GP
557 calculations are performed using GPyTorch using the default settings (so 20 Lanczos iterations
558 throughout and a CG tolerance of 1 for hyperparameter inference, and 10^{-3} for posterior predictions).
559 For all of our experiments, we utilise an AWS t3.2xlarge instance (consisting of 8 Intel Skylake
560 Processors and 32 GB of RAM).

561 **E Further simulation results**

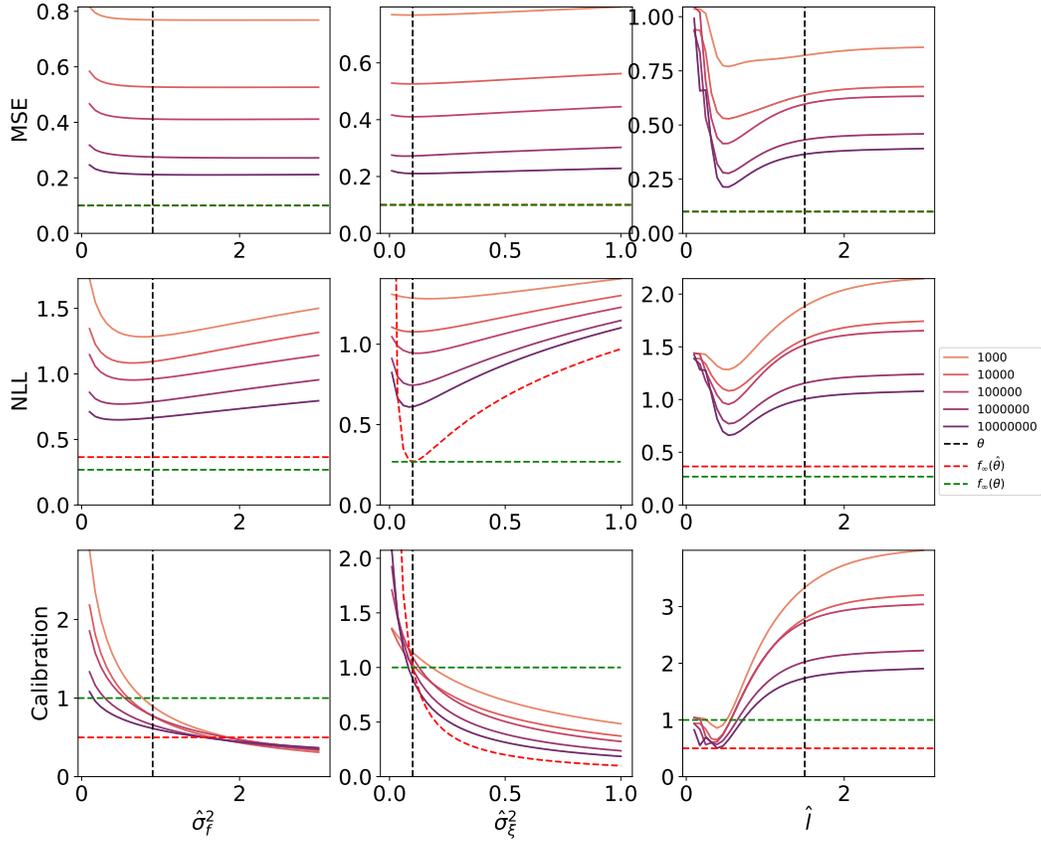


Figure 6: Behaviour of performance metrics as functions of kernel hyperparameters for increasing training set sizes n . The black dashed line denotes the true parameter value; the red dashed line shows the limiting behaviour as $n \rightarrow \infty$ and the green dashed line shows the limiting behaviour when the hyperparameters are correct. Simulations run with $d = 20, l = 0.5, \sigma_\xi^2 = 0.1, \sigma_f^2 = 0.9$. Assumed parameters when constant: $\hat{\sigma}_\xi^2 = 0.2, \hat{\sigma}_f^2 = 0.8, \hat{l} = 0.5$.

562 **F Further Results on UCI Datasets**

563 **F.1 Results for all distributed methods**

Table 3: Results for all methods on all metrics.

