

DEEP GAUSSIAN PROCESSES FOR GEOPHYSICAL PARAMETER RETRIEVAL

Daniel Heestermans Svendsen^{1,*}, Pablo Morales-Álvarez^{2,*}, Rafael Molina², Gustau Camps-Valls¹

¹Image Processing Laboratory (IPL), Universitat de València, Spain

²Computer Science and Artificial Intelligence Department, Universidad de Granada, Spain

ABSTRACT

This paper introduces deep Gaussian processes (DGPs) for geo-physical parameter retrieval. Unlike the standard full GP model, the DGP accounts for complicated (modular, hierarchical) processes, provides a efficient solution that scales well to large datasets, and improves prediction accuracy over standard full and sparse GP models. We give empirical evidence of performance for estimation of surface dew point temperature from infrared sounding data.

Index Terms— Model inversion, statistical retrieval, Deep Gaussian Processes, surface dew point temperature, IASI, infrared sounding

1. INTRODUCTION

Remote sensing applications often involve solving a complex model inversion problem to estimate bio-geo-physical parameters from observations. The problem is very challenging as it involves nonlinear input-output relations, as well as a wide diversity of noise sources and distortions. In addition, very often, the goal is to invert *metamodels*, that is, combinations of submodels that are coupled together. Radiative transfer models (RTM) describe the processes occurring at different scales (e.g. at leaf, canopy and atmospheric levels) with different complexities. The overall process is thus complicated and with different amounts of uncertainties that accumulate and propagate.

Inverting such models is complicated and several strategies exist: interpolation on look-up tables, direct physically-based inversion, or statistical approaches. Physics-based retrieval is often very costly and needs a broad knowledge about the governing mechanisms. In recent years, the remote sensing community has turned to more efficient and agnostic statistical *hybrid* approaches for model inversion [1]. After generating pairs of inputs (state vectors) and outputs (radiances) by running an RTM code, the problem boils down to performing inverse modeling via standard regression. Approximating arbitrary nonlinear functions from data is a solid field of machine learning where many successful methods are available.

In the last few years, data-driven statistical learning algorithms have attained outstanding results in the estimation of

climate variables and related geo-physical parameters at local and global scales [1, 2]. These algorithms avoid complicated assumptions and provide flexible non-parametric models that fit the observations using large heterogeneous data. In particular, during the last decade, the use of Gaussian Processes (GPs) in model inversion and emulation of physical models in the geosciences has gained popularity in general [3], and for atmospheric parameter retrieval in particular [4]. However, GPs have two important shortcomings: their high computational cost (n^3 in training and n^2 in testing), and the limited expressiveness of standard kernel functions (often restricted to the exponentiated quadratic family). On top of these, one should note the hierarchical modularity of the physical models to invert. This actually discourages the use of a plain, shallow GP model to account for all nonlinear feature relations.

To address the aforementioned problems, in this paper, we explore deep GP (DGP) regression for geophysical parameter retrieval and model inversion. DGPs were originally introduced in [5] and further analyzed in [6]. We will resort to the very recent DGP inference procedure in [7], which has provided excellent results in a wide range of regression problems. We will show that DGPs overcome the limitations of standard GPs, in particular how they: 1) account for complicated (modular, hierarchical) processes encoded in the RTM, 2) return an efficient solution that scales well to large datasets, and 3) improve prediction accuracy over standard full and sparse GP models. We illustrate performance in prediction of surface temperature and moisture from superspectral infrared sounding data [8].

2. PROBABILISTIC MODEL AND INFERENCE

Gaussian Processes [9] are non-parametric probabilistic state-of-the-art models for functions, that are successfully used in supervised and unsupervised learning. Notationally, for input-output data $\{(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i=1}^n$, a GP models the underlying dependence with latent variables $\{f_i = f(\mathbf{x}_i) \in \mathbb{R}\}_{i=1}^n$ that jointly follow a Gaussian $\mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K})$. The kernel matrix $\mathbf{K} = (k(\mathbf{x}_i, \mathbf{x}_j))_{i,j}$ depends on the kernel function k , which encodes the properties (e.g. smoothness) of the modeled functions. In regression problems, the observation model of the outputs y_i given the latent variables f_i is usually defined by the Gaussian $p(y_i|f_i, \sigma^2) = \mathcal{N}(y_i|f_i, \sigma^2)$.

