



Dropout Strikes Back: Improved Uncertainty Estimation via Diversity Sampling

Kirill Fedyanin¹(✉), Evgenii Tsymbalov^{1,2}(✉), and Maxim Panov¹(✉)

¹ Skolkovo Institute of Science and Technology (Skoltech), Moscow, Russia
{k.fedyanin,m.panov}@skoltech.ru

² Yandex, Moscow, Russia
etsymba@yandex-team.ru

Abstract. Uncertainty estimation for machine learning models is of high importance in many scenarios such as constructing the confidence intervals for model predictions and detection of out-of-distribution or adversarially generated points. In this work, we show that modifying the sampling distributions for dropout layers in neural networks improves the quality of uncertainty estimation. Our main idea consists of two main steps: computing data-driven correlations between neurons and generating samples, which include maximally diverse neurons. In a series of experiments on simulated and real-world data, we demonstrate that the diversification via *determinantal point processes*-based sampling achieves state-of-the-art results in uncertainty estimation for regression and classification tasks. An important feature of our approach is that it does not require any modification to the models or training procedures, allowing straightforward application to any deep learning model with dropout layers.

Keywords: Uncertainty estimation · Neural networks · Dropout · Determinantal point processes

1 Introduction

Uncertainty estimation (UE) recently became a very active area of research in deep learning. Neural networks usually are treated as black boxes, and in general, they are prone to overconfidence [9,15]. Uncertainty estimation methods aim to help overcome this drawback by identifying potentially erroneous predictions. This can be especially important for error-critical applications like medical diagnostics [4] or autonomous car driving [8]. Another important application for uncertainty estimation is active learning [34]. The majority of sampling criteria in active learning are based on estimates of uncertainty, which makes it important to obtain high-quality uncertainty estimates.

There are several main approaches for uncertainty estimation for deep neural networks. Bayesian neural networks (BNN) and variational inference in particular represent a natural way for uncertainty estimation due to availability of

well-defined posteriors, but they can be prohibitively slow for large-scale applications. The usage of dropout at the inference stage was shown to be good and efficient approximation to BNNs [9, 10]. The ensembles of independently trained models [25] have state-of-the-art performance in many tasks requiring uncertainty estimation [37]. Recently, forcing models in ensembles to be more diverse was shown to improve results even further [22]. The drawback of ensembles is that we need to train and use multiple models that require additional resources, i.e. more memory to store models and more computing power for training.

In this work, we aim to develop a new approach for dropout-based uncertainty estimation. Usually there are many highly correlated neurons in neural networks, which results in a slow convergence of estimates based on the standard uniform sampling in dropout layers. We propose to estimate correlations between neurons based on the data and sample the most diverse neurons in order to improve the convergence of the estimates and, as a result, the quality of uncertainty estimates. As a particular realization of the general idea, we suggest sampling dropout masks using the machinery of determinantal point processes (DPP) [24] which are known to give diverse samples.

We summarize the main contributions of the paper as follows:

- We propose two DPP-based sampling methods for neural networks with dropout. Our approach requires to train only a single model and adds only small overhead on the inference stage compared to plain MC dropout.
- We compare different dropout-based approaches for uncertainty estimation in an extensive series of experiments for real-world regression and classification datasets. The results show superior performance of proposed DPP-based approaches.
- Importantly, the proposed methods show high quality of uncertainty estimation even for very small number of stochastic passes through the network, thus opening the possibility to significantly speed up the inference stage.

The rest of the paper is organized as follows. Section 2 introduces the proposed method for DPP-based sampling from neural networks with dropout. In Sect. 3, we show the efficiency of the proposed approach in the problem of uncertainty estimation. Section 4 gives an overview of the related work on uncertainty estimation for neural networks. Section 5 concludes the study and highlights some directions for future work.

2 Methods

2.1 Neural Networks with Dropout as Implicit Ensembles

We start by considering a standard fully connected layer in a neural network

$$S_i^h = \sum_{j=1}^{N_{h-1}} w_{ij}^h O_j^{h-1}, \quad i = 1, \dots, N_h, \quad (1)$$

where $O_i^h = \sigma(S_i^h)$ is an output of the h -th layer of the neural network given by a non-linear transformation $\sigma(\cdot)$ of the corresponding pre-activation S_i^h .

