

# Deeper Connections between Neural Networks and Gaussian Processes Speed-up Active Learning

Evgenii Tsymbalov\*, Sergei Makarychev, Alexander Shapeev and Maxim Panov

Skolkovo Institute of Science and Technology (Skoltech)

{e.tsymbalov, sergei.makarychev, a.shapeev, m.panov}@skoltech.ru

## Abstract

Active learning methods for neural networks are usually based on greedy criteria, which ultimately give a single new design point for the evaluation. Such an approach requires either some heuristics to sample a batch of design points at one active learning iteration, or retraining the neural network after adding each data point, which is computationally inefficient. Moreover, uncertainty estimates for neural networks sometimes are overconfident for the points lying far from the training sample. In this work, we propose to approximate Bayesian neural networks (BNN) by Gaussian processes (GP), which allows us to update the uncertainty estimates of predictions efficiently without retraining the neural network while avoiding overconfident uncertainty prediction for out-of-sample points. In a series of experiments on real-world data, including large-scale problems of chemical and physical modeling, we show the superiority of the proposed approach over the state-of-the-art methods.

## 1 Introduction

In the modern applications of machine learning, especially in engineering design, physics, chemistry, molecular and materials modeling, the datasets are usually of limited size due to the expensive cost of computations. On the other side, such applications typically allow the calculation of additional data at points chosen by the experimentalist. Thus, there is the need to get the highest possible approximation quality with as few function evaluations as possible. *Active learning* methods help to achieve this goal by trying to select the best candidates for further target function computation using the already existing data and machine learning models based on these data.

Modern active learning methods are usually based either on ensemble-based methods [Settles, 2012] or on probabilistic models such as Gaussian process regression [Sacks et al., 1989; Burnaev and Panov, 2015]. However, the models from these classes often do not give state-of-the-art results in the downstream tasks. For example, Gaussian process-based models have very high computational complexity, which in

many cases prohibits their usage even for moderate-sized datasets. Moreover, GP training is usually done for stationary covariance functions, which is a very limiting assumption in many real-world scenarios. Thus, active learning methods applicable to the broader classes of models are needed.

In this work, we aim to develop efficient active learning strategies for neural network models. Unlike Gaussian processes, neural networks are known for relatively good scalability with the dataset size and are very flexible in terms of dependencies they can model. However, active learning for neural networks is currently not very well developed, especially for tasks other than classification. Recently some uncertainty estimation and active learning approaches were introduced for Bayesian neural networks [Gal et al., 2017; Hafner et al., 2018]. However, still one of the main complications is in designing active learning approaches which allow for sampling points without retraining the neural network too often. Another problem is that neural networks due to their parametric structure sometimes give overconfident predictions in the areas of design space lying far from the training sample. Recently, [Matthews et al., 2018] proved that deep neural networks with random weights converge to Gaussian processes in the infinite layer width limit. Our approach aims to approximate trained Bayesian neural networks by GPs and show that obtained uncertainty estimates enable to improve active learning performance significantly.

We summarize the main contributions of the paper as follows:

1. We propose to compute the approximation of trained Bayesian neural network with Gaussian process, which allows using the GPs machinery to calculate uncertainty estimates for neural networks. We propose active learning strategy based on obtained uncertainty estimates which significantly speed-up the active learning with NNs by improving the quality of selected samples and decreasing the number of times the NN is retrained.
2. The proposed active learning framework shows significant improvement over competing approaches in the wide range of real-world problems, including the cutting edge applications in chemoinformatics and physics of oil recovery.

In the next section, we discuss the problem statement and the proposed approach in detail.

\*Contact Author

## 2 Active Learning with Bayesian Neural Networks and Beyond

### 2.1 Problem Statement

We consider a regression problem with an unknown function  $f(\mathbf{x})$  defined on a subset of Euclidean space  $\mathcal{X} \subset \mathbb{R}^d$ , where an approximation function  $\hat{f}(\mathbf{x})$  should be constructed based on noisy observations

$$y = f(\mathbf{x}) + \epsilon$$

with  $\epsilon$  being some random noise.

We focus on *active learning* scenario which allows to iteratively enrich training set by computing the target function in design points specified by the experimenter. More specifically, we assume that the initial training set  $D_{\text{init}} = \{\mathbf{x}_i, y_i = f(\mathbf{x}_i) + \epsilon_i\}_{i=1}^N$  with precomputed function values is given. On top of it, we are given another set of points called “pool”  $\mathcal{P} = \{\mathbf{x}_j \in \mathcal{X}\}_{j=1}^{N_{\text{pool}}}$ , which represents unlabelled data. Each point  $\mathbf{x} \in \mathcal{P}$  can be annotated by computing  $y = f(\mathbf{x}) + \epsilon$  so that the pair  $\{\mathbf{x}, y\}$  is added to the training set.