This likelihood allows us to integrate out \mathbf{f} and compute the posterior $p(\mathbf{f}|\mathbf{y})$ in closed form solution (parameters are omitted for simplicity) [9]. However, this requires inverting

Research funded by the European Research Council (ERC) under the ERC-CoG-2014 SEDAL project (grant agreement 647423), the Spanish Ministry of Economy and Competitiveness through projects TIN2015-64210-R, DPI2016-77869-C2-2-R, and the Spanish Excellence Network TEC2016-81900-REDT. PMA received financial support through *La Caixa* Fellowship for Doctoral Studies (*La Caixa* Banking Foundation, Barcelona, Spain).

* The first two authors contributed equally.

the $n \times n$ matrix $(\mathbf{K} + \sigma^2 \mathbf{I})$, which scales cubically, $\mathcal{O}(n^3)$. This constraint makes GP prohibitive for large scale applications. Here, sparse GP approximations become the more efficient pathway.

Additionally, GPs are limited by the expressiveness of the kernel function. Although very complex kernels can be tailored to the data [9], in practice this may become unfeasible, as it requires a thorough application-specific knowledge and usually comes with a large amount of hyperparameters to estimate, which may cause overfitting. Alternatively, DGPs allow for modelling very complex data through a hierarchy of GPs that only use simple kernels with few hyperparameters.

2.1. Sparse GP approximations

The most popular approach to dealing with the computational burden of GPs is to introduce $m \ll n$ inducing points $\mathbf{u} = (u_1, \dots, u_m)$ which the inference is based on. These are GP realizations at the inducing locations $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_m\} \subset \mathbb{R}^d$, just like \mathbf{f} is at the inputs $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ [10]. One of the most widely used methods is the *Fully Independent Training Conditional (FITC)* introduced in [11].

FITC approximates the GP prior model¹ by assuming: i) conditional independence between train and test latent variables \mathbf{f}, \mathbf{f}_* given the inducing points \mathbf{u} ; and ii) a factorized (fully independent) distribution for \mathbf{f} given \mathbf{u} . The effective GP joint prior for FITC (which replaces the exact $p(\mathbf{f}, \mathbf{f}_*)$) is:

$$q(\mathbf{f}, \mathbf{f}_*) = \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \mathbf{Q}_{ff} + \text{diag}(\mathbf{K}_{ff} - \mathbf{Q}_{ff}) & \mathbf{Q}_{f*} \\ \mathbf{Q}_{*f} & \mathbf{K}_{**} \end{bmatrix} \right), \quad (1)$$

where we abbreviate $\mathbf{Q}_{ab} = \mathbf{K}_{au} \mathbf{K}_{uu}^{-1} \mathbf{K}_{ub}$. With this approximation, the observation model $p(\mathbf{y}|\mathbf{f}, \sigma^2)$ can be integrated out and the new matrix to be inverted is $(\mathbf{Q}_{ff} + \text{diag}(\mathbf{K}_{ff} - \mathbf{Q}_{ff}) + \sigma^2 \mathbf{I})$. Applying the Woodbury matrix identity, this low-rank-plus-diagonal matrix can be inverted with $\mathcal{O}(nm^2)$ complexity. Notice that this is linear in the training set size n .

2.2. Deep Gaussian Processes

In standard (single-layer) GPs, the output of the GP is directly used to model the observed response \mathbf{y} . However, this output could be used to define the input locations of another GP. If this is repeated L times, we obtain a hierarchy of GPs that is known as a Deep Gaussian Process (DGP) with L layers. This is analogous to the structure of deep neural networks, which are a cascade of generalized linear models [6, Chapter 6].