Dataset	n	d	Model	Calibration	NLL	RMSE
Bike	1.4e+04	13	BCM	1.02 ± 0.02	1.0 ± 0.0065	0.66 ± 0.0043
			GPOE	0.873 ± 0.012	1.03 ± 0.0069	0.664 ± 0.0054
			GRBCM	0.893 ± 0.014	0.977 ± 0.0057	0.634 ± 0.004
			OURS	0.974 ± 0.087	0.953 ± 0.013	0.624 ± 0.0079
			POE	1.03 ± 0.022	1.01 ± 0.0083	0.664 ± 0.0054
			RBCM	1.01 ± 0.02	1.0 ± 0.0065	0.659 ± 0.0043
			SVGP	0.898 ± 0.011	0.93 ± 0.0043	0.606 ± 0.0033
Ctslice	4.2e+04	378	BCM	5.04 ± 0.28	1.43 ± 0.13	0.311 ± 0.0052
			GPOE	0.435 ± 0.013	0.422 ± 0.0015	0.347 ± 0.0027
			GRBCM	1.13 ± 0.11	-0.159 ± 0.052	0.237 ± 0.012
			OURS	1.04 ± 0.085	-1.26 ± 0.01	0.132 ± 0.00062
			POE	6.39 ± 0.27	2.08 ± 0.12	0.347 ± 0.0027
			RBCM	4.16 ± 0.25	0.987 ± 0.11	0.28 ± 0.0048
			SVGP	0.865 ± 0.026	0.467 ± 0.016	0.384 ± 0.0064
Houseelectric	1.6e+06	8	BCM	1.27 ± 0.0046	-1.33 ± 0.0009	0.0634 ± 3.5e-05
			GPOE	0.908 ± 0.0065	-1.43 ± 0.0016	0.0638 ± 7.7e-05
			GRBCM	1.25 ± 0.011	-1.34 ± 0.0039	0.063 ± 0.00026
			OURS	1.08 ± 0.21	-1.56 ± 0.0065	0.0506 ± 0.00072
			POE	1.28 ± 0.006	-1.32 ± 0.0018	0.0638 ± 7.7e-05
			RBCM	1.24 ± 0.0054	-1.34 ± 0.0013	0.0626 ± 5.2e-05
			SVGP	0.911 ± 0.038	-1.46 ± 0.0046	0.0566 ± 0.00011
Poletele	4.6e+03	19	BCM	1.07 ± 0.029	0.00035 ± 0.019	0.243 ± 0.0048
			GPOE	0.917 ± 0.02	0.0344 ± 0.013	0.246 ± 0.0038
			GRBCM	0.872 ± 0.024	0.0091 ± 0.015	0.241 ± 0.0033
			OURS	1.03 ± 0.073	-0.214 ± 0.019	0.195 ± 0.0042
			POE	1.1 ± 0.036	0.00772 ± 0.016	0.246 ± 0.0038
			RBCM	1.08 ± 0.029	0.00309 ± 0.018	0.243 ± 0.0048
			SVGP	0.862 ± 0.035	-0.0667 ± 0.017	0.226 ± 0.0059
Protein	3.6e+04	9	BCM	1.04 ± 0.0097	1.14 ± 0.003	0.754 ± 0.0022
			GPOE	0.925 ± 0.007	1.15 ± 0.0035	0.763 ± 0.0024
			GRBCM	0.95 ± 0.012	1.11 ± 0.0051	0.733 ± 0.0038
			OURS	0.991 ± 0.029	1.01 ± 0.0016	0.666 ± 0.0014
			POE	1.07 ± 0.0088	1.15 ± 0.0033	0.763 ± 0.0024
			RBCM	1.03 ± 0.0096	1.13 ± 0.003	0.752 ± 0.0022
			SVGP	0.908 ± 0.016	1.05 ± 0.0059	0.688 ± 0.0043
Road3D	3.4e+05	2	BCM	1.01 ± 0.017	0.753 ± 0.007	0.514 ± 0.0035
			GPOE	0.756 ± 0.012	0.819 ± 0.0054	0.529 ± 0.0037
			GRBCM	0.873 ± 0.011	0.685 ± 0.0041	0.478 ± 0.0023
			OURS	0.991 ± 0.041	0.371 ± 0.004	0.351 ± 0.0014
			POE	1.07 ± 0.019	0.783 ± 0.0076	0.529 ± 0.0037
			RBCM	0.976 ± 0.016	0.735 ± 0.0066	0.505 ± 0.0034
			SVGP	0.9 ± 0.00094	0.608 ± 0.018	0.443 ± 0.008
Song	4.6e+05	90	BCM	1.56 ± 0.0063	1.32 ± 0.0012	0.851 ± 6.7e-05
			GPOE	0.926 ± 0.00049	1.27 ± 3.4e-05	0.864 ± 7.5e-05
			GRBCM	1.61 ± 0.11	1.46 ± 0.058	0.961 ± 0.035
			OURS	0.99 ± 0.037	1.18 ± 0.0045	0.787 ± 0.0045
			POE	1.61 ± 0.0067	1.34 ± 0.0013	0.864 ± 7.5e-05
			RBCM	1.56 ± 0.0062	1.31 ± 0.0011	0.851 ± 6.4e-05
			SVGP	0.991 ± 0.02	1.24 ± 0.0012	0.834 ± 0.0011

