

# Interpretable User Models via Decision-rule Gaussian Processes: Preliminary Results on Energy Storage

Danial Mohseni-Taheri

Selvaprabu Nadarajah

Theja Tulabandhula

*University of Illinois at Chicago*

SMOHSE3@UIC.EDU

SELVAN@UIC.EDU

THEJA@UIC.EDU

## 1. Introduction

Models of user behavior are critical in many decision making problems and can be viewed as decision rules that transform state information (in set  $\mathcal{S}$ ) available to the user to actions (in set  $\mathcal{A}$ ). Formally, a user model is a function  $f : \mathcal{S} \mapsto \mathcal{A}$ . Gaussian processes (GPs) employed to learn functions on the action/target space (henceforth target GPs or TGP for short) can thus be used to place a prior on user models and identify a posterior distribution over them supported by data in conjunction with approximate Bayesian inference techniques (Blei et al., 2017; Beaumont, 2019).

TGPs for user modeling would assume that user actions at a given set of finite states follow a multivariate Gaussian. To capture non-Gaussian action distributions, one could apply GPs to learn functions in a transformed space that is not the target. Examples include warped and chained GPs proposed in Snelson et al. (2004) and Saul et al. (2016), respectively. Extending this literature, we study the application of GPs in a transformed space defined by *decision rules*. Such rules are known in several applications and depend on functions themselves. Specifically, a user model based on a decision rule takes the form  $g : \prod_k \mathcal{P}_k \times \mathcal{S} \rightarrow \mathcal{A}$ , where the arguments are obtained using functions  $h_k : \mathcal{S} \mapsto \mathcal{P}_k$ ,  $k = \{1, \dots, K\}$  that map from  $\mathcal{S}$  to transformed spaces  $\mathcal{P}_k$ , possibly different from the target space  $\mathcal{A}$ . Each such function has immediate *interpretability* to a practitioner, and we model them using GPs. We refer to such a user model  $\{g, h_1, \dots, h_k\}$  as a decision-rule GP (DRGP).

To make the notion of DRGPs concrete in this short article, we focus on the problem faced by a firm providing services to store ethanol – a real application that motivated this work. Suppose capacity (in gallons) is sold via annual contracts to  $N$  users. The contract of user  $n$  specifies the maximum amount of ethanol that can be stored, denoted by  $C_n$ . User behavior corresponds to the injection of ethanol and the withdrawal of previously injected ethanol, which can be modeled as a time series. The inventory  $I_{n,t}$  in storage associated with user  $n$  at time  $t$  is the net of past injections and withdrawals. A TGP approach would employ a GP to determine the next-period storage inventory level function  $I_{n,t+1}$  directly. In contrast, we propose a DRGP that leverages a well-known decision rule based on injection and withdrawal threshold functions (Charnes et al., 1966; Secomandi, 2010). These threshold functions are learned as GPs instead of the (relatively less interpretable) inventory function.

We focus on the following research questions in the context of the ethanol storage application: (Q1) Can existing exact and approximate Bayesian inference techniques be used for inference with DRGP? and (Q2) How does DRGP perform relative to TGP?

We answer these questions by executing numerical experiments based on real data of aggregated ethanol storage injection and withdrawals. For Q1, we show that sparse vari-

ational inference (Titsias, 2009; Hensman et al., 2013), which can be applied to TGP on our data set, can also be used with DRGP, albeit heuristically, which is encouraging from an implementation standpoint. For Q2, we find that DRGP implemented in this manner leads to lesser out-of-sample error than TGP on most of our datasets, in addition to being more interpretable to practitioners. This preliminary finding is promising and suggests that applying GPs in the interpretable space of the decision rule threshold functions has potential value, which adds to the growing literature on interpretable machine learning and optimization (Letham et al., 2015; Bertsimas and Dunn, 2017). In addition, the improvements we report are based on the heuristic use of sparse variational inference with DRGPs, which bodes well for additional potential improvements from the development of new inference techniques targeting DRGPs. Finally, several applications in energy, health care, and transportation, among other domains, have known interpretable decision rules, which can be leveraged in the DRGP framework proposed here.