The standard active learning approaches rely on the greedy point selection

$$\mathbf{x}_{\text{new}} = \arg \max_{\mathbf{x} \in \mathcal{P}} A(\mathbf{x} | \hat{f}, D), \quad (1)$$

where  $A(\mathbf{x} | \hat{f}, D)$  is so-called *acquisition function*, which is usually constructed based on the current training set  $D$  and corresponding approximation function  $\hat{f}$ .

The most popular choice of acquisition function is the variance  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, D)$  of the prediction  $\hat{f}(\mathbf{x})$ , which can be easily estimated in some machine learning models such as Random forest [Mentch and Hooker, 2016] or Gaussian process regression [Rasmussen, 2004]. However, for many types of models (e.g., neural networks) the computation of prediction uncertainty becomes a nontrivial problem.

### 2.2 Uncertainty Estimation and Active Learning with Bayesian Neural Networks

In this work, we will stick to the Bayesian approach which treats neural networks as probabilistic models  $p(y | \mathbf{x}, \mathbf{w})$ . Vector of neural network weights  $\mathbf{w}$  is assumed to be a random variable with some prior distribution  $p(\mathbf{w})$ . The likelihood  $p(y | \mathbf{x}, \mathbf{w})$  determines the distribution of network output at a point  $\mathbf{x}$  given specific values of parameters  $\mathbf{w}$ . There is vast literature on training Bayesian networks (see [Graves, 2011] and [Paisley et al., 2012] among many others), which mostly targets the so-called variational approximation of the intractable posterior distribution  $p(\mathbf{w} | D)$  by some easily computable distribution  $q(\mathbf{w})$ .

The approximate posterior predictive distribution reads as:

$$q(y | \mathbf{x}) = \int p(y | \mathbf{x}, \mathbf{w}) q(\mathbf{w}) d\mathbf{w}.$$

The simple way to generate random values from this distribution is to use the Monte-Carlo approach, which allows estimating the mean:

$$\mathbb{E}_{q(y | \mathbf{x})} y \approx \frac{1}{T} \sum_{t=1}^T \hat{f}(\mathbf{x}, \mathbf{w}_t),$$




Figure 1: Contour plot of neural network variance prediction for bivariate problem and pool points (white). Five points from the pool with maximum values of variance lie in the same region of design space (upper-left corner).

where the weight values  $\mathbf{w}_t$  are i.i.d. random variables from distribution  $q(\mathbf{w})$ . Similarly, one can use Monte-Carlo to estimate the approximate posterior variance  $\hat{\sigma}^2(\mathbf{x} | \hat{f})$  of the prediction  $y$  at a point  $\mathbf{x}$  and use it as an acquisition function:

$$A(\mathbf{x} | \hat{f}, D) = \hat{\sigma}^2(\mathbf{x} | \hat{f}).$$

We note that the considered acquisition function formally doesn't depend on the dataset  $D$  except for the fact that  $D$  was used for training the neural network  $\hat{f}$ .

Let us note that the general greedy active learning approach (1) by design gives one candidate point per active learning iteration. If one tries to use (1) to obtain several samples with the same acquisition function, it usually results in obtaining several nearby points from the same region of design space, see Figure 1. Such behaviour is typically undesirable as nearby points are likely to have very similar information about the target function. Moreover, neural network uncertainty predictions are sometimes overconfident in out-of-sample regions of design space.

There are several approaches to overcome these issues each having its drawbacks:

1. One may retrain the model after each point addition which may result in a significant change of the acquisition function and lead to the selection of a more diverse set of points. However, such an approach is usually very computationally expensive, especially for neural network-based models.
2. One may try to add distance-based heuristic, which explicitly prohibits sampling points which are very close to each other and increase values of acquisition function for points positioned far from the training sample. Such an approach may give satisfactory results in some cases, however usually requires fine-tuning towards particular application (like the selection of specific distance function or choice the parameter value which determines whether two points are near or not), while its performance may degrade in high dimensional problems.
3. One may treat specially normalized vector of acquisition function values at points from the pool as a probabilistic distribution and sample the desired number of points




Figure 2: Bivariate distribution plots for the stochastic NN output at points  $\mathbf{x}_1$ ,  $\mathbf{x}_2$  and  $\mathbf{x}_3$ , where  $\mathbf{x}_1$  is much closer to  $\mathbf{x}_2$  in feature space than to  $\mathbf{x}_3$ . Both univariate and bivariate distributions are Gaussian-like, while the correlation between function values is much higher for closer points.

based on their probabilities (the higher acquisition function value, the point is more likely to be selected). This approach usually improves over greedy baseline procedure. However, it still gives many nearby points.