Table 4: Results on the UCI datasets using different kernel choices for our method and demonstrating the apparent superiority of the exponential kernel in these cases.

Dataset	n	d	Calibration				
			Distributed	Ours (Exp)	Ours (Matérn)	Ours (RBF)	SVGP
Poletele	4.6e+03	19	0.872 ± 0.024	0.994 ± 0.15	0.971 ± 0.13	1.03 ± 0.073	0.862 ± 0.035
Bike	1.4e+04	13	0.893 ± 0.014	0.988 ± 0.098	0.971 ± 0.086	0.974 ± 0.087	0.898 ± 0.011
Protein	3.6e+04	9	0.95 ± 0.012	0.995 ± 0.038	0.993 ± 0.031	0.991 ± 0.029	0.908 ± 0.016
Ctslice	4.2e+04	378	1.13 ± 0.11	0.912 ± 0.071	1.04 ± 0.082	1.04 ± 0.085	0.865 ± 0.026
Road3D	3.4e+05	2	0.873 ± 0.011	1.09 ± 0.065	1.0 ± 0.054	0.991 ± 0.041	0.9 ± 0.00094
Song	4.6e+05	90	1.56 ± 0.0063	0.995 ± 0.033	0.994 ± 0.035	0.99 ± 0.037	0.991 ± 0.02
Houseelectric	1.6e+06	8	1.24 ± 0.0054	1.11 ± 0.29	1.08 ± 0.27	1.08 ± 0.21	0.911 ± 0.038

Dataset	n	d	RMSE				
			Distributed	Ours (Exp)	Ours (Matérn)	Ours (RBF)	SVGP
Poletele	4.6e+03	19	0.241 ± 0.0033	0.169 ± 0.0076	0.17 ± 0.0076	0.195 ± 0.0042	0.226 ± 0.0059
Bike	1.4e+04	13	0.634 ± 0.004	0.565 ± 0.0036	0.6 ± 0.0044	0.624 ± 0.0079	0.606 ± 0.0033
Protein	3.6e+04	9	0.733 ± 0.0038	0.58 ± 0.0068	0.629 ± 0.004	0.666 ± 0.0014	0.688 ± 0.0043
Ctslice	4.2e+04	378	0.237 ± 0.012	0.123 ± 0.004	0.126 ± 0.0024	0.132 ± 0.00062	0.384 ± 0.0064
Road3D	3.4e+05	2	0.478 ± 0.0023	0.0976 ± 0.013	0.27 ± 0.01	0.351 ± 0.0014	0.443 ± 0.008
Song	4.6e+05	90	0.851 ± 6.7e-05	0.776 ± 0.004	0.778 ± 0.0045	0.787 ± 0.0045	0.834 ± 0.0011
Houseelectric	1.6e+06	8	0.0626 ± 5.2e-05	0.045 ± 0.00025	0.0485 ± 0.0004	0.0506 ± 0.00072	0.0566 ± 0.0001

