
Doubly Stochastic Variational Inference for Neural Processes with Hierarchical Latent Variables

Qi Wang¹ Herke van Hoof¹

Abstract

Neural processes (NPs) constitute a family of variational approximate models for stochastic processes with promising properties in computational efficiency and uncertainty quantification. These processes use neural networks with latent variable inputs to induce predictive distributions. However, the expressiveness of vanilla NPs is limited as they only use a global latent variable, while target-specific local variation may be crucial sometimes. To address this challenge, we investigate NPs systematically and present a new variant of NP model that we call Doubly Stochastic Variational Neural Process (DSVNP). This model combines the global latent variable and local latent variables for prediction. We evaluate this model in several experiments, and our results demonstrate competitive prediction performance in multi-output regression and uncertainty estimation in classification.

1. Introduction

In recent decades, increasingly attention has been focused on deep neural networks, and the success of deep learning in computer vision, natural language processing and robotics control etc. can be attributed to the great potential of function approximation with high-capacity models (LeCun et al., 2015). Despite this, there still remain some limitations which incur doubts from industry when applying these models to real world scenarios. Among them, uncertainty quantification is long-standing and challenging, and instead of point estimates we prefer probabilistic estimates with meaningful confidence values in predictions. With uncertainty estimates at hand, we can relieve some risk and make relatively conservative choices in cost-sensitive

¹Amsterdam Machine Learning Lab, University of Amsterdam, Amsterdam, the Netherlands. Correspondence to: Qi Wang <q.wang3@uva.nl>.

decision-making (Gal & Ghahramani, 2016).

Faced with such reality, Bayesian statistics provides a plausible schema to reason about subjective uncertainty and stochasticity, and marrying deep neural networks and Bayesian approaches together satisfies practical demands. Traditionally, Gaussian processes (GPs) (Rasmussen, 2003) as typical non-parametric models can be used to handle uncertainties by placing Gaussian priors over functions. The advantage of introducing distributions over functions lies in the characterization of the underlying uncertainties from observations, enabling more reliable and flexible decision-making. For example, the uncertainty-aware dynamics model enjoys popularity in model-based reinforcement learning, and GPs deployed in PILCO enable propagation of uncertainty in forecasting future states (Deisenroth & Rasmussen, 2011). Another specific instance can be found in demonstration learning; higher uncertainty in prediction would suggest the learning system to query new observations to avoid dangerous behaviors (Thakur et al., 2019). During the past few years, a variety of models inspired by GPs and deep neural networks have been proposed (Salimbeni & Deisenroth, 2017; Snelson & Ghahramani, 2006; Titsias, 2009; Titsias & Lawrence, 2010).

However, GP induced predictive distributions are met with some concerns. One is high computational complexity in prediction due to the matrix inversion, and another is less flexibility in function space. Recognized as an explicit stochastic process model, the vanilla GP strongly depends on the assumption that the joint distribution is Gaussian, and such a unimodal property makes it tough to scale to more complicated cases. These issues facilitate the birth of adaptations or approximate variants for GP related models (Garnelo et al., 2018a;b; Louizos et al., 2019), which incorporate latent variables in modeling to account for uncertainties. Among them, Neural Processes (NPs) are relatively representative with advantages like uncertainty-aware prediction and efficient computations.

In this paper, we investigate NP related models and explore more expressive approximations towards general stochastic processes (*SPs*). *The main focus is on a novel variational approximate model for NPs to solve learning problems in high-dimensional cases.* To improve flexibility in predictive

distributions, hierarchical latent variables are considered as part of the model structure. Our primary contributions can be summarized as follows.

- We systematically revisit NPs, \mathcal{SP} s and other related models from a unified perspective with an implicit Latent Variable Model (LVM). Both GPs and NPs are studied in this hierarchical LVM.
- The Doubly Stochastic Variational Neural Process (DSVNP) is proposed to enrich the NP family, where both global and target specific local latent variables are involved in a predictive distribution.
- Experimental results demonstrate effectiveness of the proposed Bayesian model in high dimensional domains, including regression with multiple outputs and uncertainty-aware image classification.

2. Deep Latent Variable Model as Stochastic Processes

Generally, a stochastic process places a distribution over functions and any finite collections of variables can be associated with an implicit probability distribution. Here, we naturally formulate an implicit LVM (Rezende et al., 2014; Kingma & Welling, 2014) to characterize General Stochastic Function Processes (GSFPs). The conceptual generation paradigm for this LVM can be depicted in the following equations,

$$z_i = \phi(x_i) + \epsilon(x_i) \quad (1)$$

$$y_i = \varphi(x_i, z_i) + \zeta_i \quad (2)$$

where terms ϵ and ζ respectively indicate the stochastic component in the latent space and random noise in observations. To avoid ambiguity in notation, the stochastic term ϵ is declared as an index dependent random variable $\epsilon(x_i)$, and ζ_i is observation noise in the environment. Also, the transformations ϕ and φ are assumed to be Borel measurable, and the latent variables in Eq. (1) are not restricted. *Note that they can be some set of random variables with statistical correlations without loss of generality.* When Kolmogorov Extension Theorem (Oksendal, 2013) is satisfied for $\epsilon(x_i)$, a latent \mathcal{SP} can be induced. Eq. (1) decomposes the process into a deterministic component and a stochastic component in some models. The transformation φ in Eq. (2) is directly connected to the output. Such a generative process can simultaneously inject *aleatoric uncertainty* and *epistemic uncertainty* in modelling (Hofer et al., 2002), but inherent correlations in examples make the exact inference intractable mostly.

Another principal issue is about prediction with permutation invariance, which learns a conditional distribution in \mathcal{SP} models. With the context $\mathcal{C} = \{(x_i, y_i) | i = 1, 2, \dots, N\}$

and input variables of the target x_T , we seek a stochastic function f_θ mapping from X to Y and formalize the distribution as $p_\theta(y_T | x_C, y_C, x_T)$ ¹ invariant to the order of context observations. The definitions about permutation invariant functions (PIFs) and permutation equivariant functions (PEFs) are included in Appendix A.