DGPs were first introduced in [5], where the authors performed approximate variational inference analytically. However, in each layer they use a set of latent variables that ends up inducing *independence across layers* in the posterior distribution. This uncorrelated posterior fails to express the complexity of the deep model, and is not realistic in practice. To

¹Among the sparse GP approximations based on inducing points, two families are distinguished: those that approximate the GP prior and perform exact inference (where FITC is the most popular), and those that utilize the exact prior but perform approximate inference (VFE [12] is the main example). Both families are compared in detail in [10]. A sparse GP approximation without inducing points was recently applied to remote sensing data [13].




Fig. 1. Graphical representation of the DGP model used here.

overcome this problem, we follow the recent inference procedure in [7], which keeps a strong conditional independence in the posterior by marginalizing out the aforementioned set of latent variables. In exchange, analytic tractability is sacrificed. However, we will see that the structure of the posterior allows one to efficiently sample from it and use Monte Carlo approximations. As will be justified later, this approach is called *Doubly Stochastic Variational Inference* [7].

DGPs can be used for regression by placing a Gaussian likelihood after the last layer. For notation simplicity, in this work the dimensions of the hidden layers will be fixed to one². However, exact inference in DGP is intractable (not only computationally expensive as in GPs), as it involves integrating out latent variables that are used as inputs in the next layer (i.e. they appear inside a complex kernel matrix). To overcome this, m inducing points \mathbf{u}^l at inducing locations \mathbf{z}^{l-1} are introduced at each layer l . Interestingly, we will see that this sparse formulation also makes DGP scale well to large datasets. For observed $\{\mathbf{X}, \mathbf{y}\}$, the regression joint model is

$$p(\mathbf{y}, \{\mathbf{f}^l, \mathbf{u}^l\}_{l=1}^L) = p(\mathbf{y}|\mathbf{f}^L) \prod_{l=1}^L p(\mathbf{f}^l|\mathbf{u}^l; \mathbf{f}^{l-1}, \mathbf{z}^{l-1}) p(\mathbf{u}^l; \mathbf{z}^{l-1}). \quad (2)$$

Here, $\mathbf{f}^0 = \mathbf{X}$, and each factor in the product is the joint distribution over $(\mathbf{f}^l, \mathbf{u}^l)$ of a GP in the inputs $(\mathbf{f}^{l-1}, \mathbf{z}^{l-1})$, but rewritten with the conditional probability given \mathbf{u}^{l-1} . Figure 1 shows a graphical representation of the described model.

2.2.1. Doubly Stochastic Variational Inference

The approach followed in [7] uses variational inference, where a lower bound of the marginal log-likelihood is maximized to find the optimal posterior among a family of distributions. The proposed posterior is

$$q(\{\mathbf{f}^l, \mathbf{u}^l\}_{l=1}^L) = \prod_{l=1}^L p(\mathbf{f}^l|\mathbf{u}^l; \mathbf{f}^{l-1}, \mathbf{z}^{l-1}) q(\mathbf{u}^l). \quad (3)$$

The first factor is the prior conditional of eq. (2) and keeps correlations between layers. The second is taken Gaussian with mean \mathbf{m}^l and full covariance \mathbf{S}^l (variational parameters to be estimated). With this posterior, the next *evidence lower*

²This can be generalized straightforwardly, see both approaches [5, 7].