## 2. Related Work

Snelson et al. (2004) show that modeling data using a warped GP, which is a non-linear transformation (aka warping) of a GP, can enhance predictive performance. Inference using a warped GP can be performed in closed-form provided the warping function satisfies certain properties, such as being invertible. Lázaro-Gredilla (2012) consider the case where the warping function is not fixed a priori. DRGPs differ from warped GPs as they are based on a potentially *non-invertible* transformation of *multiple* GPs.

Chained GPs by Saul et al. (2016) extend warped GPs by considering a likelihood function that factorizes across the data and is a general nonlinear transformation of multiple latent functions, each modeled as a GP. Exact inference of chained GPs is not tractable in general and thus approximate inference techniques are used instead. See Lázaro-Gredilla and Titsias (2011); Vanhatalo et al. (2013); and Moreno-Muñoz et al. (2018) for alternative approaches to handle multiple GPs.

Recent work has focused on finding a balance between the modeling generality (restrictiveness) of chained (warped) GPs and its associated challenging (straightforward) inference procedure. For example, Tobar and Rios, 2019 extend a warped GP using a composition of simple functions and retain closed form inference. DRGP is similar to a chained GP because its underlying decision rule is a nonlinear transformation  $g(\cdot)$  of multiple GPs that model functions  $h_1, \dots, h_K$ . However, unlike a chained GP, each function  $h_k$  is *interpretable* and *not necessarily latent*, which simplifies inference (see §3 for details). For instance, in our energy storage application (where  $K = 2$ ), the functions  $h_1$  and  $h_2$  correspond to injection and withdrawal threshold functions, respectively, and are fully or partially observable.

## 3. DRGPs for Energy Storage

For each user  $n \in N$ , the most basic inventory update model capturing temporal dependencies can be written as:  $I_{n,t+1} = f_n(I_{n,t}, X_{n,t}) + \epsilon_{n,t}$ , where  $f_n$  is the user specific transition function,  $X_{n,t}$  is an exogenous variable with information such as commodity price at time  $t$  and other observable user characteristics (e.g., contract size  $C_n$ ), and  $\epsilon_{n,t}$  is an i.i.d. zero mean Gaussian noise variable. We assume that the exogenous state evolves in a Marko-

vian manner. Given sufficient historical inventory usage data for each user, we can infer a posterior on  $f_n$  for each user  $n$  separately (this is TGP).

While TGP can capture rich user behavior patterns, it is relatively less interpretable because the relationship between the previous inventory level (and other inputs) and the next inventory level can turn out be highly nonlinear, and using the corresponding posterior belief in downstream overbooking decisions may become cumbersome. To alleviate this, we enhance the interpretability by incorporating findings from prior literature (Charnes et al., 1966). In particular, it is known that a user (e.g., a merchant operator) makes injection-withdrawal decisions using a two threshold decision rule structure (also called a double base-stock policy) under reasonable assumptions on the stochasticity of the exogenous variable  $X_{n,t}$ :

$$I_{n,t+1} - I_{n,t} = \begin{cases} \min\{f_{n,t}^a(I_{n,t}, X_{n,t}) - I_{n,t}, G\} & \text{if } I_{n,t} \leq f_{n,t}^a(I_{n,t}, X_{n,t}) \\ 0 & \text{if } f_{n,t}^b(I_{n,t}, X_{n,t}) > I_{n,t} > f_{n,t}^a(I_{n,t}, X_{n,t}) \\ -\min\{I_{n,t} - f_{n,t}^b(I_{n,t}, X_{n,t}), G\} & \text{if } f_{n,t}^b(I_{n,t}, X_{n,t}) \leq I_{n,t}, \end{cases}$$

where  $f^a, f^b$  are two threshold functions and  $G$  is a known operational parameter. Because this two-threshold structure for user behavior is interpretable (user injects if below a given threshold, withdraws if above another threshold, and holds still in between), we use this to define DRGP as follows:

$$I_{n,t+1} = g(f_n^1, f_n^2, X_{n,t}) = \begin{cases} f_n^1(I_{n,t}, X_{n,t}) & \text{if } I_{n,t} \leq f_n^1(I_{n,t}, X_{n,t}) \\ I_{n,t} & \text{if } f_n^2(I_{n,t}, X_{n,t}) > I_{n,t} > f_n^1(I_{n,t}, X_{n,t}) \\ f_n^2(I_{n,t}, X_{n,t}) & \text{if } f_n^2(I_{n,t}, X_{n,t}) \leq I_{n,t}, \end{cases}$$

where GP beliefs are placed on the threshold functions  $f^1$  and  $f^2$  (with noise terms associated with each function suppressed to ease notation). Note that this composition of two functions  $f^1$  and  $f^2$  is non-invertible.