In the next section, we propose the approach to deal with problems above by considering the Gaussian process approximation of the neural network.

### 2.3 Gaussian Process Approximation of Bayesian Neural Network

Effectively, the random function

$$\hat{f}(\mathbf{x}, \mathbf{w}) = \mathbb{E}_{p(y|\mathbf{x}, \mathbf{w})} y$$

is the stochastic process indexed by  $\mathbf{x}$ . The covariance function of the process  $\hat{f}(\mathbf{x}, \mathbf{w})$  is given by

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{q(\mathbf{w})} (\hat{f}(\mathbf{x}, \mathbf{w}) - m(\mathbf{x})) (\hat{f}(\mathbf{x}', \mathbf{w}) - m(\mathbf{x}')),$$

where  $m(\mathbf{x}) = \mathbb{E}_{q(\mathbf{w})} \hat{f}(\mathbf{x}, \mathbf{w})$ .

As was shown in [Matthews et al., 2018; Lee et al., 2017] neural networks with random weights converge to Gaussian processes in the infinite layer width limit. However, one is not limited to asymptotic properties of purely random networks as Bayesian neural networks trained on real-world data exhibit near Gaussian behaviour, see the example on Figure 2.

We aim to make the Gaussian process approximation  $\hat{g}(\mathbf{x} | \hat{f})$  of the stochastic process  $\hat{f}(\mathbf{x}, \mathbf{w})$  and compute its posterior variance  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X)$  given the set of anchor points  $X = \{\mathbf{x}_i\}_{i=1}^N$ . Typically,  $X$  is a subset of the training sample. Given  $X$ , Monte-Carlo estimates  $\hat{k}(\mathbf{x}', \mathbf{x}'')$  of the covariance function  $k(\mathbf{x}', \mathbf{x}'')$  for every pair of points  $\mathbf{x}', \mathbf{x}'' \in X \cup \mathbf{x}$  allow computing

$$\hat{\sigma}^2(\mathbf{x} | \hat{f}, X) = \hat{k}(\mathbf{x}, \mathbf{x}) - \hat{\mathbf{k}}^T(\mathbf{x}) \hat{K}^{-1} \hat{\mathbf{k}}(\mathbf{x}), \quad (2)$$

where  $\hat{K} = [\hat{k}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^N$  and  $\hat{\mathbf{k}}(\mathbf{x}) = (\hat{k}(\mathbf{x}_1, \mathbf{x}), \dots, \hat{k}(\mathbf{x}_N, \mathbf{x}))^T$ .

We note that only the trained neural network  $\hat{f}(\mathbf{x}, \mathbf{w})$  and the ability to sample from the distribution  $q(\mathbf{w})$  are needed to compute  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X)$ .

### 2.4 Active Learning Strategies

The benefits of the Gaussian process approximation and the usage of the formula (2) are not evident as one might directly estimate the variance of neural network prediction  $\hat{f}(\mathbf{x}, \mathbf{w})$  at any point  $\mathbf{x}$  by sampling from  $q(\mathbf{w})$  and use it as acquisition function. However, the approximate posterior variance  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X)$  of Gaussian process  $\hat{g}(\mathbf{x} | \hat{f})$  has an important property that is has large values for points  $\mathbf{x}$  lying far from the points from the training set  $X$ . Thus, out-of-sample points are likely to be selected by the active learning procedure.

Moreover, the function  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X)$  depends solely on covariance function values for points from a set  $X$  (and not on the output function values). Such property allows updating uncertainty predictions by just adding sample points to the set  $X$ . More specifically, if we decide to sample some point  $\mathbf{x}'$ , then the updated posterior variance  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X')$  for  $X' = X \cup \mathbf{x}'$  can be easily computed:

$$[c] \hat{\sigma}^2(\mathbf{x} | \hat{f}, X') = \hat{\sigma}^2(\mathbf{x} | \hat{f}, X) - \frac{\hat{k}^2(\mathbf{x}, \mathbf{x}' | \hat{f}, X)}{\hat{\sigma}^2(\mathbf{x}' | \hat{f}, X)},$$

where  $\hat{k}(\mathbf{x}, \mathbf{x}' | \hat{f}, X) = \hat{k}(\mathbf{x}, \mathbf{x}') - \hat{\mathbf{k}}^T(\mathbf{x}) \hat{K}^{-1} \hat{\mathbf{k}}(\mathbf{x}')$  is the posterior covariance function of the process  $\hat{g}(\mathbf{x} | \hat{f})$  given  $X$ .