2.1. Gaussian Processes in the Implicit LVM

Let us consider a trivial case in the LVM when the operation φ is an identity map, ζ is Gaussian white noise, and the latent layer follows a multivariate Gaussian distribution. This degenerated case indicates a GP, and terms ϕ , ϵ are respectively the mean function and the zero-mean GP prior. Meanwhile, recall that the prediction at target input x_T in GPs relies on a predictive distribution $p(y_T | x_C, y_C, x_T)$, where the mean and covariance matrix are inferred from the context $[x_C, y_C]$ and target input x_T .

$$\begin{aligned} \mu(x_T; x_C, y_C) &= \phi_\theta(x_T) + \Sigma_{T,C} \Sigma_{C,C}^{-1} (y_C - \phi_\theta(x_C)) \\ \Sigma(x_T; x_C, y_C) &= \Sigma_{T,T} - \Sigma_{T,C} \Sigma_{C,C}^{-1} \Sigma_{C,T} \end{aligned} \quad (3)$$

Here ϕ_θ and Σ in Eq. (3) are vectors of mean functions and covariance matrices. For additive terms, they embed context statistics and connect them to the target sample x_T . Furthermore, two propositions are drawn, which we prove in Appendix B.

Proposition 1. The statistics of GP predictive distributions, such as mean and (co)-variances, for a specific point x_* are PIFs, while those in $p(y_T | x_C, y_C, x_T)$ are PEFs.

2.2. Neural Processes in the Implicit LVM

In non-GP scenarios, inference and prediction processes for the LVM can be non-trivial, and NPs are the family of approximate models for implicit \mathcal{SP} s. Also, relationship between GPs and NPs can be explicitly established with deep kernel network (Rudner et al., 2018). Note that *NPs translate some properties of GPs to predictive distributions, especially permutation invariance of context statistics*, which is highlighted in **Proposition 1**. Here three typical models are investigated, respectively conditional neural process (CNP) (Garnelo et al., 2018a), vanilla NP (Garnelo et al., 2018b) and attentive neural process (AttnNP) (Kim et al., 2019).

When approximate inference is used in NP family with Latent Variables, a preliminary evidence lower bound (ELBO) for the training process can be derived, which aims at pre-

¹For brief notations, the inputs, outputs of the context and the target are respectively denoted as $x_C = x_{1:N}$, $y_C = y_{1:N}$, $x_T = x_{1:N+M}$, $y_T = y_{1:N+M}$. Only in CNP, $x_T = x_{N+1:N+M}$, $y_T = y_{N+1:N+M}$. And $[x_*, y_*]$ refers to any instance in the target.

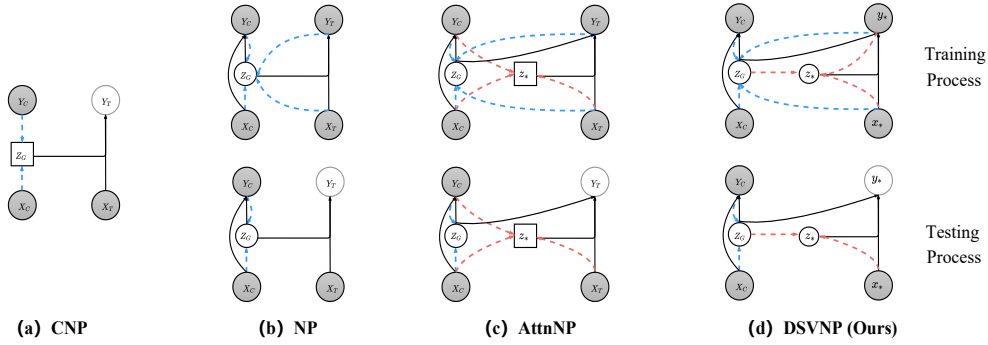


Figure 1. Probabilistic Graphs for CNP, vanilla NP, Attentive NP and DSVNP. The blue dotted lines characterize the inference towards global latent variable z_G , while the pink dotted lines are for target specific local latent variables z_* . The ones in first row are training cases, while those in second row are testing cases.

dictive distributions for most NP related models.

$$\begin{aligned} \ln [p(y_T|x_C, y_C, x_T)] &\geq \mathbb{E}_{q_\phi(z_T)} \ln [p_\theta(y_T|x_T, z_T)] \\ -D_{KL}[q_\phi(z_T|x_C, y_C, x_T, y_T) \parallel p(z_T|x_C, y_C, x_T)] \end{aligned} \quad (4)$$

To ensure context information invariant to orders of points, CNP embeds the context point $[x_C, y_C]$ in an elementwise way and then aggregates them with a permutation invariant operation \oplus , such as mean or max pooling.

$$r_i = h_\theta(x_i, y_i), \quad r_C = \bigoplus_{i=1}^N r_i \quad (5)$$

The latent variable in CNP is a deterministic embedding of the form $p_\theta(z_C|x_C, y_C) = r_C(x_C, y_C)$. Followed with Eq. (2), CNP decodes statistics as the mean and the variance for a predictive distribution.

For vanilla NPs, the encoder structure resembles that of CNPs, and the learned embedding variable r in Eq. (5) is no longer a function but a Gaussian variable after amortized transformations. For the graphical structure of this LVM in vanilla NPs (Refer to Figure 1 (b)), all latent variables are degraded to a global Gaussian latent variable, which accounts for the consistency.

AttnNP further improves the expressiveness of context information in NP, leaving the latent variable as the combination of a global variable and a local variable. Especially, the attention network uses self-attention or dot-product attention to enable transformations of context points and the extraction of hierarchical dependencies between context points and target points. For the graphical model of this LVM in AttnNP, the context information is instance-specific. The latent variable of AttnNP in Eq. (6) is the concatenation of attention embedding z_{attn} from element-wise context embedding s_i and a global latent variable z_G drawn from

an amortized distribution parameterized with Eq. (5).

$$z_{attn} = \bigoplus_{i=1}^N w(x_i, x_*) s_i, \quad z = [z_{attn}, z_G] \quad (6)$$

As a summary, AttnNP boosts performance with attention networks, which implicitly seeks more flexible functional translations for each target.

2.3. Connection to Other Models

In some scenarios, when the latent layer in Eq. (1) is specified as a Markovian chain, the LVM degrades to classical state space model. If random variables in the latent layer of the LVM are independent, the resulted neural network is similar to the conditional variational auto-encoder (Sohn et al., 2015) and no context information is utilized for prediction. Instead, the existence of correlations between latent variables in the hidden layer increases the model capacity. The induced \mathcal{SP} in Eq. (2) is a warped GP when the latent \mathcal{SP} is a GP and the transformation φ is nonlinear monotonic (Snelson et al., 2004). In addition, several previous works integrate this idea in modelling as well, and representative examples are deep GPs (Dai et al., 2016) and hierarchical GPs (Tran et al., 2016).