An application of dropout to neurons results in the following formula for the pre-activations:

$$S_i^h = \sum_{j=1}^{N_{h-1}} \frac{1}{1-p} m_j^h w_{ij}^h O_j^{h-1}, \quad i = 1, \dots, N_h, \quad (2)$$

where m_j^h are Bernoulli random variables with a probability of 0 equal to p . The outputs O_i^h of the h -th layer remain to be computed by $O_i^h = \sigma(S_i^h)$. Note that if an input variable of neural network is denoted by \mathbf{x} , then output of every layer is a function of \mathbf{x} , i.e., $O_i^h = O_i^h(\mathbf{x})$.

Let us denote the vector of dropout weights m_j^h for the h -th layer by $\mathbf{m}_h = (m_1^h, \dots, m_{N_h}^h)^\top$ and the full set of dropout weights by $\mathbf{M} = (\mathbf{m}_1, \dots, \mathbf{m}_K)$. Thus, any neural network $\hat{f}(\mathbf{x})$ with dropout layers essentially has two sets of parameters: the full set of learnable weights \mathbf{W} and the set of dropout weights \mathbf{M} :

$$\hat{f}(\mathbf{x}) = \hat{f}(\mathbf{x} \mid \mathbf{W}, \mathbf{M}).$$

Let us have a neural network with dropout, which was trained on some dataset giving weight estimates $\hat{\mathbf{W}}$. Then dropout weights \mathbf{M} can be considered as free parameters and require selection at the time of inference

$$\hat{f}(\mathbf{x} \mid \mathbf{M}) = \hat{f}(\mathbf{x} \mid \hat{\mathbf{W}}, \mathbf{M}).$$

The originally proposed [18] and currently the standard choice is to take $\hat{\mathbf{M}} = (1-p) \cdot \mathbf{E}$, where \mathbf{E} is the matrix of all ones of the corresponding shape. Such an approach gives the fixed function $\hat{f}(\mathbf{x} \mid \hat{\mathbf{M}})$, which is known to give reasonably good performance in practice. The main intuition behind such choice is the replacement of the stochastic pre-activations S_i^h given by (2) with their expectations, which are exactly equal to (1).

Recently, it was proposed to consider dropout as a variational approximation in a specially chosen Bayesian model, see [10]. Within this approach, one can sample T i.i.d. realizations $\mathbf{M}_1, \dots, \mathbf{M}_T \sim \text{Bernoulli}(1-p)$ and compute approximate posterior mean and variance

$$\bar{f}_T(\mathbf{x}) = \frac{1}{T} \sum_{i=1}^T \hat{f}(\mathbf{x} \mid \mathbf{M}_i), \quad \bar{\sigma}_T^2(\mathbf{x}) = \frac{1}{T} \sum_{i=1}^T (\hat{f}(\mathbf{x} \mid \mathbf{M}_i) - \bar{f}_T(\mathbf{x}))^2.$$

The approximate posterior variance $\bar{\sigma}_T^2(\mathbf{x})$ is a natural choice for the uncertainty estimate and was successfully used in the variety of applications such as out-of-distribution detection [40] and active learning [11].

In this paper, we suggest a different approach, namely we treat $\hat{f}(\mathbf{x} \mid \mathbf{M})$ as an ensemble of models indexed by dropout masks \mathbf{M} . Such a view allows us to decouple inference from training and pose an intuitive question: what set of masks $\mathbf{M}_1, \dots, \mathbf{M}_T$ should one choose in order to obtain the best uncertainty estimate $\bar{\sigma}_T^2(\mathbf{x})$?

Importantly, here we do not limit the selection of masks to be samples from standard dropout distribution, which, in principle, should allow us to obtain better estimates. However, the design of mask selection procedure is a non-trivial problem, which we discuss below in detail.

Remark 1. The standard approach in the literature is to consider an ensemble of models trained on different subsets of the data set or just from different random initializations giving the set of parameter estimates $\hat{\mathbf{W}}_1, \dots, \hat{\mathbf{W}}_T$ and corresponding approximations $\hat{f}(\mathbf{x} \mid \hat{\mathbf{W}}_i, \hat{\mathbf{M}})$, $i = 1, \dots, T$. Similarly, one can compute the variance $\bar{\sigma}_T^2(\mathbf{x})$, which was shown to be a reasonable uncertainty estimate in practice [5, 36]. The main drawback of this approach is the need to train and store T different models, which might be very costly both in terms of computation and storage needed.