³Notice that a semicolon is used to specify the inputs of the GP.

bound (ELBO) for the marginal log-likelihood is obtained⁴:

$$\begin{aligned} \log p(\mathbf{y}) &= \log \int \frac{q(\{\mathbf{f}^l, \mathbf{u}^l\})}{q(\{\mathbf{f}^l, \mathbf{u}^l\})} p(\mathbf{y}, \{\mathbf{f}^l, \mathbf{u}^l\}) d\mathbf{f}^l d\mathbf{u}^l \geq \\ &\sum_{i=1}^n \mathbb{E}_{q(f_i^L)} [\log p(y_i | f_i^L)] - \sum_{l=1}^L \text{KL}(q(\mathbf{u}^l) \| p(\mathbf{u}^l; \mathbf{z}^{l-1})). \end{aligned} \quad (4)$$

The second term is tractable, as the KL divergence between Gaussians is known [9]. However, the expectation involves the marginals of the posterior at the last layer, $q(f_i^L)$. Next we see that, whereas this distribution is analytically intractable, it can be sampled efficiently using univariate Gaussians.

Marginalizing out the inducing points in eq. (3), the posterior for the GP layers $\{\mathbf{f}^l\}$ is

$$q(\{\mathbf{f}^l\}) = \prod_{l=1}^L q(\mathbf{f}^l | \mathbf{m}^l, \mathbf{S}^l; \mathbf{f}^{l-1}, \mathbf{z}^{l-1}) = \prod_{l=1}^L \mathcal{N}(\mathbf{f}^l | \tilde{\boldsymbol{\mu}}^l, \tilde{\boldsymbol{\Sigma}}^l), \quad (5)$$

where the vector $\tilde{\boldsymbol{\mu}}^l$ is given by $[\tilde{\boldsymbol{\mu}}^l]_i = \mu_{\mathbf{m}^l, \mathbf{z}^{l-1}}(f_i^{l-1})$ and the $n \times n$ matrix $\tilde{\boldsymbol{\Sigma}}^l$ by $[\tilde{\boldsymbol{\Sigma}}^l]_{ij} = \Sigma_{\mathbf{S}^l, \mathbf{z}^{l-1}}(f_i^{l-1}, f_j^{l-1})$. The specific form of the functions $\mu_{\mathbf{m}^l, \mathbf{z}^{l-1}}$ and $\Sigma_{\mathbf{S}^l, \mathbf{z}^{l-1}}$ can be found in [7, Eqs. (7-8)]. The key point here is to observe that, although the distribution in eq. (5) is fully coupled between layers (and thus the posterior in the last layer is analytically intractable), the i -th marginal at each layer $\mathcal{N}(f_i^l | [\tilde{\boldsymbol{\mu}}^l]_i, [\tilde{\boldsymbol{\Sigma}}^l]_{ii})$ only depends on the corresponding i -th input of the previous layer. This allows one to recursively sample $\hat{f}_i^1 \rightarrow \hat{f}_i^2 \rightarrow \dots \rightarrow \hat{f}_i^L$ from all the layers up to the last one by means of univariate Gaussians. Specifically, $\varepsilon_i^l \sim \mathcal{N}(0, 1)$ is first sampled and then for $l = 1, \dots, L$:

$$\hat{f}_i^l = \mu_{\mathbf{m}^l, \mathbf{z}^{l-1}}(\hat{f}_i^{l-1}) + \varepsilon_i^l \cdot \sqrt{\Sigma_{\mathbf{S}^l, \mathbf{z}^{l-1}}(\hat{f}_i^{l-1}, \hat{f}_i^{l-1})}. \quad (6)$$

Now, the expectation in the ELBO (recall eq. (4)) can be approximated with a Monte Carlo sample generated with eq. (6). This provides the first source of stochasticity. Since the ELBO factorizes across data points and the samples can be drawn independently for each point i , scalability is achieved through sub-sampling the data in mini-batches. This is the second source of stochasticity. The ELBO is maximized w.r.t. the variational parameters $\mathbf{m}^l, \mathbf{S}^l$, the inducing locations \mathbf{z}^l , and the kernel hyperparameters $\boldsymbol{\theta}^l$ (which, to alleviate the notation, have not been included in the equations). The complexity to evaluate the ELBO and its gradients is $\mathcal{O}(nm^2L)$. The code is integrated within GPflow (a GP framework built on top of Tensorflow) and is publicly available⁵.