3. Neural Process under Doubly Stochastic Variational Inference

In the last section, we gain more insights about mechanism of GPs and NPs and disentangle these models with the implicit LVM. A conclusion can be drawn that the posterior inference conditioned on the context requires both *approximate distributions with permutation invariance* and *some bridge to connect observations and the target in latent space*. Note that the global induced latent variable may be insufficient to describe dependencies, and critical challenge comes from *non-stationarity* and *locality*, which are even crucial

Table 1. Structure Summary over NP Related Models on Training Dataset. Here f corresponds to some functional operation. Global latent variable in CNP only governs points to predict, while that in NP works for the whole points. And local latent variables in AttnNP and DSVNP are distinguished, with the latter as a latent random variable.

NP FAMILY	RECOGNITION MODEL	GENERATIVE MODEL	PRIOR DISTRIBUTION	LATENT VARIABLE
CNP	$z_C = f(x_C, y_C)$	$p(y_T z_C, x_T)$	NULL	GLOBAL
NP	$q(z_G x_C, y_C, x_T, y_T)$	$p(y_T z_G, x_T)$	$p(z_G x_C, y_C)$	GLOBAL
ATTNP	$q(z_G x_C, y_C, x_T, y_T),$ $z_* = f(x_C, y_C, x_*)$	$p(y_* z_G, z_*, x_*)$	$p(z_G x_C, y_C)$	GLOBAL +LOCAL
DSVNP (OURS)	$q(z_G x_C, y_C, x_T, y_T),$ $q(z_* z_G, x_*, y_*)$	$p(y_* z_G, z_*, x_*)$	$p(z_G x_C, y_C),$ $p(z_* z_G, x_*)$	GLOBAL +LOCAL

in high-dimensional cases.

Hence, we present a hierarchical way to modify NPs, and the trick is to involve auxiliary latent variables for NPs and derive a new evidence lower bound for different levels of random variables with doubly stochastic variational inference (Salimbeni & Deisenroth, 2017; Titsias & Lázaro-Gredilla, 2014). The original intention of involving auxiliary latent variables is to improve the expressiveness of approximate posteriors, and it is commonly used in deep generative models (Maaløe et al., 2016). So, as displayed in Table (1), DSVNP considers a global latent variable and a local latent variable to convey context information at different levels. Our work is also consistent with the hierarchical implicit Bayesian neural networks (Tran et al., 2016; 2017), which distinguish the role of latent variables and induce more powerful posteriors. Without exception, the local latent variable z_* refers to any data point (x_*, y_*) for prediction in DSVNP in the remainders of this paper.

3.1. Neural Process with Hierarchical Latent Variables

To extract hierarchical context information for the predictive distribution, we distinguish the global latent variable and the local latent variable in the form of a Bayesian model, and the induced LVM is DSVNP. This variant shares the same prediction structure with AttnNP. The global latent variable is shared across all observations, and the role of context points resembles inducing points in sparse GP (Snelson & Ghahramani, 2006), summarizing general statistics in the task. As for the local latent variable in our proposed DSVNP, it is an auxiliary latent variable responsible mainly for variations of instance locality. From another perspective, DSVNP combines the global latent variable in vanilla NPs with the local latent variable in conditional variational autoencoder (C-VAE). This implementation in model construction separates the global variations and sample specific variations and theoretically increases the expressiveness of the neural network.

As illustrated in Figure 1 (d), the target to predict is governed by these two latent variables. Both the global latent variable z_G and the local latent variable z_* contribute to prediction.

Formally, the generative model as a \mathcal{SP} is described as follows, where exact inferences for latent variables z_G and z_* are infeasible.

$$\rho_{x_{1:N+M}}(y_{1:N+M}) = \iint \prod_{i=1}^{N+M} p(y_i | z_G, z_i, x_i) p(z_i | x_i, z_G) p(z_G) dz_{1:N+M} dz_G \quad (7)$$

Meanwhile, we emphasize that this generation method naturally induces an *exchangeable stochastic process* (Bhattacharya & Waymire, 2009). (The proof is given in Appendix C.)

3.2. Approximate Inference and ELBO

With the relationship between these variables clarified, we can characterize the inference process for DSVNP and then a new ELBO is presented. Distinguished from AttnNP, we need to infer both global and local latent variables with evidence of collected dataset. Posteriors of the global and local latent variables on training dataset are approximated with distributions like vanilla NPs, mapping Eq. (5) to means and variances. And inference towards local latent variables requires target information in the approximate posterior,

$$q_{\phi_{1,1}} = \mathcal{N}(z_G | \mu(x_C, y_C, x_T, y_T), \Sigma(x_C, y_C, x_T, y_T)) \quad (8)$$

$$q_{\phi_{2,1}} = \mathcal{N}(z_* | \mu(z_G, x_*, y_*), \Sigma(z_G, x_*, y_*)) \quad (9)$$

where $q_{\phi_{1,1}}$ and $q_{\phi_{2,1}}$ are approximate posteriors in training process. The generative process is reflected in Eq. (10) for DSVNP, where g_θ indicates a decoder in a neural network.

$$p(y_* | x_C, y_C, x_*) = g_\theta(z_G, z_*, x_*) \quad (10)$$

Consequently, this difference between vanilla NP and DSVNP leads to another ELBO or negative variational free energy \mathcal{L} as the right term,

$$\begin{aligned} \ln [p(y_* | x_C, y_C, x_*)] &\geq \mathbb{E}_{q_{\phi_{1,1}}} \mathbb{E}_{q_{\phi_{2,1}}} \ln [p(y_* | z_G, z_*, x_*)] \\ &\quad - \mathbb{E}_{q_{\phi_{1,1}}} [D_{KL}[q_{\phi_{2,1}}(z_* | z_G, x_*, y_*) \parallel p_{\phi_{2,2}}(z_* | z_G, x_*)]] \\ &\quad - D_{KL}[q_{\phi_{1,1}}(z_G | x_C, y_C, x_T, y_T) \parallel p_{\phi_{1,2}}(z_G | x_C, y_C)] \end{aligned} \quad (11)$$

where $p_{\phi_{1,2}}(z_G|x_C, y_C)$ and $p_{\phi_{2,2}}(z_*|z_G, x_*)$ parameterized with neural networks are used as prior distributions. Here we no longer employ standard normal distributions with zero prior information, and instead these are parameterized with two diagonal Gaussians for the sake of simplicity and learned in an amortized way.