2.2 Data-Driven Mask Generation Under General Sampling Distributions

In practice, many neurons in the network are highly correlated. For example, consider a correlation matrix of neurons in a hidden layer of a fully-connected neural network, trained on the regression dataset (see Fig. 1a). The correlation matrix was computed on the test set and clearly shows groups of highly correlated neurons. Sampling masks for this layer uniformly at random might result in a high variance of pre-activations (2). As a result, the estimates for the whole network may require a significant number of samples (stochastic passes through the NN) T to converge. We illustrate this behaviour on Fig. 1b, where several hundreds of simple MC dropout estimates are required for the convergence of the log-likelihood values. It is clearly seen that a larger number of samples improves the values of log-likelihood, yet may impose computational cost too large to be used in real-world applications. However, one may expect that the knowledge about the correlations between neurons can help to sample more diverse neurons and improve the estimates.

In what follows, we consider the probabilistic generation of masks \mathbf{m}_h from some distribution $P^{(h)}$ with possibly non-i.i.d. distributions of components. Similarly to the case of dropout, we suggest using an unbiased estimate of the layer-wise mean. Our main motivation is to approximately preserve the average performance of the trained network. The construction of the unbiased estimator is non-trivial and is given by celebrated Horvitz-Thompson (HT) estimator [19]:

$$S_i^h = \sum_{j=1}^{N_h-1} \frac{1}{\pi_j^h} m_j^h w_{ij}^h O_j^{h-1}, \quad i = 1, \dots, N_h, \quad (3)$$

where π_j^h is the marginal probability of value 1 for the random variable m_j^h .

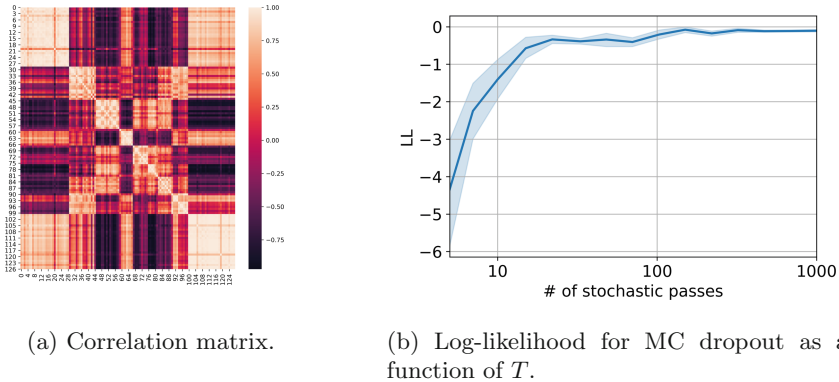


Fig. 1. (a) Correlation matrix C between the outputs of the neurons in a hidden layer of the NN trained on the *naval propulsion* dataset. (b) For the same dataset log-likelihood computed via MC dropout increases with increase of the number of stochastic passes T . More than 100 samples are needed to reach convergence.

2.3 Diversity Sampling Approaches

Let us consider h -th hidden layer of the neural network with dropout. Assume that we have access to the correlations

$$C_{ij}^{(h)} = \text{corr}_{\mathbf{x}}\{O_i^h(\mathbf{x}), O_j^h(\mathbf{x})\}, \quad i, j = 1, \dots, N_h.$$

In practice, we compute an empirical correlation based on some set of points, which represents the data distribution well enough. As a result, we obtain the correlation matrix $C^{(h)} \in \mathbb{R}^{N_h \times N_h}$ between the neurons of the h -th hidden layer. Below we discuss several approaches to sampling neurons in a way that the correlation between sampled neurons is as small as possible. We note that instead of the correlation matrix $C^{(h)}$ one may consider the covariance matrix $K^{(h)}$ in any of the approaches described below. The properties of the methods significantly depend on the choice of the matrix, and we will perform the empirical evaluation of the methods based on each of them in the experiments.

Leverage Score Sampling. A basic approach for non-uniform sampling of rows and columns in kernel matrices is the so-called *leverage score sampling* [1]. In this approach, the neurons are sampled independently with different probabilities π_j^h :

$$\pi_j^h \sim \ell_\lambda^{(h)}(j) = \left[C^{(h)}(C^{(h)} + \lambda I)^{-1} \right]_{jj}, \quad j = 1, \dots, N_h,$$

where the quantities $\ell_\lambda^{(h)}(j)$ are called *leverage scores*. This approach makes neurons from large and highly correlated clusters to be sampled less frequently. In Sect. 3, we show that leverage score sampling indeed allows obtaining better

uncertainty estimates for out-of-distribution data in regression tasks compared to MC dropout. However, its performance for in-domain data is even inferior to uniform sampling. In the next section, we propose a more complex approach, which allows to significantly improve the quality of uncertainty estimation.