## 4. Computational Experiments

### 4.1. Data

We use *aggregate* inventory level data ( $\sim 100$  observations over 2 years) provided by a US ethanol storage operator. The aggregate values are log-transformed and split into separate inventory levels for four users based on three different heuristics to simulate different types of injection-withdrawal behavior (see Appendix A). As a result, we obtain three datasets with low, medium, and high variability of injection and withdrawal patterns. We also vary the number of data points across all users,  $T$ , between 200 and 400.

These data sets also include information about the exogenous state vector  $X_{n,t}$  that includes: (i) the lease capacity of each user; (ii) the spot and prompt-week futures prices for ethanol; and (iii) the prompt-week futures prices for corn and natural gas. We obtain price data from Bloomberg.

### 4.2. Inference

At any time step, a user may inject, withdraw, or do nothing. When a user injects, the inventory level  $X_{n,t+1}$  reached as a result of this injection is  $f_1$  at  $X_{n,t}$ , and as a result,

this threshold value is observed but the withdrawal threshold is not. Similarly, if there is a withdrawal action,  $f_2$  at  $X_{n,t}$  is observed while  $f_1$  is not. In other words,  $f_1$  and  $f_2$  are partially observable over time. To avoid handling partial observability during inference, we partition the dataset based on when users inject and withdraw and learn the functions  $f_1$  and  $f_2$ , respectively, on the resulting subsets. When computing posteriors in this manner, the ordering of  $f_1$  and  $f_2$  may not satisfy the condition  $f_1 \leq f_2$  that is implicitly assumed in the DRGP model. To overcome this issue, we train a classifier to first predict if a user’s decision is either injection or withdrawal and then employ the corresponding threshold to determine the next stage inventory level.

We compute posteriors on  $f_1$  and  $f_2$  using sparse (GP) variational inference (Titsias, 2009) with 10 inducing points and an Automatic Relevance Determination kernel (note that while one can also use exact GP regression here as an alternative, we chose the former for future scalability). We use a gradient boosting decision tree based classifier.

Both TGP and DRGP can be combined with transfer learning by assuming a common component across users and a user specific latent variable. We also consider such models and label them TGP-TL and DRGP-TL. Details of these models and their accompanying inference procedures can be found in Appendix B.

### 4.3. Results

In the first experiment, we answer the question (Q2) laid out in the Introduction, which seeks to relate the empirical performance of DRGP when compared to TGP. In order to do so, we perform a training-validation partition of each dataset based on a 70% – 30% split. The training data is then used to obtain the posteriors, for instance on  $f_n^1$  and  $f_n^2$  in the case of DRGP, for each user  $n = 1, \dots, 4$ . Subsequently these posteriors are used to predict the inventory levels in the validation data. The mean and standard deviation of the root mean squared errors (RMSEs) for TGP and DRGP are displayed in Figure 4.3 for two values of dataset size  $T$ . When  $T$  equals 200, the RMSE of DRGP is smaller than TGP across all datasets. As the number of data points  $T$  is increased to 400, this trend continues to be true for datasets 1 and 3 but is reversed for dataset 2. Overall, we can conclude that DRGP obtains a lower RMSE than TGP in most cases, while also buying us interpretability.

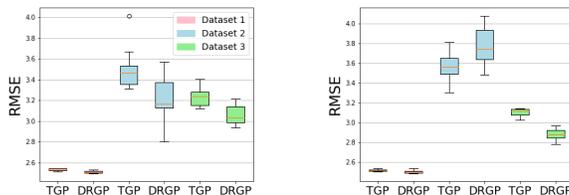


Figure 1: Out of sample RMSE for TGP and DRGP with  $T = 200$  (left) and  $T = 400$  (right). In each panel, there are 6 box plots corresponding to three pairs of datasets in order.