3.3. Scalable Training and Uncertainty-aware Prediction

Based on the inference process in DSVNP and the corresponding ELBO in Eq. (11), the Monte Carlo estimation for the lower bound is derived, in which we wish to maximize,

$$\begin{aligned} \mathcal{L}_{MC} = & \frac{1}{K} \sum_{k=1}^K \left[\frac{1}{S} \sum_{s=1}^S \ln[p(y_*|x_*, z_*^{(s)}, z_G^{(k)})] \right. \\ & - D_{KL}[q(z_*|z_G^{(k)}, x_*, y_*) \parallel p(z_*|z_G^{(k)}, x_*)] \\ & \left. - D_{KL}[q(z_G|x_C, y_C, x_T, y_T) \parallel p(z_G|x_C, y_C)] \right] \end{aligned} \quad (12)$$

where latent variables are sampled as $z_G^{(k)} \sim q_{\phi_{1,1}}(z_G|x_C, y_C)$ and $z_*^{(s)} \sim q_{\phi_{2,1}}(z_*|z_G^{(k)}, x_*, y_*)$. And the resulted Eq. (12) is employed as the objective function in the training process. To reduce variance in sampling, the reparameterization trick (Kingma & Welling, 2014) is used for all approximate distributions and the model is optimized using Stochastic Gradient Variational Bayes (Kingma & Welling, 2014). More details can be found in Algorithm (1).

The predictive distribution is of our interest. For DSVNP, prior networks as $p(z_G|x_C, y_C)$ and $p(z_*|z_G, x_*)$ are involved in prediction, and this leads to the integration over both global and local latent variables here as revealed in Eq. (13).

$$\begin{aligned} p(y_*|x_C, y_C, x_*) = & \iint p(y_*|z_G, z_*, x_*) p_{\phi_{1,2}}(z_G|x_C, y_C) \\ & p_{\phi_{2,2}}(z_*|z_G, x_*) dz_G dz_* \end{aligned} \quad (13)$$

For uncertainty-aware prediction, there exist different approaches for Bayesian neural networks. Generally, once the model is well trained, the conditional distribution in neural networks can be derived. The accuracy can be evaluated through deterministic inference over latent variables, i.e., $\tilde{z}_G = \mathbb{E}[z_G|x_C, y_C]$, $\tilde{z}_* = \mathbb{E}[z_*|z_G, x_*]$, $y_* = \arg \max_y p(y|\tilde{z}_*, \tilde{z}_G, x_*)$.

The Monte Carlo estimation over Eq. (13), which is commonly used for prediction, can be written in the following equation,

$$p(y_*|x_C, y_C, x_*) \approx \frac{1}{KS} \sum_{k=1}^K \sum_{s=1}^S p_{\theta}(y_*|x_*, z_*^{(s)}, z_G^{(k)}) \quad (14)$$

Algorithm 1 Variational Inference for DSVNP in Training.

Input: Dataset \mathcal{D} , Maximum context points N_{max} , Learning rate α , Batch size B and Epoch number m .

Output: Model parameters ϕ_1, ϕ_2 and θ .

Initialize parameters ϕ_1, ϕ_2, θ of recognition model and generative model in Eq. (8), (9) and (10).

for $i = 1$ **to** m **do**

Draw some context number $N_C \sim U[1, N_{max}]$;

Draw mini-batch instances and formulate context-target pairs $\{(x_C, y_C, x_T, y_T)_{bs}\}_{bs=1}^B \sim \mathcal{D}$;

Feedforward the mini-batch instances to recognition model q_{ϕ_1} :

Draw sample of latent variable $z_G \sim q_{\phi_{1,1}}$ in

Eq. (8);

Draw sample of latent variable $z_* \sim q_{\phi_{2,1}}$ in

Eq. (9);

Feedforward the latent variables to discrimination model p_{θ} :

Compute conditional probability distribution in Eq. (10);

Update parameters by Optimizing Eq. (12):

$\phi_1 \leftarrow \phi_1 + \alpha \nabla_{\phi_1} \mathcal{L}_{MC} \triangleright \phi_1 = [\phi_{1,1}, \phi_{1,2}]$

$\phi_2 \leftarrow \phi_2 + \alpha \nabla_{\phi_2} \mathcal{L}_{MC} \triangleright \phi_2 = [\phi_{2,1}, \phi_{2,2}]$

$\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}_{MC}$

end for

where the global and local latent variables are sampled in prior networks through ancestral sampling as $z_G^{(k)} \sim p_{\phi_{1,2}}(z_G|x_C, y_C)$ and $z_*^{(s)} \sim p_{\phi_{2,2}}(z_*|z_G^{(k)}, x_*)$.

3.4. More Insights and Implementation Tricks

The global latent variable and local latent variables govern different variations in prediction and sample generation. This is a part of motivations for AttnNP and DSVNP. Interestingly, the inference for our induced \mathcal{SP} integrates the aspects of vanilla NPs (Eslami et al., 2018) and C-VAEs (Sohn et al., 2015).

Similar to β -VAE (Higgins et al., 2017), we rewrite the right term in Eq. (11) with constraints and these restrict the search for variational distributions. Equivalently, tuning the weights of divergence terms in Eq. (11) leads to varying balance between global and local information.

$$\begin{aligned} & \max_{\phi_1, \phi_2, \theta} \mathbb{E}_{q_{\phi_{1,1}}} \mathbb{E}_{q_{\phi_{2,1}}} \ln[p_{\theta}(y_*|z_G, z_*, x_*)] \\ & D_{KL}[q(z_G|x_C, y_C, x_T, y_T) \parallel p(z_G|x_C, y_C)] < \epsilon_G \\ & \mathbb{E}_{q_{\phi_{1,1}}} [D_{KL}[q(z_*|z_G, x_*, y_*) \parallel p(z_*|z_G, x_*)]] < \epsilon_L \end{aligned} \quad (15)$$

Here, a more practical objective in implementations derived

by weight calibrations in Eq. (16).

$$\begin{aligned} \mathcal{L}_{MC}^W = & \frac{1}{K} \sum_{k=1}^K \left[\frac{1}{S} \sum_{s=1}^S \ln[p(y_* | x_*, z_*^{(s)}, z_G^{(k)})] \right. \\ & - \beta_1 D_{KL}[q(z_* | z_G^{(k)}, x_*, y_*) \parallel p(z_* | z_G^{(k)}, x_*)] \\ & \left. - \beta_2 D_{KL}[q(z_G | x_C, y_C, x_T, y_T) \parallel p(z_G | x_C, y_C)] \right] \end{aligned} \quad (16)$$

Also, training stochastic model with multiple latent variables is non-trivial, and there exist several works about KL divergence term annealing (Sønderby et al., 2016) or dynamically adapting for the weights. Importantly, the target specific KL divergence term is sometimes suggested to assign more penalty to guarantee the consistency between approximate posterior and prior distribution (Kohl et al., 2018; Sohn et al., 2015).