Importantly,  $\hat{\sigma}^2(\mathbf{x} | \hat{f}, X')$  for points  $\mathbf{x}$  in some vicinity of  $\mathbf{x}'$  will have low values, which guarantees that further sampled points will not lie too close to  $\mathbf{x}'$  and other points from the training set  $X$ . The resulting NNGP active learning procedure is depicted in Figure 3.

## 3 Experiments

Our experimental study is focused on real-world data to ensure that the algorithms are indeed useful in real-world scenarios. We compare the following algorithms:

- pure random sampling;
- sampling based on the variance of NN stochastic output from  $\hat{f}$ , which we refer to as MCDUE (see [Gal et al., 2017; Tsymbalov et al., 2018]);
- the proposed GP-based approaches (NNGP).

Following [Gal and Ghahramani, 2016] we use the Bayesian NN with the Bernoulli distribution on the weights, which is equivalent to using the dropout on the inference stage.

### 3.1 Airline Delays Dataset and NCP Comparison

We start the experiments from comparing the proposed approach to active learning with the one based on uncertainty estimates obtained from a Bayesian neural network with *Noise Contrastive Prior (NCP)*, see [Hafner et al., 2018]. Following this paper, we use the airline delays dataset (see [Hensman et al., 2013]) and NN consisting of two layers with 50 neurons each, leaky ReLU activation function, and trained with respect to NCP-based loss function.

We took a random subset of 50,000 data samples from the data available on the January – April of 2008 as a training set and we chose 100,000 random data samples from May

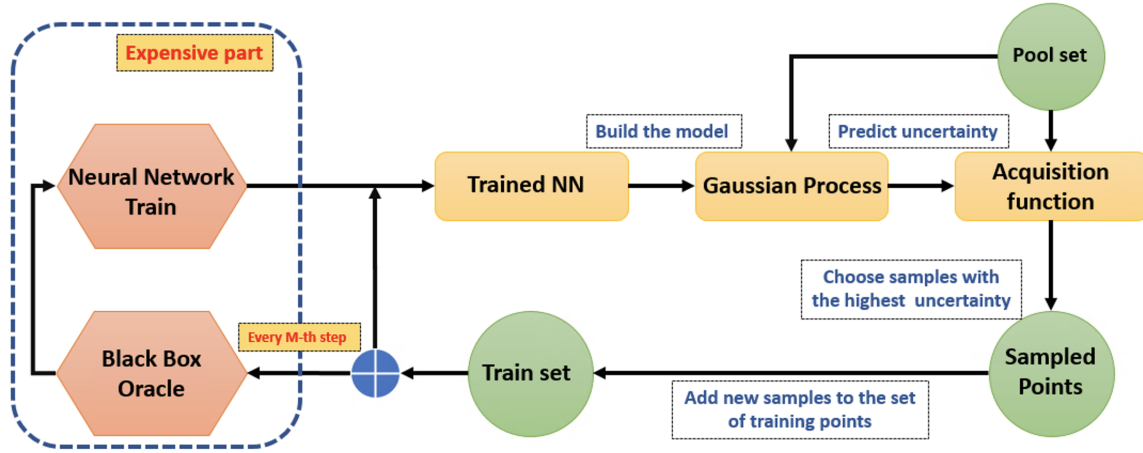


Figure 3: Schematic representation of the NNGP approach to active learning. GP is fitted on the data from a stochastic output of NN, and the posterior variance of GP is used as an acquisition function for sampling. The most computationally expensive part (function evaluation at sampled points and neural network retraining) is done only every  $M$  steps of sampling, while all the intermediate iterations are based solely on trained neural network and corresponding GP approximation.