In the second experiment, we investigate the value of transfer learning (where users share common priors). Figure 2(a-c) compares the models with and without transfer learning ( $T =$

400). We observe that incorporating transfer learning produces mixed results, suggesting that these datasets may lack a common user behavioral pattern that can be exploited. Further, in figure 3 (a-d), we illustrate the quality of one-step predictions of the transfer learning models (for all users in Dataset 3, with  $T = 400$ ) as a function of one of the exogenous variables (spot price) in the validation data. We observe that DRGP-TL can predict the out-of-sample log-inventory levels with higher accuracy and low uncertainty when compared to TGP-TL.

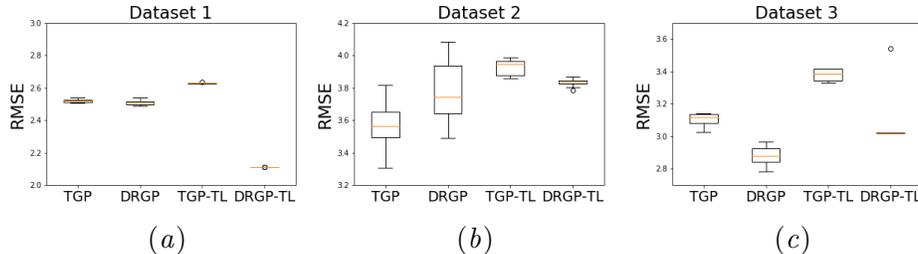


Figure 2: Predictive performance with & without transfer learning across the three datasets.

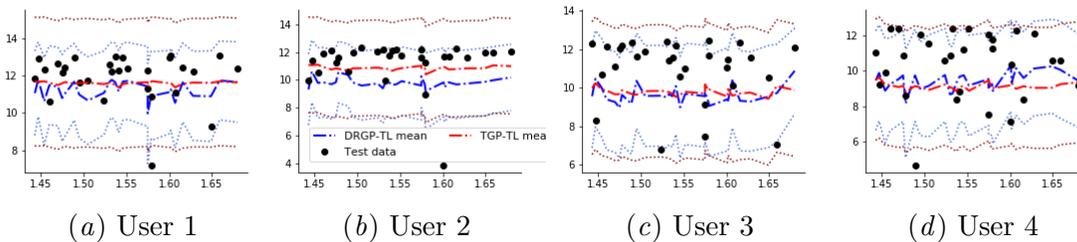


Figure 3: One step prediction of  $\log(\text{inventory})$  level for dataset 3. The means are shown in solid, and the standard deviations around them are shown using dotted curves.

### 5. Future Research

Our initial study of DRGPs shows that there is promise in leveraging decision rules to define non-linear transformations of GPs for user modeling in the ethanol storage application. Extending this investigation to other real-world applications and developing inference procedures tailored to DRGP would be valuable. For instance, we are in the process of developing an inference procedure that directly handles the partial observability of the thresholds and thus benefits from the full dataset, as opposed to partitioning the dataset as we did in this paper. Other research directions being explored include: (i) the interplay between the structure of decision rules in a class of applications, their interpretability, and how this can be leveraged within inference procedures for DRGPs; and (ii) robust inference techniques for DRGP, where parameters are computed by optimizing a metric other than the (exact/approximate) likelihood function.