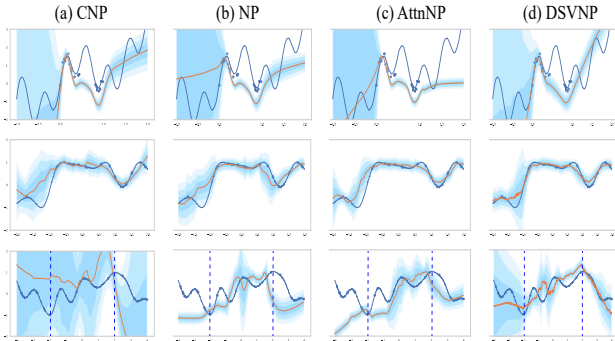


Figure 2. Function Prediction in Interpolation and Extrapolation. Blue curves are ground truth with dotted units as context points, and orange ones are predicted results. Rows from up to down respectively indicate cases: single function with noise, interpolation and extrapolation towards realizations from stochastic process. The shadow regions are ± 3 standard deviations from the mean.

4. Experiments

In this section, we start with learning predictive functions on several toy dataset, and then high-dimensional tasks, including system identification on physics engines, multioutput regression on real-world dataset as well as image classification with uncertainty quantification, are performed to evaluate properties of NP related models. The dot-product attention is implemented in all AttnNPs here. All implementation details are attached in Appendix E.

4.1. Synthetic Experiments

We initially investigate the *epidemic* uncertainty captured by NP related models on a 1-d regression task, and the function (Osband et al., 2016) is characterized as $y = x + \epsilon + \sin(4(x + \epsilon)) + \sin(13(x + \epsilon))$. Observations as the training set include 12 points and 8 points respectively uniformly

Table 2. Average Negative Log-likelihoods over all target points on realizations from Synthetic Stochastic Process. (Figures in brackets are variances.)

PREDICTION	CNP	NP	ATTNNP	DSVNP
INTER	-0.802 (1E-6)	-0.958 (2E-5)	-1.149 (8E-6)	-0.975 (2E-5)
EXTRA	1.764 (1E-1)	8.192 (7E1)	8.091 (7E2)	4.203 (9E0)

drawn from intervals $U[0, 0.6]$ and $U[0.8, 1.0]$, with the noise drawn from $\epsilon \sim N(0, 0.003^2)$. As illustrated in the first row of Fig. (2), we can observe CNP and DSVNP better quantify variance outside the interval $[0, 1.0]$, while AttnNP either overestimates or underestimates the uncertainty to show higher or lower standard deviations in regions with less observations. All models share similar properties with GPs in predictive distributions, displaying lower variances around observed points. As for the gap in interval $[0.6, 0.8]$, the revealed uncertainty is consistent to that in (Sun et al., 2019; Hernández-Lobato & Adams, 2015) with intermediate variances.

Further, we conduct curve fitting tasks in \mathcal{SP} . The \mathcal{SP} initializes with a zero mean Gaussian Process $y^{(0)} \sim \mathcal{GP}(0, k(\cdot, \cdot))$ indexed in the interval $x \in [-2.0, 2.0]$, where the radial basis kernel $k(x, x') = \sigma^2 \exp(-(x - x')^2 / 2l^2)$ is used with $l=1$ 0.4 and σ 1.0. Then the transformation is performed to yield $y = \sin(y^{(0)}(x) + x)$. The training process follows that in NP (Garnelo et al., 2018b). Predicted results are visualized in the second and the third rows of Fig. (2). Note that CNP only predicts points out of the context in default settings. More evidence is reported in Table (2), where 2000 realizations are independently sampled and predicted for both interpolation and extrapolation. After several repetitive observations, we find in terms of the interpolation accuracy, DSVNP works better than vanilla NP but the improvement is not as significant as that in AttnNP, which is also verified in visualizations. All (C)NPs show higher uncertainties around index 0, where less context points are located, and variances are relatively close in other regions. For extrapolation results, since all models are trained in the dotted column lines restricted regions, it is tough to scale to regions out of training interval and all negative log-likelihoods (NLLs) are higher. When there exist many context points located outside the interval, the learned context variable may deteriorate predictions for all (C)NPs, and observations confirm findings in (Gordon et al., 2020). Interestingly, DSVNP tends to overestimate uncertainties out of the training interval but predicted extrapolation results mostly fall into the one σ confident region, this property is similar to CNP. On the other hand, vanilla NP and AttnNP tend to underestimate the uncertainty sometimes.

4.2. System Identification on Physics Engines

Capturing dynamics in systems is crucial in control related problems, and we extend synthetic experiments on a classical simulator, Cart-Pole systems, which is detailed in (Gal et al., 2016). As shown in Fig. (3), the original intention is to conduct actions to reach the goal with the end of a pole, but here we focus on dynamics and the state is a vector of the location, the angle and their first-order derivatives. Specifically, the aim is to forecast the transited state $[x_c, \theta, x'_c, \theta']$ in time step $t + 1$ based on the input as a state action pair $[x_c, \theta, x'_c, \theta', a]$ in time step t . To generate a variety of trajectories under a random policy for this experiment, the mass m_c and the ground friction coefficient f_c are varied in the discrete choices $m_c \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$ and $f_c \in \{0.06, 0.08, 0.1, 0.12\}$. Each pair of $[m_c, f_c]$ values specifies a dynamics environment, and we formulate all pairs of $m_c \in \{0.3, 0.5, 0.7\}$ and $f_c \in \{0.08, 0.12\}$ as training environments with the rest 16 pairs of configurations as the testing environments.




Figure 3. Cart-Pole Dynamical Systems. The cart and the pole are with masses m_c and m_p , and the length of the pole is l . And the configuration of the simulator is up to parameters of the cart-pole mass and the ground friction coefficient here with other hyper-parameters fixed in this experiment.