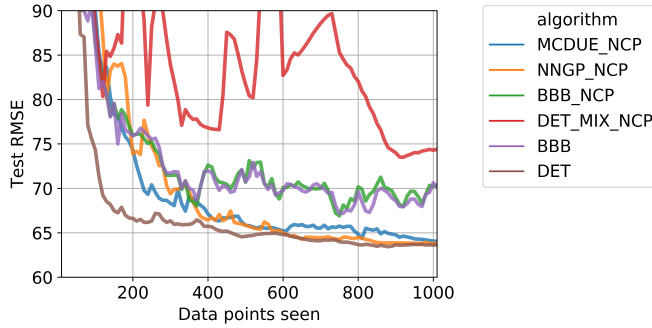


Figure 4: Root mean squared errors as functions of active learning iteration for different methods on the Airline delays data set. Plots show median of the errors over 25 runs. NNGP initially has a much higher error, but shows the rapid improvement and becomes the best NCP-based method near iteration 400.

of 2008 as a test set. We used the following variables as input features *PlaneAge*, *Distance*, *CRSDepTime*, *AirTime*, *CRSArrTime*, *DayOfWeek*, *DayofMonth*, *Month* and *ArrDelay + DepDelay* as a target.

The results for the test set are shown in Figure 4. The proposed NNGP approach demonstrates a comparable error with respect to the previous results and outperforms (on average) other methods in the continuous active learning scenario.

### 3.2 Experiments on UCI Datasets

We conducted a series of experiments with active learning performed on the data from the UCI ML repository<sup>1</sup>. All the datasets represent real-world regression problems with 15+ dimensions and 30000+ samples, see Table 1 for details.

<sup>1</sup> D. Dua, E. Taniskidou. UCI Machine Learning Repository [http://archive.ics.uci.edu/ml],

| Dataset name                                       | Samples | Attributes | Target                    |
|--|---------|------------|---------------------------|
| BlogFeedback [Buza, 2014]                          | 60021   | 281        | Number of comments        |
| SGEMM GPU [Nugteren and Codreanu, 2015]            | 241600  | 18         | Calculation time (median) |
| YearPredictionMSD [Bertin-Mahieux et al., 2011]    | 515345  | 90         | Year                      |
| Relative location of CT slices [Graf et al., 2011] | 53500   | 386        | Relative location         |
| Online News Popularity [Fernandes et al., 2015]    | 39797   | 61         | Number of shares          |
| KEGG Network [Shannon et al., 2003]                | 53414   | 24         | Clustering coefficient    |

Table 1: Summary of the datasets used in experiments with UCI data.

For every experiment, data are shuffled and split in the following proportions: 10% for the training set  $D_{train}$ , 5% for the test set  $D_{test}$ , 5% for the validation set  $D_{val}$  needed for early-stopping and 80% for the pool  $\mathcal{P}$ .

We used a simple neural network with three hidden layers of sizes 256, 128 and 128. We performed 16 active learning iterations with 200 points picked at each iteration.

To compare the performance of the algorithms across the different datasets and different choices of training samples, we will use so-called *Dolan-More curves*. Let  $q_a^p$  be an error measure of the  $a$ -th algorithm on the  $\mathcal{P}$ -th problem. Then, determining the performance ratio  $r_a^p = \frac{q_a^p}{\min_x(q_x^p)}$ , we can define the Dolan-More curve as a function of the performance ratio factor  $\tau$ :

$$\rho_a(\tau) = \frac{\#\{p: r_a^p \leq \tau\}}{n_p},$$

where  $n_p$  is the total number of evaluations for the problem  $p$ . Thus,  $\rho_a(\tau)$  defines the fraction of problems in which the  $a$ -th algorithm has the error not more than  $\tau$  times bigger than the best competitor in the chosen performance metric. Note that  $\rho_a(1)$  is the ratio of problems on which the  $a$ -th algorithm

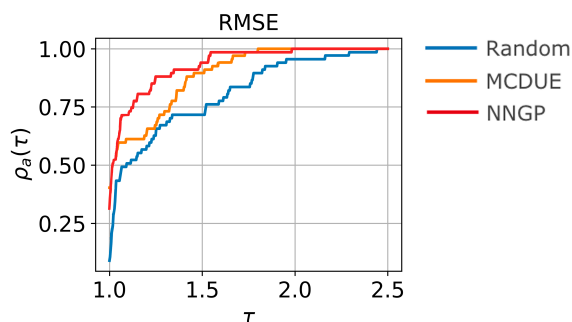


Figure 5: Dolan-More curves for UCI datasets and different active learning algorithms after 16 active learning iterations. Root mean squared error (RMSE) on independent test set is considered. NNGP-based algorithm show better performance compared to MCDUE and random sampling.

performance was the best, while in general, the higher curve means the better performance of the algorithm.