To predict in a new x_* , eq. (6) is used to sample S times⁶ from the posterior up to the $(L-1)$ -th layer using the test location as initial input. This yields a set $\{f_*^{L-1}(s)\}_{s=1}^S$ with S samples. Then, the density over f_*^L is given by the Gaussian

mixture (recall that all the terms in eq. (5) are Gaussians):

$$q(f_*^L) = \frac{1}{S} \sum_{s=1}^S q(f_*^L | \mathbf{m}^L, \mathbf{S}^L; f_*^{L-1}(s), \mathbf{z}^{L-1}).$$

3. EXPERIMENTAL RESULTS

We compare prediction error and training times with a standard full GP and FITC, the most widely used sparse GP approximation.

3.1. Data description

We focus on assessing the performance of the DGP for predicting relative humidity at surface level, measured as dew point temperature T_{dew} [K], using data from the Infrared Atmospheric Sounding Interferometer (IASI). The IASI data are point measurements of approximately 25 km diameter with 8461 spectral components, ranging in the infrared emission spectra from 645 to 2760 cm^{-1} with 0.25 cm^{-1} resolution. We used the simulated database in [14], which consists of 13456 datapoints generated using the optimal spectral sampling (OSS) infrared radiative transfer code [15]. First, we remove certain bands from the spectrum that do not contain useful information for retrieval, reducing the data to 4699 spectral components, following the same scheme proposed in [4]. Dimensionality reduction is then performed with principal component analysis (PCA), projecting the spectra down to a 13 dimensional subspace, where 99.96% of the data variance is retained.

3.2. Predicting dew point temperature

The spectra were randomly split into training sets of sizes 300, 1000, 4000, 8000 and 10000, and the rest are left for testing. The resulting root mean squared error (RMSE) is shown in Fig. 2, which is averaged over several runs. We compare the following GP-based models:

GP: The standard full GP, using an exponentiated quadratic (EQ) kernel function. The computational cost of training scales like $\mathcal{O}(n^3)$ with the number of training points.

FITC: Introduced in Section 2.1, it is among the most popular sparse GP approximations. The EQ kernel is used, and the code is taken from GPflow <https://github.com/GPflow>. The cost of training scales like $\mathcal{O}(nm^2)$, and the number inducing points is set to $m = 100$.

DGP1-4: DGP described in Section 2.2 with 1-4 layers and $m = 100$ inducing inputs per layer. The number of hidden units per layer is equal to the dimension of the input data (13). Note that DGP1 is just a sparse GP using the inference scheme of Sec. 2.2.1. The computational cost is $\mathcal{O}(nm^2L)$.

As theoretically expected, the full GP consistently outperforms its sparse approximations FITC and DGP1. The DGP2-4 take advantage of their hierarchical structure and achieve even lower RMSE than the full GP. Moreover, as the amount of training data increases, the gap between DGP2 and DGP3,4 widens, indicating that there is a complex structure in the data

⁴The prior conditionals of eq. (3) smartly cancel with those of eq. (2).

⁵<https://github.com/ICL-SML/Doubly-Stochastic-DGP>

⁶Results become stable after a few samples. Here, S was set to 200.




Fig. 2. Performance of the compared methods as a function of the training set size. DGP2-4 outperform the full GP, which in turn is better than its sparse approximations DGP1 and FITC.

which can be learned when combining a sufficiently complex model (3 or more layers) with sufficient amounts of data.

Interestingly, we see that, for a small amount of training data ($n = 300$), the 2-, 3- and 4-layer DGPs, do not overfit and perform better than the single-layer sparse models and the full GP. This is in spite of having many parameters to fit compared to the amount of data, showing that the Bayesian training scheme adequately penalizes complexity.