For each configuration of the simulator including training and testing environments, we sample 400 trajectories of horizon as 10 steps using a random controller, and more details refer to Appendix E. The training process follows Algorithm (1) with the maximum number of context points as 100. During the testing process, 100 state transition pairs are randomly selected for each configuration of the environment, working as the maximum context points to identify the configuration of dynamics. And the collected results are reported in Table (3), where prediction performance on 5600 trajectories from 14 configurations of environments are revealed. As can be seen, the negative log-likelihood values are not consistent with those of mean square errors, and DSVNP shows both better uncertainty quantification with lowest NLLs and approximation errors in MSEs. AttnNP improves NP in both metrics, while CNP shows relatively better NLLs but the approximation error is a bit higher than others.

Table 3. Predictive Negative Log-Likelihoods and Mean Square Errors on Cart-Pole State Transition Testing Dataset. (Figures in brackets are variances.)

METRICS	CNP	NP	ATTNPNP	DSVNP
NLL	-2.014 (9E-4)	-1.537 (1E-3)	-1.821 (7E-3)	-2.145 (9E-4)
MSE	0.096 (3E-4)	0.074 (2E-4)	0.067 (1E-4)	0.036 (2.1E-5)

4.3. Multi-Output Regression on Real-world Dataset

Further, more complicated scenarios are considered when the regression task relates to multiple outputs. As investigated in (Moreno-Muñoz et al., 2018; Bonilla et al., 2008), distributions of output variables are implicit, which means no explicit distributions are appropriate to be used in parameterizing the output. We evaluate the performance of all models on dataset, including SARCOS², Water Quality (WQ) (Džeroski et al., 2000) and SCM20d (Spyromitros-Xioufis et al., 2016). Details about these dataset and neural architectures for all models are included in Appendix E. Furthermore, Monte-Carlo Dropout is included for comparisons. Similar to NP (Garnelo et al., 2018b), the variance parameter is not learned and the objective in optimization is pointwise mean square errors (MSEs) after averaging all dimensions in the output. Each dataset is randomly split into 2-folds as training and testing sets. The training procedure in (C)NPs follows that in Algorithm (1), and some context points are randomly selected in batch samples. In the testing stage, we randomly select 30 instances as the context and then perform predictions with (C)NPs. The weights of data likelihood and KL divergence terms in models are not tuned here.

During training, ELBOs in NP related models are optimized, while MSEs are used as evaluation metric in testing (Dez-fouli & Bonilla, 2015)³. The predictive results on testing dataset are reported in Table (4). All MSEs are averaged after 10 independent experiments. We observe DSVNP outperforms other models, and deterministic context information in CNP hardly increases performance. Compared with NP models, MC-NN is relatively satisfying on Sarcos and WQ, and AttnNP works not well in these cases. A potential reason can be that deterministic context embedding with dot product attention is less predictive for output with multiple dimensions, while the role of local latent variable in DSVNP not only bridges the gap between input and output, but also extracts some correlation information among variables in outputs. As a comparison to synthetic experiments, the attention mechanism is more suitable to extract

²<http://www.gaussianprocess.org/gpml/data/>

³Directly optimizing Gaussian log-likelihoods does harm to performance based on experimental results.

Table 4. Predictive MSEs on Multi-Output Dataset. CNP’s results are for target points. D records (input,output) dimensions, and N is the number of samples. MC-Dropout runs 50 stochastic forward propagation and average results for prediction in each data point. (Figures in brackets are variances.)

DATASET	N	D	MC-DROPOUT	CNP	NP	ATTNP	DSVNP
SARCOS	48933	(21,7)	1.215(3E-3)	1.437(2.9E-2)	1.285(1.2E-1)	1.362(8.4E-2)	0.839(1.5E-2)
WQ	1060	(16,14)	0.007(9.6E-8)	0.015(2.4E-5)	0.007(5.2E-6)	0.01(8.5E-6)	0.006(1.6E-6)
SCM20D	8966	(61,16)	0.017(2.4E-7)	0.037(6.7E-5)	0.015(7.1E-8)	0.015(8.1E-7)	0.007(2.3E-7)

local information when the output dimension is lower.

4.4. Classification with Uncertainty Quantification

Here image classification is performed with NP models and MC-Dropout, and out of distribution (o.o.d.) detection is chosen to measure the goodness of uncertainty quantification. We respectively train models on MNIST and CIFAR10. The dimensions for latent variables are 64 on MNIST and 128 on CIFAR10. The training process for NP related models follows Algorithm (1), with the number of context images randomly selected in each batch update. For the testing process, we randomly select 100 instances from the domain dataset as the context for (C)NP models. The commonly used measure for uncertainty in K-class o.o.d. detection is entropy (Lakshminarayanan et al., 2017), $\mathbb{H}[y^*|x^*] = -\sum_{c=1}^K p_w(y_c^*|x^*; \mathcal{D}_{tr}) \ln p_w(y_c^*|x^*; \mathcal{D}_{tr})$, where data point (x^*, y^*) comes from either domain test dataset \mathcal{D}_{te} or o.o.d. dataset \mathcal{D}_{ood} .

For classification performance with NP related models, we observe the difference is extremely tiny on MNIST with all accuracies around 99%, while on CIFAR10 DSVNP beats all baselines with highest accuracy 86.3% and lowest in-distribution entropies (Refer to Table (2) in Appendix E). The involvement of a deterministic path does not improve much, and in contrast, MC-Dropout and CNP achieve intermediate performance. A possible cause can be implicit kernel information captured by attention network in images is imprecise. The cumulative distributions about predictive entropies are reported in Fig. (4). For models trained on MNIST, we observe no significant difference on domain dataset, but DSVNP achieves best results on FMNIST/KMNIST and MC-Dropout performs superior on Uniform/Gaussian noise dataset. Interestingly, AttnNP tends to underestimate uncertainty on FMNIST/KMNIST and the measure is close to the neural network without dropout. Those trained on CIFAR10 are different from observations in the second row of Fig. (4). It can be noticed DSVNP shows lowest uncertainty on domain dataset (CIFAR10) and medium uncertainty on SVHN/Gaussian/Uniform Dataset. MC-Dropout and AttnNP seem to not work so well overall, but CNP well measures uncertainty on Gaussian/Uniform dataset. Results again verify SVHN as tough dataset for the task (Nalisnick et al., 2019). Also note that entropy distri-

butions on Rademacher Dataset are akin to that on domain dataset, which means the Rademacher noise is more risky for CIFAR10 classification, and DSVNP is a better choice to avoid such adversarial attack in this case. Those evidences tell us that the deterministic path in AttnNP does not boost classification performance on domain dataset but weakens the ability of o.o.d. detection mostly, while local latent variables in DSVNP improve both performance. Maybe deterministic local latent variables require more practical attention information, but here only dot-product attention information is included. As a comparison, the local latent variable in DSVNP captures some target specific information during the training process and improves detection performance with it.