Sampling with Determinantal Point Processes. Determinantal Point Processes (DPPs) [24] are specific probability distributions over configurations of points that encode diversity through a kernel function. They were introduced in [28] for the needs of statistical physics and were used for a number of ML applications, see [24] for an overview. DPP can be seen as a probabilistic MaxVol algorithm [14] of finding a maximal-volume submatrix.

We use correlation matrix $C^{(h)}$ as the likelihood kernel for DPP. Then, given a set S of selected points for a mask distribution $\mathbf{m}_h \sim DPP(C^{(h)})$, we obtain

$$\mathbb{P}[\mathbf{m}_h = S] = \frac{\det[C_S^{(h)}]}{\det[C^{(h)} + I]}, \quad h = 1, \dots, K,$$

where $C_S^{(h)} = [C_{ij}^{(h)}, i, j \in S]$, i.e., a square submatrix of $C^{(h)}$ obtained by keeping only rows and columns indexed by S .

To better understand the DPP, let us come back to the correlation matrix depicted in Fig. 1a. The probability for DPP to take highly correlated neurons into the sample S is low as, in this case, the corresponding determinant $\det C_S^{(h)}$ will have a small value. Thus, DPP tends to sample neurons from different clusters, increasing an overall diversity.

From computational point of view, DPP-sampling requires $O(N_h^3)$ operations for generating each sample. It is quite expensive but completely viable even for modern large networks which usually have up to 1024 neurons in fully-connected layers. Importantly, masks can be precomputed once, and then the same masks are used on the inference stage for every test sample with no additional overhead. Also, computations in last fully-connected layers with dropout usually require only few percents of the total computational budget in ImageNet-size networks. Therefore, a computational overhead caused by the DPP-sampling does not have a significant impact on the inference time.

K-DPP. The k-DPP [24] is a variation of the DPP, conditioned to produce samples of fixed size $|S| = k$. With the cost of introducing an additional parameter, it allows us to tune the sampling procedure as the choice of k apparently has a significant influence on the result. In this work, we use for the h -th layer $k^{(h)} = (1 - p)N_h$, so that the number of neurons in the sample is equal to the mean number of neurons in the sample of MC-Dropout. In the case of k-DPP, the computation of the marginal probabilities π_j^h for HT-estimator (3) is non-trivial and requires the separate optimization procedure, see the details in [2].

2.4 Diversification for Uncertainty Estimation in Classification

For regression, the variance of prediction is a standard uncertainty measure. However, uncertainty estimation for classification is, in some sense, more challenging than for regression as there is no obvious candidate for uncertainty measure.

Let us define the average probability for the class prediction by ensemble members $\bar{p}_T(y = c | \mathbf{x}) = \frac{1}{T} \sum_{i=1}^T p(y = c | \mathbf{x}, \mathbf{M}_i)$. The standard uncertainty measure usually considered in the literature is

$$s(\mathbf{x}) = 1 - \max_c \bar{p}_T(y = c | \mathbf{x}),$$

which is based solely on the mean probabilities predicted by the ensemble. While providing good results in practice [3, 37] it doesn't use the information about the variation of predictions between ensemble members.

In our work, we consider *BALD* [20] uncertainty measure and combine it with different sampling schemes considered above. BALD is equal to the mutual information between outputs and model parameters:

$$I(\mathbf{x}) = H(\mathbf{x}) - \frac{1}{T} \sum_{c=1}^C \sum_{i=1}^T -p(y = c | \mathbf{x}, \mathbf{M}_i) \log(p(y = c | \mathbf{x}, \mathbf{M}_i)),$$

where $H(\mathbf{x}) = -\sum_{c=1}^C \bar{p}_T(y = c | \mathbf{x}) \log(\bar{p}_T(y = c | \mathbf{x}))$ is an entropy of the ensemble mean. Importantly, BALD values are directly linked with the diversity of the ensemble members, and therefore are well suited for combination with our approach.