The Dolan-More curves for the errors of approximation for considered problems after the 16th iteration of the active learning procedure are presented in Figure 5. We see that the NNGP procedure is superior in terms of RMSE compared to MCDUE and random sampling.

### 3.3 SchNet Training

To demonstrate the power of our approach, we conducted a series of numerical experiments with the state-of-the-art neural network architecture in the field of chemoinformatics ‘‘SchNet’’ [Schütt et al., 2017]. This network takes information about an organic molecule as an input, and, after special preprocessing and complicated training procedure, outputs some properties of the molecule (like energy). Despite its complex structure, SchNet contains fully connected layers, so it is possible to use a dropout in between them.

We tested our approach on the problem of predicting the internal energy of the molecule at 0K from the QM9 data set [Ramakrishnan et al., 2014]. We used a Tensorflow implementation of a SchNet with the same architecture as in original paper except for an increased size of hidden layers (from 64 and 32 units to 256 and 128 units, respectively) and dropout layer placed in between of them and turned on during an inference only.

In our experiment, we separate the whole dataset of 133 885 molecules into the initial set of 10 000 molecules, the test set of 5 000 molecules, and the rest of the data allocated as the pool. On each active learning iteration, we perform 100 000 training epochs and then calculate the uncertainty estimates using either MCDUE or NNGP approach. We then select 2 500 molecules with the highest uncertainty from the pool, add them to the training set and perform another active learning iteration.

The results are shown in Figure 6. The NNGP approach demonstrates the most steady decrease in error, with the 25% accuracy increase in RMSE. Such improvement is very significant in terms of the time savings for the computationally expensive quantum-mechanical calculations. For example, to reach the RMSE of 2 kcal/mol starting from the SchNet

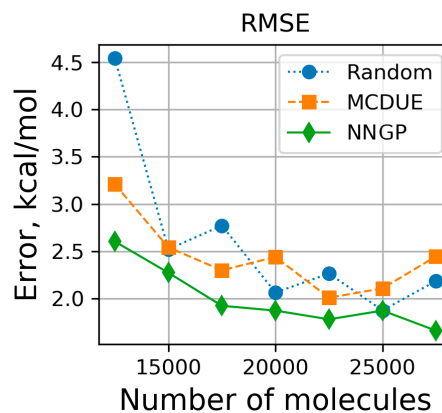


Figure 6: Training curves for the active learning scenario for SchNet: starting from 10 000 random molecules pick 2 500 based on the uncertainty estimate. NNGP-based algorithm results in 25% decrease in RMSE. Simple dropout-based approach (MCDUE) does not demonstrate a difference from the random sampling in terms of accuracy.

trained on 10 000 molecules, one need to additionally sample 15 000 molecules in case of random sampling or just 7 500 molecules using the NNGP uncertainty estimation procedure.

### 3.4 Hydraulic Simulator

In the oil industry, to determine the optimal parameters and control the drilling process, engineers carry out hydraulic calculations of the well’s circulation system, that usually are based on either empirical formulas or semi-analytical solutions of the hydrodynamic equations. However, such semi-analytical solutions are known just for few individual cases, while in the other ones, only very crude approximations are usually available (see [Podryabinkin et al., 2013] for details). As a result, such calculations have relatively low precision. On the other hand, the full-scale numerical solution of the hydrodynamic equations describing the flow of drilling fluids can provide a sufficient level of accuracy, but it requires significant computational resources and subsequently is very costly and time-consuming. The possible solution to this problem is the use of a surrogate model.

We used a surrogate model for the fluid flow in the well-bore while drilling. The oracle, in this case, is a numerical solver for the hydrodynamic equations [Podryabinkin et al., 2013], which, given six main parameters of the drilling process as an input, outputs the unit-less hydraulic resistance coefficient that characterizes the drop in the pressure.

In this experiment, we used a two-layer neural network with 50 neurons per each layer and LeakyReLU activation function. Initial training and pool sets had 50 and 20000 points respectively. We completed 10 active learning iterations, adding 50 points per each iteration. The results are shown in Figure 7. Clearly, NNGP is superior in terms of RMSE and MAE, while maximum error for NNGP is 10 times lower.




Figure 7: Training curves for the active learning scenario for the hydraulic simulator case. The NNGP sampling algorithm outperforms the random sampling and MCDUE with a large margin in terms of maximal error.