5. Related Works

Scalability and Expressiveness in Stochastic Process.

GPs are the most well known member of \mathcal{SP} s family and have inspired a lot of extensions, such as deep kernel learning (Wilson et al., 2016a;b) and sparse GPs (Snelson & Ghahramani, 2006) with better scalability. Especially, the latter incorporated sparse prior in function distribution and utilized a small proportion of observations in predictions. In multi-task cases, several GP variants were proposed (Moreno-Muñoz et al., 2018; Bonilla et al., 2008; Zhao & Sun, 2016). Other works also achieve sparse effect but with variational inference, approximating the posterior in GPs and optimizing ELBO (Hensman et al., 2015; Salimbeni et al., 2019; Titsias & Lawrence, 2010). Another branch is about directly capturing uncertainties with deep neural networks, which is revealed in NP related models. Other extensions include generative query network (Eslami et al., 2018), sequential NP (Singh et al., 2019) and convolutional conditional NP (Gordon et al., 2020). Variational implicit process (Ma et al., 2019) targeted at more general \mathcal{SP} s and utilized GPs in latent space as approximation. Sun et al. proposed functional variational Bayesian neural networks (Sun et al., 2019), and variational distribution over functions of measurement set was used to represent \mathcal{SP} s. The more recently proposed functional NPs (Louizos et al., 2019) characterized a novel family of exchangeable stochastic processes, placing more flexible distributions over latent variables and constructing directed acyclic graphs with la-

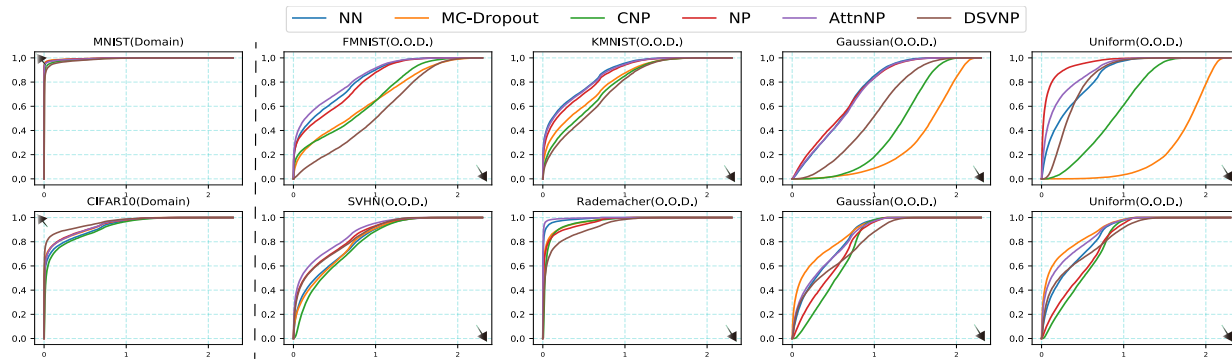


Figure 4. Cumulative Distribution Functions of Entropies in O.O.D. Detection Tasks. Values in X-axis are example entropies ranging from 0 to 2.3, and y-axis records cumulative probabilities. The first row corresponds to the predictive result with models trained on MNIST, while the second is with models trained on CIFAR10. NN means the baseline neural network without dropout layer. Curves in CDFs closer to triangular arrows reveal better uncertainty quantification.

tent affinities of instances in inference and prediction.

Uncertainty Quantification and Computational Complexity. GPs can well characterize aleatoric uncertainty and epistemic uncertainty through kernel function and Gaussian noise. But SP s with non-Gaussian marginals are crucial in modelling. Apart from GPs, there exist some other techniques such as Dropout (Gal & Ghahramani, 2016) or other variants of Bayesian neural networks (Louizos et al., 2017) to quantify uncertainty. In (Depeweg et al., 2018), uncertainties were further decomposed in Bayesian neural network. DSVNP can theoretically capture both uncertainties as an approximate prediction model for general SP s and approaches the problem in a Bayesian way. For the computational cost in prediction, the superior sparse GPs with K-rank covariance matrix approximations (Burt et al., 2019) are with the complexity $\mathcal{O}((M+N)K^2)$, while the variants of CNPs or NPs mostly reduce the complexity $\mathcal{O}((N+M)^3)$ in GPs to $\mathcal{O}(M+N)$ in prediction process. And those for AttnNP and DSVNP are $\mathcal{O}((M+N)N)$.

6. Discussion and Conclusion

In this paper, we present a novel exchangeable stochastic process as DSVNP, which is formalized as a latent variable model. DSVNP integrates latent variables hierarchically and improves the expressiveness of the vanilla NP model. Experiments on high-dimensional tasks demonstrate better capability in prediction and uncertainty quantification. Since this work mainly concentrates on latent variables and associated inference methods, future directions can be the enhancement in the representation of latent variables, such as the use of more flexible equivariant transformations over the context or the dedicated selection of proper context points.

Acknowledgements

We express thanks to Christos Louizos and Bastiaan Veeling for helpful discussions. We would like to thank Daniel Worral, Patrick Forré, Xiahan Shi, Sindy Löwe and anonymous reviewers for precious feedback on the initial manuscript. Q. Wang gratefully acknowledges consistent and sincere support from Prof. Max Welling and was also supported by China Scholarship Council during his pursuing Ph.D at AMLAB. Also, we gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan V GPU.

References

- Bhattacharya, R. N. and Waymire, E. C. *Stochastic processes with applications*. SIAM, 2009.
- Bonilla, E. V., Chai, K. M., and Williams, C. Multi-task gaussian process prediction. In *Advances in Neural Information Processing Systems*, pp. 153–160, 2008.
- Burt, D., Rasmussen, C. E., and Van Der Wilk, M. Rates of convergence for sparse variational gaussian process regression. In *International Conference on Machine Learning*, pp. 862–871, 2019.
- Dai, Z., Damianou, A. C., González, J., and Lawrence, N. D. Variational auto-encoded deep gaussian processes. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016*, 2016.
- Deisenroth, M. and Rasmussen, C. E. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 465–472, 2011.