Dataset	n	d	NLL				
			Distributed	Ours (Exp)	Ours (Matérn)	Ours (RBF)	SVGP
Poletele	4.6e+03	19	0.0091 ± 0.015	-0.397 ± 0.028	-0.346 ± 0.032	-0.214 ± 0.019	-0.0667 ± 0.017
Bike	1.4e+04	13	0.977 ± 0.0057	0.854 ± 0.004	0.915 ± 0.0077	0.953 ± 0.013	0.93 ± 0.0043
Protein	3.6e+04	9	1.11 ± 0.0051	0.853 ± 0.013	0.95 ± 0.0061	1.01 ± 0.0016	1.05 ± 0.0059
Ctslice	4.2e+04	378	-0.159 ± 0.052	-1.05 ± 0.027	-1.31 ± 0.017	-1.26 ± 0.01	0.467 ± 0.016
Road3D	3.4e+05	2	0.685 ± 0.0041	-0.931 ± 0.14	0.109 ± 0.039	0.371 ± 0.004	0.608 ± 0.018
Song	4.6e+05	90	1.32 ± 0.0012	1.16 ± 0.0046	1.17 ± 0.0051	1.18 ± 0.0045	1.24 ± 0.0012
Houseelectric	1.6e+06	8	-1.34 ± 0.0013	-1.95 ± 0.028	-1.62 ± 0.0095	-1.56 ± 0.0065	-1.46 ± 0.0046

565 **G Overall Computational Expenditure**

566 Our distributed and variational method experiments were conducted using cloud computing resources.
 567 Experiments using our own method have been carried out on an author’s laptop. SVGP experiments
 568 were run using a SageMaker virtual machine on a single Nvidia Tesla V100 GPU with 16GB memory.
 569 Distributed method experiments were run using eight Intel Xeon Platinum 8000 CPU cores (t3.2xlarge
 570 EC2 instances).

571 Below we will attempt to give reasonable indications of the amount of computational work expended
 572 to obtain the results in this paper, though note that we are neglecting the work expended in the
 573 development and research stages that did not directly contribute to the runs in the paper. As such,
 574 the costs presented are representative of the costs of replicating our paper, not repeating the research
 575 from scratch. Instead of reporting costs in dollars, we will report approximate computing hours for
 576 each instance type. The reader can then estimate their own costs using the current instance costs in
 577 the region of their choice, or under other cloud providers or even using on-premise compute.

Dataset	Billed hours (1 GPU, 3 runs)	Billed hours (8 CPUs, 3 runs of 5 methods)
bike	0.027	0.222
ctslice	0.082	0.546
houseelectric	3.713	256.776
poletele	0.010	0.079
protein	0.068	0.680
road3d	0.635	31.674
song	0.904	12.237

578 This gives a total of around 5.4 hours of compute time on a 1 GPU VM and 302.2 hours on an 8 CPU
 579 VM.

580 **References**

581 [1] F. Bachoc, N. Durrande, D. Rullière, and C. Chevalier. Properties and Comparison of Some
 582 Kriging Sub-model Aggregation Methods. *Mathematical Geosciences*, 2022.

583 [2] Y. Cao and D. J. Fleet. Generalized product of experts for automatic and principled fusion of
 584 gaussian process predictions. *arXiv preprint arXiv:1410.7827*, 2014.

585 [3] K. Chalupka, C. K. Williams, and I. Murray. A framework for evaluating approximation
 586 methods for gaussian process regression. *Journal of Machine Learning Research*, 14:333–350,
 587 2013.

588 [4] S. Cohen, R. Mbuva, T. Marwala, and M. Deisenroth. Healing products of Gaussian process
 589 experts. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference*
 590 *on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages
 591 2068–2077. PMLR, 2020-13.

592 [5] A. Datta, S. Banerjee, A. O. Finley, and A. E. Gelfand. Hierarchical Nearest-Neighbor Gaus-
 593 sian Process Models for Large Geostatistical Datasets. *Journal of the American Statistical*
 594 *Association*, 111(514):800–812, 2016.