3 Experiments

3.1 Uncertainty Estimation for Regression


Models and Metrics. For the experiments, we consider MC dropout as a baseline and all the proposed UE methods discussed in the Sect. 2.3: leverage score sampling, DPP and k-DPP. We present the results for leverage score sampling and DPP based on correlation matrix and k-DPP based on covariance matrix as such a choices give consistently better results compared to an alternative. For leverage score sampling we deliberately choose $\lambda = 1$ to make it working with de-facto the same matrix as DPP-based methods. All the regression models were trained with RMSE as a loss function. We used feed-forward NNs with 3 hidden layers (128-128-64 neurons) and leaky ReLU activation function [27]. For DPP-based methods, we use the DPPy implementation provided in [13].

We should note that we do not compare with fully Bayesian approaches as we are focusing on the solutions applicable to the standard dropout-based models without changing model architecture and training procedure. Following [17, 22], we use log-likelihood of Gaussian distribution with mean and variance computed by different methods as a quality measure.

On top of single models, we also consider a straightforward ensemble approach with NNs trained exactly the same way as single models but from different

Table 1. Summary of the UCI datasets used in experiments, see [7].

Dataset name	Naval propulsion	Concrete	Boston housing	Kin8nm	Ccpp	Red wine
Samples	11934	1030	506	8192	9568	1599
Features	16	8	13	8	4	11

**Fig. 2.** Log-likelihood metric across various UCI datasets for NN UE models with different number of stochastic passes $T = 10, 30, 100$. DPP and k-DPP give better results compared to other methods with DPP working well already for $T = 10$ and consistently showing lower variance.

random initializations. Our experiments show that uncertainty estimates based on ensembling of networks without sampling in individual networks doesn't work for well for the considered regression datasets.

Experiments on Regression Datasets. Similarly to [22], we run a series of experiments on various regression datasets, see Table 1 for the full list of datasets. We start with in-domain uncertainty estimation: for each dataset, random 50% of points were used for training and other 50% for testing. The log-likelihood values are averaged over testing set. Multiple experiments are done via 5 random train-test splits, 2-fold cross-validation and 5 runs of the training procedures for every model (resulting in 50 average log-likelihood values contributing to each boxplot). Uncertainty estimates were computed for different number of stochastic passes $T = 10, 30$ and 100 for every model.

We show the resulting distributions of log-likelihood values for each dataset on Fig. 2. We observe that either DPP or k-DPP always show the best results. Most importantly, DPP works very well already for small number of stochastic passes $T = 10$ and consistently has low variance which is extremely important for practical usage.

We also performed an experiment with out-of-distribution (OOD) data. To generate OOD data we pick a random feature and split the data into the train set and OOD set by the median value on this feature. The experiments were run for 5 different splits. For OOD data good uncertainty estimates should have on average higher values compared to in-domain data. Table 2 provides for *concrete* dataset the percentages of OOD points with UE values higher than α percentile of UE distribution for training data ($\alpha = 80\%, 90\%, 95\%$). The resulting numbers

Table 2. Percentages of OOD points with UE values higher than specified percentile of UE distribution for training data for *concrete* dataset. DPP and k-DPP show the best results based on average values (top-2 average values are put in **bold**). For all the methods $T = 100$.

Percentile	MC dropout	Leverage	DPP	k-DPP
80	55.0 ± 27.6	61.3 ± 27.7	70.4 ± 26.0	71.9 ± 28.0
90	46.0 ± 30.7	52.9 ± 30.8	59.6 ± 30.1	60.8 ± 33.7
95	40.6 ± 32.1	46.5 ± 33.1	52.1 ± 32.9	51.8 ± 36.3

should be considered with a significant grain of salt due to their high variance but still DPP and k-DPP show the best results based on average values.

3.2 Uncertainty Estimation for Classification

Data, Models and Metrics. In this section, we aim to show the applicability of the proposed methods to the classification tasks. We take *BALD* [20] as an uncertainty estimate. We consider three datasets: MNIST, which is a toy dataset of handwritten digits [26], CIFAR-10, which is a 10-class image dataset with simple objects [23], and ImageNet [6], the large scale image classification dataset. Importantly, for MNIST we use only 500 train samples, otherwise the models would have too good accuracy and uncertainty estimation for in-domain data would not be relevant. For CIFAR-10 we use 50'000 samples for training and 10'000 for testing. For the MNIST dataset, we use a simple convolutional neural network with two convolutional layers, max-pooling and two fully connected layers. For the CIFAR-10 we use a more powerful network with 6 convolutional layers and batch normalization. Finally, for ImageNet we use the pre-trained ResNet-18 neural network [16] from PyTorch [33]. Dropout with rate $p = 0.5$ is used before the last fully-connected layer in all the cases. $T = 100$ stochastic passes were made for every model. The experiments are repeated three times with different seeds for the models.