## References

- Mark A Beaumont. Approximate bayesian computation. *Annual Review of Statistics and its Application*, 6:379–403, 2019.
- Dimitris Bertsimas and Jack Dunn. Optimal classification trees. *Machine Learning*, 106(7):1039–1082, 2017.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Abraham Charnes, Jacques Dreze, and Merton Miller. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica: Journal of the Econometric Society*, pages 307–330, 1966.
- Zhenwen Dai, Mauricio Alvarez, and Neil Lawrence. Efficient modeling of latent information in supervised learning using gaussian processes. In *Advances in Neural Information Processing Systems*, pages 5131–5139, 2017.
- Andreas Damianou and Neil D Lawrence. Semi-described and semi-supervised learning with gaussian processes. *arXiv preprint arXiv:1509.01168*, 2015.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- Miguel Lázaro-Gredilla. Bayesian warped gaussian processes. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1619–1627. Curran Associates, Inc., 2012.
- Miguel Lázaro-Gredilla and Michalis K Titsias. Variational heteroscedastic gaussian process regression. In *International Conference on Machine Learning*, pages 841–848, 2011.
- Benjamin Letham, Cynthia Rudin, Tyler H McCormick, David Madigan, et al. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics*, 9(3):1350–1371, 2015.
- Pablo Moreno-Muñoz, Antonio Artés, and Mauricio Álvarez. Heterogeneous multi-output gaussian process prediction. In *Advances in Neural Information Processing Systems*, pages 6711–6720, 2018.
- Steindór Sæmundsson, Katja Hofmann, and Marc Peter Deisenroth. Meta reinforcement learning with latent variable gaussian processes. *arXiv preprint arXiv:1803.07551*, 2018.
- Alan D Saul, James Hensman, Aki Vehtari, and Neil D Lawrence. Chained gaussian processes. In *Artificial Intelligence and Statistics*, pages 1431–1440, 2016.
- Nicola Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467, 2010.

- Edward Snelson, Zoubin Ghahramani, and Carl E Rasmussen. Warped gaussian processes. In *Advances in Neural Information Processing Systems*, pages 337–344, 2004.
- Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. volume 5 of *Proceedings of Machine Learning Research*, pages 567–574. PMLR, 16–18 Apr 2009.
- Felipe Tobar and Gonzalo Rios. Compositionally-warped gaussian processes. *arXiv preprint arXiv:1906.09665*, 2019.
- Jarno Vanhatalo, Jaakko Riihimäki, Jouni Hartikainen, Pasi Jylänki, Ville Tolvanen, and Aki Vehtari. Gpstuff: Bayesian modeling with gaussian processes. *Journal of Machine Learning Research*, 14(Apr):1175–1179, 2013.
- Chunyi Wang and Radford M Neal. Gaussian process regression with heteroscedastic or non-gaussian residuals. *arXiv preprint arXiv:1212.6246*, 2012.

## Appendix A. Datasets

The ethanol application dataset contains the daily aggregate inventory level of a storage tank in the US, and the daily price of ethanol over a period of two years. We consider weekly inventory levels to model the behavior of users, as suggested by practitioners. There were 39 companies signed up in the system, with various contract sizes. Assuming that users cluster into groups that have similar injection and withdrawal patterns, we created four users (essentially user types/groups) by assigned these companies to each group based on their contract size. We break down the aggregate inventory levels to four user levels based on three heuristics. These three heuristics are designed to test the performance of the four approaches we have for user modeling; and they capture low, medium, and high variance of injection-withdrawal patterns of users.

The first dataset is created by assigning fractions of the aggregate inventory to each user proportional to their contracted capacity, and simulates a system where the users have low variability. In the second dataset, we simulate a setting where users have medium variability when interacting with the system. This is captured by ensuring that the users do not change their inventory levels with probability 0.5, and change their inventory levels randomly between 0 and their rented capacity, again with probability 0.5. Finally, to simulate a system where users interact with the system with high volatility, we make the users change their inventory level randomly from 0 to their rented capacity in every period, such that the aggregate of these individual inventory levels is equal to the aggregate inventory level.

## Appendix B. Approximations for Bayesian Inference

**Sparse GP for Scalability:** For TGP and DRGP, we rely on variational sparse Gaussian process based inference procedure [Titsias \(2009\)](#); [Hensman et al. \(2013\)](#). Inducing

point methods involve introducing  $M \ll T$  *inducing points* at locations  $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^M$  with corresponding function values given by  $\mathbf{u}_i = f(\mathbf{z}_i)$  such that:

$$p(\mathbf{f}|\{I_{n,t}, X_{n,t}\}_{t=1}^T, \mathbf{Z}) = \int p(\mathbf{f}|\{I_{n,t}, X_{n,t}\}_{t=1}^T, \mathbf{u})p(\mathbf{u}|\mathbf{Z})d\mathbf{u},$$

where  $\mathbf{f}$  is the vector of function evaluations at the  $T$  observation points. Using this approach, we are able to approximate the posterior GP with a variational distribution that only depends on the inducing points by obtaining a lower bound on the marginal likelihood.