Fig. 3 shows the training times for the different models. As theoretically expected, all models scale linearly except for the GP, which is $\mathcal{O}(n^3)$. This makes it impossible to train a GP on a dataset of magnitude $n \sim 10^5$ [9]. However, DGPs can be trained within a reasonable time span on big datasets. Thus, they constitute a very powerful alternative to classical sparse GP approximations (such as FITC), especially for datasets with a complex underlying structure.




Fig. 3. Training time of the compared methods for increasing n . We see how the DGP and FITC scale linearly with n while the full GP scales cubically, making it unfit for big data.

4. CONCLUSIONS

We introduced deep GPs and the doubly stochastic variational inference procedure for remote sensing applications. We showed how DGP benefits from its hierarchical structure and consistently outperforms a full GP in superspectral

infrared sounding data. Moreover, unlike the full GP, its sparse formulation makes it scale linearly with the training set size. Thus, DGP is further compared against FITC, the most widely used classical sparse GP approximation. The experimental results suggest that DGP performs better than current state-of-the-art GP-based methods in large scale remote sensing datasets, especially when there is a complex underlying process generating the data.

5. REFERENCES

- [1] G. Camps-Valls, D. Tuia, L. Gómez-Chova, and J. Malo, Eds., *Remote Sensing Image Processing*, Morgan & Claypool, Sept 2011.
- [2] G. Camps-Valls and L. Bruzzone, Eds., *Kernel methods for Remote Sensing Data Analysis*, Wiley & Sons, UK, Dec 2009.
- [3] G. Camps-Valls, J. Verrelst, J. Muñoz-Marí, V. Laparra, F. Mateo-Jimenez, and J. Gomez-Dans, "A survey on gaussian processes for earth observation data analysis: A comprehensive investigation," *IEEE Geoscience and Remote Sensing Magazine*, , no. 6, June 2016.
- [4] G. Camps-Valls, J. Muñoz-Marí, L. Gómez-Chova, L. Guanter, and X. Calbet, "Nonlinear statistical retrieval of atmospheric profiles from MetOp-IASI and MTG-IRS infrared sounding data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 5 Part 2, pp. 1759–1769, 2012.
- [5] A. Damianou and N. Lawrence, "Deep Gaussian processes," in *Artificial Intelligence and Statistics*, 2013, pp. 207–215.
- [6] A. Damianou, *Deep Gaussian processes and variational propagation of uncertainty*, Ph.D. thesis, University of Sheffield, 2015.
- [7] H. Salimbeni and M. Deisenroth, "Doubly stochastic variational inference for deep Gaussian processes," in *Advances in Neural Information Processing Systems*, 2017, pp. 4591–4602.
- [8] F Aires, "A regularized neural net approach for retrieval of atmospheric and surface temperatures with the IASI instrument," *Journal of Applied Meteorology*, vol. 41, pp. 144–159, 2002.
- [9] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, New York, 2006.
- [10] M. Bauer, M. van der Wilk, and C.E. Rasmussen, "Understanding probabilistic sparse Gaussian process approximations," in *Advances in Neural Information Processing Systems*, 2016, pp. 1533–1541.
- [11] E. Snelson and Z. Ghahramani, "Sparse Gaussian processes using pseudo-inputs," in *Advances in neural information processing systems*, 2006, pp. 1257–1264.
- [12] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Artificial Intelligence and Statistics*, 2009, pp. 567–574.
- [13] P. Morales-Álvarez, A. Pérez-Suay, R. Molina, and G. Camps-Valls, "Remote sensing image classification with large-scale Gaussian processes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. PP, no. 99, pp. 1–12, 2017.
- [14] F. Chevallier, *Sampled Databases of 60-level Atmospheric Profiles from the ECMWF Analyses: Version 1.0*, EUMETSAT, 2002.
- [15] J.L. Moncet, G. Uymin, A.E. Lipton, and H.E. Snell, "Infrared radiance modeling by optimal spectral sampling," *Journal of the Atmospheric Sciences*, vol. 65, no. 12, pp. 3917–3934, 2008.