595 [6] A. Datta, S. Banerjee, A. O. Finley, and A. E. Gelfand. On nearest-neighbor Gaussian process
 596 models for massive spatial data. *Wiley Interdisciplinary Reviews: Computational Statistics*,
 597 8(5):162–171, 2016.

598 [7] M. Deisenroth and J. W. Ng. Distributed gaussian processes. In *International Conference on*
 599 *Machine Learning*, pages 1481–1490. PMLR, 2015.

600 [8] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson. Gpytorch: Blackbox
 601 matrix-matrix gaussian process inference with gpu acceleration. *Advances in neural information*
 602 *processing systems*, 31, 2018.

603 [9] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric*
 604 *Regression*. Springer Series in Statistics, 2010.

- 605 [10] K. Hayashi, M. Imaizumi, and Y. Yoshida. On random subsampling of gaussian process
606 regression: A graphon-based analysis. In *International Conference on Artificial Intelligence*
607 *and Statistics*, pages 2055–2065. PMLR, 2020.
- 608 [11] J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. *arXiv preprint*
609 *arXiv:1309.6835*, 2013.
- 610 [12] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural*
611 *computation*, 14(8):1771–1800, 2002.
- 612 [13] H. Liu, J. Cai, Y. Wang, and Y. S. Ong. Generalized robust bayesian committee machine for
613 large-scale gaussian process regression. In *International Conference on Machine Learning*,
614 pages 3131–3140. PMLR, 2018.
- 615 [14] H. Liu, Y.-S. Ong, X. Shen, and J. Cai. When gaussian process meets big data: A review of
616 scalable gps. *IEEE transactions on neural networks and learning systems*, 31(11):4405–4423,
617 2020.
- 618 [15] J. E. Oakley and A. O’Hagan. Probabilistic sensitivity analysis of complex models: A bayesian
619 approach. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 66(3):751–
620 769, 2004.
- 621 [16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel,
622 P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher,
623 M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine*
624 *Learning Research*, 12:2825–2830, 2011.
- 625 [17] B. Schölkopf. The Kernel Trick for Distances. In T. Leen, T. Dietterich, and V. Tresp, editors,
626 *Advances in Neural Information Processing Systems*, volume 13. MIT Press, 2000.
- 627 [18] H. Song, T. Diethe, M. Kull, and P. Flach. Distribution calibration for regression. *36th*
628 *International Conference on Machine Learning, ICML 2019*, 2019-June:10347–10356, 2019.
- 629 [19] M. L. Stein, Z. Chi, and L. J. Welty. Approximating likelihoods for large spatial data sets.
630 *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 66(2):275–296, 2004.
- 631 [20] A. Stephenson, R. Allison, and E. Pyzer-Knapp. Provably reliable large-scale sampling from
632 gaussian processes. *arXiv preprint arXiv:2211.08036*, 2022.
- 633 [21] S. Surjanovic and D. Bingham. Virtual library of simulation experiments: Test functions and
634 datasets. Retrieved May 5, 2023, from <http://www.sfu.ca/~ssurjano>.
- 635 [22] M. Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial*
636 *intelligence and statistics*, pages 567–574. PMLR, 2009.
- 637 [23] G.-L. Tran, D. Miliotis, P. Michiardi, and M. Filippone. Sparse within sparse gaussian processes
638 using neighbor information. In M. Meila and T. Zhang, editors, *Proceedings of the 38th Inter-*
639 *national Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning*
640 *Research*, pages 10369–10378. PMLR, 2021-18.
- 641 [24] V. Tresp. A bayesian committee machine. *Neural computation*, 12(11):2719–2741, 2000.
- 642 [25] K. Wang, G. Pleiss, J. Gardner, S. Tyree, K. Q. Weinberger, and A. G. Wilson. Exact gaussian
643 processes on a million data points. *Advances in Neural Information Processing Systems*, 32,
644 2019.
- 645 [26] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT
646 press Cambridge, MA, 2006.
- 647 [27] L. Wu, G. Pleiss, and J. Cunningham. Variational Nearest Neighbor Gaussian Process. *Proceed-*
648 *ings of the 39th International Conference on Machine Learning*, 2022.