Experimental Results. For in-domain uncertainty estimation the results are presented via UE-accuracy curve, see Fig. 3. It assumes that samples with lower uncertainty will be classified with a higher average accuracy. It can be clearly seen that DPP significantly outperforms all the competitors on every dataset. We should emphasize that the superiority of DPP is especially strong for ImageNet, where the usage of DPP required only 2% computational overhead compared to MC dropout according to our experiments.

We also consider detection of out-of-distribution samples which is one of the important problems for the uncertainty estimation. As OOD samples we use fashion-MNIST [41] and SVHN images [32] for MNIST and CIFAR-10 correspondingly. We use count-vs-uncertainty curve and expect there should be few




Fig. 3. UE-accuracy curve (the higher curve – the better). We select the samples with low uncertainty to assure that the accuracy is higher for them.




Fig. 4. Count-vs-uncertainty curve for out-of-distribution data (the lower curve – the better).

points with the low uncertainty for good uncertainty estimation methods. The results are presented in Fig. 4. We see that DPP-based approach allows to detect the OOD samples better for the both considered datasets.

4 Related Work

Dropout [18, 38] has emerged in recent years as a technique to prevent the overfitting in deep and overparametrized neural networks. Over the years, it obtained theoretical explanations as an averaged ensembling technique [38], a Bernoulli realization of the corresponding Bayesian neural network [10] and a latent variable model [30]. It was shown in [9, 31] that using dropout at the prediction stage (i.e., stochastic forward passes of the test samples through the network, also referred to as *MC dropout*) leads to unbiased Monte-Carlo estimates of the mean and variance for the corresponding Bayesian neural network trained using variational inference. These uncertainty estimates were shown to be efficient in different scenarios [9, 39].

Training an ensemble of models and uncertainty estimation by their disagreement is another common approach [25]. It is shown that with few models in an ensemble, you can get robust and useful calibrated results [5], outperforming MC dropout in active learning and error detection. The main disadvantage of ensembles is the necessity to train multiple model instances. However, it was addressed in recent works [12, 21, 29] which consider different strategies for speeding up ensemble construction. Recently, it was shown that improving diversity of ensemble members improves the quality of the resulting uncertainty estimates [22]. We also mention recent works which thoroughly investigate in-domain [3] and out-of-domain [37] uncertainty estimation in classification for the case of maximum probability uncertainty estimate.

5 Conclusions

We have proposed a new approach that strengthens the dropout-based uncertainty estimation for neural networks. Instead of randomly sampling the dropout masks on the inference stage, we sample special sets of diverse neurons via determinantal point processes that utilize the information about the correlations of neurons in the inner layers. Numerical experiments on a wide range of regression and classification tasks show that uncertainty estimates based our approach outperform the MC dropout and other baselines with a significant margin. A combination of dropout-based inference with ensembling of several models allows to further improve the quality of the proposed uncertainty estimates and achieve state-of-the-art performance. From the practical perspective, our method is simple to implement as it does not require any modifications to the neural network architecture and the training process. Importantly, the proposed uncertainty estimates have high quality even for a small number of stochastic passes through the network making the inference stage even faster in practice.

We expect that the proposed methods of dropout mask sampling may also be used on the training stage, leading to more robust and efficient models. Another compelling direction of further research is approximate DPP sampling, which may increase the sampling speed of the proposed approaches, making them more production-friendly, as in [35].

The code reproducing the experiments is available at Github¹.

Acknowledgements. The research was carried out at Skoltech and supported by the Russian Science Foundation (project no. 21-11-00373). The authors want to thank Nikita Mokrov for useful discussions. M.P. and K.F. acknowledge the use of “Zhores” supercomputer [42] for obtaining the part of results presented in this paper.

References

1. Alaoui, A., Mahoney, M.W.: Fast randomized kernel ridge regression with statistical guarantees. In: NIPS, pp. 775–783 (2015)

¹ <https://github.com/stat-ml/dpp-dropout-uncertainty/>.