### 4 Related Work

Active learning [Settles, 2012] (also known as *adaptive design of experiments* [Forrester et al., 2008] in statistics and engineering design) is a framework which allows the additional data points to be annotated by computing target function value and then added to the training set. The particular points to sample are usually chosen as the ones that maximize so-called *acquisition function*. The most popular approaches to construct acquisition function are committee-based methods [Seung et al., 1992], also known as *query-by-committee*, where an ensemble of models is trained on the same or various parts of data, and Bayesian models, such as Gaussian Processes [Rasmussen, 2004], where the uncertainty estimates can be directly inferred from the model [Sacks et al., 1989; Burnaev and Panov, 2015]. However, the majority of existing approaches are computationally expensive and sometimes even intractable for large sample sizes and input dimensions.

Active learning for neural networks is not very well developed as neural networks usually excel in applications with large datasets already available. Ensembles of neural networks (see [Li et al., 2018] for a detailed review) often boil down to an independent training of several models, which works well in some applications [Beluch et al., 2018], but is computationally expensive for the large-scale applications. Bayesian neural networks provide uncertainty estimates, which can be efficiently used for active learning procedures [Gal et al., 2017; Hafner et al., 2018]. The most popular uncertainty estimation approach for BNNs, *MC-Dropout*, is based on classical dropout first proposed as a technique for a neural network regularization, which was recently interpreted as a method for approximate inference in Bayesian neural networks and shown to provide unbiased Monte-Carlo estimates predictive variance [Gal and Ghahramani, 2016]. However, it was shown [Beluch et al., 2018] that uncertainty estimates based on MC-Dropout are usually less efficient for active learning than those based on ensembles. [Pop and Fulp, 2018] suggests that it is partially due to the mode collapse effect, which leads to overconfident out of sample predictions of uncertainty.

Connections between neural networks and Gaussian processes recently gain significant attention, see [Matthews et al., 2018; Lee et al., 2017] which study random, untrained NNs and show that such networks can be approximated by Gaus-

sian processes in infinite network width limit. Another direction is incorporation of the GP-like elements into the NN structure, see [Sun et al., 2018; Garnelo et al., 2018] for some recent contributions. In [Bradshaw et al., 2017], authors propose an NN+GP pipeline that may be trained end-to-end; however, the GP part slows down both training and inference whereas we use the GP for inference only.

### 5 Summary and Discussion

We have proposed a novel dropout-based method for the uncertainty estimation for deep neural networks, which uses the approximation of neural network by Gaussian process. Experiments on different architectures and real-world problems show that the proposed estimate allows to achieve state-of-the-art results in context of active learning. Importantly, the proposed approach works for any neural network architecture involving dropout (as well as other Bayesian neural networks), so it can be applied to very wide range of networks and problems without a need to change neural network architecture.

It is of interest whether the proposed approach can be efficient for other applications areas, such as image classification. We also plan to study the applicability of modern methods for GP speed-up in order to improve the scalability of proposed approach.

### Acknowledgements

E.T. and A.S. were supported by the Skoltech NGP Program No. 2016-7/NGP (a Skoltech-MIT joint project).

### References

[Beluch et al., 2018] Beluch, W. H., Genewein, T., Nürnberger, A., and Köhler, J. M. (2018). The power of ensembles for active learning in image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9368–9377.

[Bertin-Mahieux et al., 2011] Bertin-Mahieux, T., Ellis, D. P., Whitman, B., and Lamere, P. (2011). The million song dataset. In *Ismir*, volume 2, page 10.

[Bradshaw et al., 2017] Bradshaw, J., Matthews, A. G. d. G., and Ghahramani, Z. (2017). Adversarial examples, uncer-

