

<sup>1</sup>Aalto University

Finnish Center for Artificial Intelligence

# Fantasizing with Dual GPs in Bayesian **Optimization and Active Learning**

Paul E. Chang<sup>1</sup> Prakhar Verma<sup>1</sup> ST John<sup>1</sup> Victor Picheny<sup>2</sup> Henry Moss<sup>2</sup> Arno Solin<sup>1</sup>





## TL;DR

- Our dual conditioning efficiently updates variational parameters of a sparse GP in sequential/streaming settings.
- Unifying previous methods based on dual variables, we approximate the full-data posterior without revisiting previous batches.
- Our method works well in sequential decision making, particularly in batch acquisition of new data.

## Sequential data classification





(c) OVC (previous work by Maddox et al., 2021)

Figure 1: Sequential banana data set: our one-step updates on batches give similar solution to multiple optimization steps on full training data.

#### Bayesian optimization (BO) and active learning (AL)

- BO: black-box optimization. AL: learning a decision boundary.
- Both techniques combine an acquisition function  $\alpha(\cdot)$  and a surrogate model to encode a trade-off between exploration and exploitation.
- Standard BO/AL usually find one query points, however in high data settings it may be beneficial to find a **batch** of query points.

#### Variational sparse surrogate GP model

- Sparse GP methods swap computations on the full training set X with a sparser set of inducing points  $\mathbf{Z} \coloneqq (\mathbf{z}_j)_{j=1}^m$ ,  $(\mathbf{u})_j = f(\mathbf{z}_j)$ , with the approximate posterior  $q(\mathbf{u}; \mathbf{m}, \mathbf{V})$ .
- Common approach: use the variational ELBO to learn  $(\mathbf{m}, \mathbf{V})$ .
- [1] show that optimal variational parameters decompose as:

$$\mathbf{m}^* \equiv \mathbf{V}^* \boldsymbol{\lambda}^*$$
 and  $\mathbf{V}^* \equiv [\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1} + \boldsymbol{\Lambda}^*]^{-1}$ . (1)

We therefore reparameterize our variational parameters as  $\lambda$  and  $\Lambda$ .

## **Model & Methods**

- We show how to use **dual conditioning** for batch BO/AL. We build a batch greedily by 'fantasizing' for each point an observation value and updating the posterior [2, 3].
- Our dual conditioning updates are as follows:

- We apply our batch method to the lunar landing problem where the goal is to land a rocket successfully on a specific region of a lunar surface under stochasticity.
- The acquisition function is a product of Expected Improvement of the regression model and the predictive mean of the classification model. Therefore, needing both a Gaussian and non-Gaussian surrogate model.
- Our method shows an improvement over the non-batch baseline (see Figure 2).



Figure 2: We test our method on lunar landing problem where we need to find optimal parameters of landing a rocket with wind stochasticity.

In summary, we present dual conditioning, a method to condition data quickly into a sparse GP posterior. We show its usefulness for BO and AL

 $\boldsymbol{\lambda}_{\text{new}}^* \leftarrow \boldsymbol{\lambda}_{\text{old}}^* + \nabla_{\boldsymbol{\mu}^{(1)}} \mathbb{E}_{q_{\mathbf{u}}(\mathbf{f}_{\text{new}})}[\log p(\mathbf{y}_{\text{new}} \mid \mathbf{f}_{\text{new}})],$  $\boldsymbol{\Lambda}_{\text{new}}^* \leftarrow \boldsymbol{\Lambda}_{\text{old}}^* + \nabla_{\boldsymbol{\mu}^{(2)}} \mathbb{E}_{q_{\mathbf{u}}(\mathbf{f}_{\text{new}})}[\log p(\mathbf{y}_{\text{new}} \mid \mathbf{f}_{\text{new}})],$ 

where  $q_{\mathbf{u}}(\cdot)$  is the posterior process and  $\boldsymbol{\mu}$  are the expectation parameters of  $q(\mathbf{u})$ .

- The updates Eq. (2) are exact for Gaussian likelihoods and for non-Gaussian likelihoods recreate the posterior well compared to the offline solution (see Figure 1).
- Novelty of dual conditioning: does not need an optimization loop [4] or use previous data when conditioning the posterior on new information.
- We use Eq. (2) to iteratively add 'fantasized' data points from our  $\alpha(\cdot)$ to build a batch of query points.
- The method is independent of the acquisition function and can be used to perform batch BO/AL for any simple acquisition function.

(2)

### More details...

See the paper:



arxiv.org/abs/ 2211.01053

#### References

- [1] V. Adam, P. Chang, M. E. E. Khan, and A. Solin, "Dual parameterization of sparse variational Gaussian processes," Advances in Neural Information Processing Systems, 2021.
- D. Ginsbourger, R. L. Riche, and L. Carraro, "Kriging is well-suited to [2] parallelize optimization," in *Computational intelligence in expensive* optimization problems, pp. 131–162, Springer, 2010.
- [3] T. Desautels, A. Krause, and J. W. Burdick, "Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization," Journal of Machine Learning Research, 2014.
- [4] W. J. Maddox, S. Stanton, and A. G. Wilson, "Conditioning sparse variational Gaussian processes for online decision-making," Advances in Neural Information Processing Systems, 2021.

paul.chang@aalto.fi