- Depeweg, S., Hernandez-Lobato, J.-M., Doshi-Velez, F., and Udluft, S. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In *International Conference on Machine Learning*, pp. 1184–1193, 2018.
- Dezfouli, A. and Bonilla, E. V. Scalable inference for gaussian process models with black-box likelihoods. In *Advances in Neural Information Processing Systems*, pp. 1414–1422, 2015.
- Džeroski, S., Demšar, D., and Grbović, J. Predicting chemical parameters of river water quality from bioindicator data. *Applied Intelligence*, 13(1):7–17, 2000.
- Eslami, S. A., Rezende, D. J., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., Ruderman, A., Rusu, A. A., Danihelka, I., Gregor, K., et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pp. 1050–1059, 2016.
- Gal, Y., McAllister, R., and Rasmussen, C. E. Improving pilco with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, volume 4, pp. 34, 2016.
- Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. A. Conditional neural processes. In *International Conference on Machine Learning*, pp. 1704–1713, 2018a.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018b.
- Gordon, J., Bruinsma, W. P., Foong, A. Y. K., Requeima, J., Dubois, Y., and Turner, R. E. Convolutional conditional neural processes. In *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- Hensman, J., Matthews, A., and Ghahramani, Z. Scalable variational gaussian process classification. In *Artificial Intelligence and Statistics*, pp. 351–360, 2015.
- Hernández-Lobato, J. M. and Adams, R. Probabilistic back-propagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869, 2015.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2(5):6, 2017.
- Hofer, E., Kloos, M., Krzykacz-Hausmann, B., Peschke, J., and Woltereck, M. An approximate epistemic uncertainty analysis approach in the presence of epistemic and aleatory uncertainties. *Reliability Engineering & System Safety*, 77(3):229–238, 2002.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, S. M. A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Kingma, D. P. and Welling, M. Stochastic gradient vb and the variational auto-encoder. In *Second International Conference on Learning Representations, ICLR*, volume 19, 2014.
- Kohl, S., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J. R., Maier-Hein, K., Eslami, S. A., Rezende, D. J., and Ronneberger, O. A probabilistic u-net for segmentation of ambiguous images. In *Advances in Neural Information Processing Systems*, pp. 6965–6975, 2018.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pp. 6402–6413, 2017.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Louizos, C., Ullrich, K., and Welling, M. Bayesian compression for deep learning. In *Advances in Neural Information Processing Systems*, pp. 3288–3298, 2017.
- Louizos, C., Shi, X., Schutte, K., and Welling, M. The functional neural process. In *Advances in Neural Information Processing Systems*, pp. 8743–8754, 2019.
- Ma, C., Li, Y., and Hernandez-Lobato, J. M. Variational implicit processes. In *International Conference on Machine Learning*, pp. 4222–4233, 2019.
- Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. Auxiliary deep generative models. In *International Conference on Machine Learning*, pp. 1445–1453, 2016.
- Moreno-Muñoz, P., Artés, A., and Álvarez, M. Heterogeneous multi-output gaussian process prediction. In *Advances in Neural Information Processing Systems*, pp. 6711–6720, 2018.
- Nalisnick, E. T., Matsukawa, A., Teh, Y. W., Görür, D., and Lakshminarayanan, B. Do deep generative models know what they don’t know? In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Oksendal, B. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep exploration via bootstrapped dqn. In *Advances in Neural Information Processing Systems*, pp. 4026–4034, 2016.
- Rasmussen, C. E. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pp. 63–71. Springer, 2003.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286, 2014.
- Rudner, T. G., Fortuin, V., Teh, Y. W., and Gal, Y. On the connection between neural processes and gaussian processes with deep kernels. In *Workshop on Bayesian Deep Learning, NeurIPS*, 2018.
- Salimbeni, H. and Deisenroth, M. Doubly stochastic variational inference for deep gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 4588–4599, 2017.
- Salimbeni, H., Dutordoir, V., Hensman, J., and Deisenroth, M. Deep gaussian processes with importance-weighted variational inference. In *International Conference on Machine Learning*, pp. 5589–5598, 2019.
- Singh, G., Yoon, J., Son, Y., and Ahn, S. Sequential neural processes. In *Advances in Neural Information Processing Systems*, pp. 10254–10264, 2019.
- Snelson, E. and Ghahramani, Z. Sparse gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pp. 1257–1264, 2006.
- Snelson, E., Ghahramani, Z., and Rasmussen, C. E. Warped gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 337–344, 2004.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, pp. 3483–3491, 2015.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. How to train deep variational autoencoders and probabilistic ladder networks. In *33rd International Conference on Machine Learning (ICML 2016)*, 2016.
- Spyromitros-Xioufis, E., Tsoumakas, G., Groves, W., and Vlahavas, I. Multi-target regression via input space expansion: treating targets as inputs. *Machine Learning*, 104(1):55–98, 2016.
- Sun, S., Zhang, G., Shi, J., and Grosse, R. B. Functional variational bayesian neural networks. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Thakur, S., van Hoof, H., Higuera, J. C. G., Precup, D., and Meger, D. Uncertainty aware learning from demonstrations in multiple contexts using bayesian neural networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 768–774. IEEE, 2019.
- Titsias, M. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Titsias, M. and Lawrence, N. D. Bayesian gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 844–851, 2010.
- Titsias, M. and Lázaro-Gredilla, M. Doubly stochastic variational bayes for non-conjugate inference. In *International Conference on Machine Learning*, pp. 1971–1979, 2014.
- Tran, D., Ranganath, R., and Blei, D. M. Variational gaussian process. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016*, 2016.
- Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pp. 5523–5533, 2017.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. In *Artificial Intelligence and Statistics*, pp. 370–378, 2016a.
- Wilson, A. G., Hu, Z., Salakhutdinov, R. R., and Xing, E. P. Stochastic variational deep kernel learning. In *Advances in Neural Information Processing Systems*, pp. 2586–2594, 2016b.
- Zhao, J. and Sun, S. Variational dependent multi-output gaussian process dynamical systems. *The Journal of Machine Learning Research*, 17(1):4134–4169, 2016.