- tainty, and transfer testing robustness in gaussian process hybrid deep networks. *arXiv:1707.02476*.
- [Burnaev and Panov, 2015] Burnaev, E. and Panov, M. (2015). Adaptive design of experiments based on gaussian processes. In *International Symposium on Statistical Learning and Data Sciences*, pages 116–125. Springer.
- [Buza, 2014] Buza, K. (2014). Feedback prediction for blogs. In *Data analysis, machine learning and knowledge discovery*, pages 145–152. Springer.
- [Fernandes et al., 2015] Fernandes, K., Vinagre, P., and Cortez, P. (2015). A proactive intelligent decision support system for predicting the popularity of online news. In *Portuguese Conference on Artificial Intelligence*, pages 535–546. Springer.
- [Forrester et al., 2008] Forrester, A., Keane, A., et al. (2008). *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.
- [Gal and Ghahramani, 2016] Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proc. ICML'16*, pages 1050–1059.
- [Gal et al., 2017] Gal, Y., Islam, R., and Ghahramani, Z. (2017). Deep bayesian active learning with image data. *arXiv:1703.02910*.
- [Garnelo et al., 2018] Garnelo, M., Rosenbaum, D., Maddison, C. J., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D. J., and Eslami, S. (2018). Conditional neural processes. *arXiv:1807.01613*.
- [Graf et al., 2011] Graf, F., Kriegel, H.-P., Schubert, M., Pölsterl, S., and Cavallaro, A. (2011). 2d image registration in ct images using radial image descriptors. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 607–614. Springer.
- [Graves, 2011] Graves, A. (2011). Practical variational inference for neural networks. In *Advances in neural information processing systems*, pages 2348–2356.
- [Hafner et al., 2018] Hafner, D., Tran, D., Irpan, A., Lillcrap, T., and Davidson, J. (2018). Reliable uncertainty estimates in deep neural networks using noise contrastive priors. *arXiv:1807.09289*.
- [Hensman et al., 2013] Hensman, J., Fusi, N., and Lawrence, N. D. (2013). Gaussian processes for big data. *arXiv:1309.6835*.
- [Lee et al., 2017] Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. (2017). Deep neural networks as gaussian processes. *arXiv:1711.00165*.
- [Li et al., 2018] Li, H., Wang, X., and Ding, S. (2018). Research and development of neural network ensembles: a survey. *Artificial Intelligence Review*, 49(4):455–479.
- [Matthews et al., 2018] Matthews, A. G. d. G., Rowland, M., Hron, J., Turner, R. E., and Ghahramani, Z. (2018). Gaussian process behaviour in wide deep neural networks. *arXiv:1804.11271*.
- [Mentch and Hooker, 2016] Mentch, L. and Hooker, G. (2016). Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *The Journal of Machine Learning Research*, 17(1):841–881.
- [Nugteren and Codreanu, 2015] Nugteren, C. and Codreanu, V. (2015). Cltune: A generic auto-tuner for opencl kernels. In *Embedded Multicore/Many-core Systems-on-Chip (MCSoc), 2015 IEEE 9th International Symposium on*, pages 195–202. IEEE.
- [Paisley et al., 2012] Paisley, J., Blei, D. M., and Jordan, M. I. (2012). Variational bayesian inference with stochastic search. In *Proc. ICML'12*, pages 1363–1370.
- [Podryabinkin et al., 2013] Podryabinkin, E., Rudyak, V., Gavrilov, A., and May, R. (2013). Detailed modeling of drilling fluid flow in a wellbore annulus while drilling. In *Proceedings of the International Conference on Offshore Mechanics and Arctic Engineering - OMAE*, volume 6.
- [Pop and Fulop, 2018] Pop, R. and Fulop, P. (2018). Deep ensemble bayesian active learning: Addressing the mode collapse issue in monte carlo dropout via ensembles. *arXiv preprint arXiv:1811.03897*.
- [Ramakrishnan et al., 2014] Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. (2014). Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1:140022.
- [Rasmussen, 2004] Rasmussen, C. E. (2004). Gaussian processes in machine learning. In *Advanced lectures on machine learning*, pages 63–71. Springer.
- [Sacks et al., 1989] Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical science*, pages 409–423.
- [Schütt et al., 2017] Schütt, K. T., Arbabzadah, F., Chmiela, S., Müller, K. R., and Tkatchenko, A. (2017). Quantum-chemical insights from deep tensor neural networks. *Nature communications*, 8:13890.
- [Settles, 2012] Settles, B. (2012). Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114.
- [Seung et al., 1992] Seung, H. S., Opper, M., and Sompolinsky, H. (1992). Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294. ACM.
- [Shannon et al., 2003] Shannon, P., Markiel, A., Ozier, O., Baliga, N. S., Wang, J. T., Ramage, D., Amin, N., Schwikowski, B., and Ideker, T. (2003). Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research*, 13(11):2498–2504.
- [Sun et al., 2018] Sun, S., Zhang, G., Wang, C., Zeng, W., Li, J., and Grosse, R. (2018). Differentiable compositional kernel learning for gaussian processes. *arXiv:1806.04326*.
- [Tsymbalov et al., 2018] Tsymbalov, E., Panov, M., and Shapeev, A. (2018). Dropout-based active learning for regression. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 247–258. Springer.