**Transfer Learning:** TGP-TL modifies TGP by assuming that user specific latent variables and a common target function together drive the inventory updates of all users (Wang and Neal, 2012; Damianou and Lawrence, 2015; Dai et al., 2017). That is,  $I_{n,t+1} = f(I_{n,t}, X_{n,t}, \gamma_n) + \epsilon_{n,t}$ , where  $f$  is a common target function across users that maps the triple  $(I_{n,t}, X_{n,t}, \gamma_n)$  to the next inventory level  $I_{n,t+1}$ , and  $\gamma_n$  is a user specific latent variable. We can jointly infer a posterior belief on  $f$  (we fix this to be a GP) and  $\gamma_n$  (which we take to be Gaussian distributed) using *LVMOGP* Dai et al. (2017). Common temporal patterns of all users can now be captured by  $f$ , while idiosyncratic aspects of each user can be captured using  $\gamma_n$ . Similarly, DRGP can be extended to DRGP-TL, where we have common threshold functions for all users and user-specific latent variables to capture user heterogeneity. The graphical models for TGP-TL and DRGP-TL are illustrated in Figure 4.

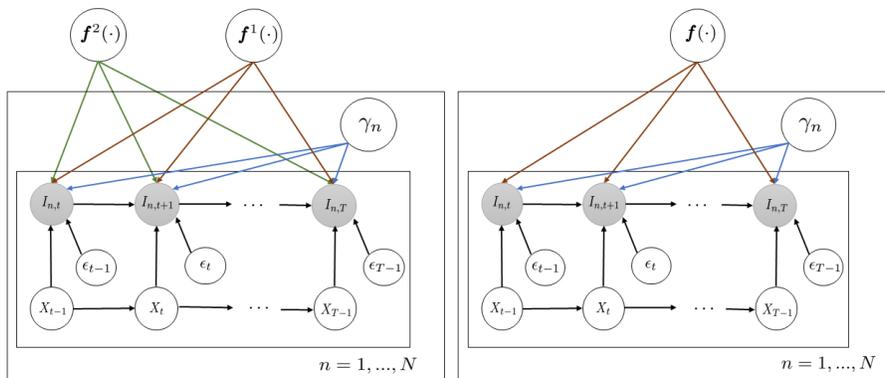


Figure 4: The DRGP-TL and TGP-TL models capturing common user behaviors.

The inference procedure for transfer learning extensions of TGP-TL and DRGP-TL involves handling the joint distribution with respect to the latent variables  $\Gamma = \{\gamma_1, \dots, \gamma_N\}$  and the common function (two functions in DRGP-TL). The following independence assumption is made in the variational approximation for tractability:

$$p(\mathbf{f}, \Gamma) \approx q(\mathbf{f})q(\Gamma),$$

where  $q(\mathbf{f})$  is a GP and  $q(\Gamma) = \prod_{n=1}^N \mathcal{N}(\gamma_n | \mu_n, \Sigma_n)$ . Below, we show the evidence lower bound (ELBO) in TGP-TL:

$$\begin{aligned} \log p(\{\{I_{n,t}\}_{t=1}^T\}_{n=1}^N) &\geq \mathbb{E}_{p(\Gamma, \mathbf{f})} \left[ \log \frac{p(\{\{I_{n,t}\}_{t=1}^T\}_{n=1}^N, \Gamma, \mathbf{f})}{p(\Gamma, \mathbf{f})} \right] \\ &\geq \sum_{n=1}^N \sum_{t=1}^T \mathbb{E}_{q(\mathbf{f})} [\log p(I_{n,t} | \mathbf{f})] \\ &\quad - KL(q(\mathbf{u}) || p(\mathbf{u})) - KL(q(\Gamma) || p(\Gamma)), \end{aligned}$$

where we use Figure 4 in the second inequality. Following Hoffman et al. (2013) and other prior works, we maximize the evidence lower bound (ELBO) which provides a lower bound for the log-marginal likelihood of observed data, and jointly optimize with respect to the model hyper-parameters and the variational parameters as suggested in Sæmundsson et al. (2018).