2. Amblard, P.O., Barthelmé, S., Tremblay, N.: Subsampling with k determinantal point processes for estimating statistics in large data sets. In: 2018 IEEE Statistical Signal Processing Workshop (SSP), pp. 313–317. IEEE (2018)
3. Ashukha, A., Lyzhov, A., Molchanov, D., Vetrov, D.: Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. In: ICLR (2019)
4. Begoli, E., Bhattacharya, T., Kusnezov, D.: The need for uncertainty quantification in machine-assisted medical decision making. *Nat. Mach. Intell.* **1**(1), 20–23 (2019)
5. Beluch, W.H., Genewein, T., Nürnberger, A., Köhler, J.M.: The power of ensembles for active learning in image classification. In: CVPR, pp. 9368–9377 (2018)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: CVPR, pp. 248–255. IEEE (2009)
7. Dua, D., Taniskidou, E.K.: UCI machine learning repository (2017). <http://archive.ics.uci.edu/ml>
8. Feng, D., Rosenbaum, L., Dietmayer, K.: Towards safe autonomous driving: capture uncertainty in the deep neural network for lidar 3D vehicle detection. In: ITSC, pp. 3266–3273 (2018)
9. Gal, Y.: Uncertainty in deep learning. University of Cambridge (2016)
10. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: ICML, pp. 1050–1059 (2016)
11. Gal, Y., Islam, R., Ghahramani, Z.: Deep Bayesian active learning with image data. In: ICML, pp. 1183–1192 (2017)
12. Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D.P., Wilson, A.G.: Loss surfaces, mode connectivity, and fast ensembling of DNNs. In: NeurIPS, pp. 8789–8798 (2018)
13. Gautier, G., Polito, G., Bardenet, R., Valko, M.: DPPy: DPP sampling with Python. *JMLR* **20**(180), 1–7 (2019)
14. Goreinov, S., Oseledets, I., Savostyanov, D., Tyrtshnikov, E., Zamarashkin, N.: How to find a good submatrix. In: *Matrix Methods: Theory, Algorithms and Applications: Dedicated to the Memory of Gene Golub*, pp. 247–256 (2010)
15. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *Proceedings ICML*, pp. 1321–1330 (2017)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
17. Hernández-Lobato, J.M., Adams, R.: Probabilistic backpropagation for scalable learning of Bayesian neural networks. In: ICML, pp. 1861–1869 (2015)
18. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. *arXiv arXiv:1207.0580* (2012)
19. Horvitz, D.G., Thompson, D.J.: A generalization of sampling without replacement from a finite universe. *JASA* **47**(260), 663–685 (1952)
20. Houlshby, N., Huszár, F., Ghahramani, Z., Lengyel, M.: Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745* (2011)
21. Izmailov, P., Maddox, W., Kirichenko, P., Garipov, T., Vetrov, D., Wilson, A.G.: Subspace inference for Bayesian deep learning (2019)
22. Jain, S., Liu, G., Mueller, J., Gifford, D.: Maximizing overall diversity for improved uncertainty estimates in deep ensembles. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 4264–4271 (2020)
23. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Technical report TR-2009, University of Toronto (2009)
24. Kulesza, A., Taskar, B., et al.: Determinantal point processes for machine learning. *Found. Trends® Mach. Learn.* **5**(2–3), 123–286 (2012)

25. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: NIPS, pp. 6402–6413 (2017)
26. LeCun, Y.: The MNIST database of handwritten digits (1998). <https://yannlecun.com/exdb/mnist/>
27. Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. In: Proceedings of the ICML 2013, vol. 30, p. 3 (2013)
28. Macchi, O.: The coincidence approach to stochastic point processes. *Adv. Appl. Probab.* **7**(1), 83–122 (1975)
29. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for Bayesian uncertainty in deep learning. In: NeurIPS, pp. 13132–13143 (2019)
30. Maeda, S.: A Bayesian encourages dropout. arXiv [arXiv:1412.7003](https://arxiv.org/abs/1412.7003) (2014)
31. Nalisnick, E., Hernandez-Lobato, J.M., Smyth, P.: Dropout as a structured shrinkage prior. In: International Conference on Machine Learning, pp. 4712–4722 (2019)
32. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
33. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: NeurIPS, pp. 8024–8035 (2019)
34. Settles, B.: Active Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 6, no. 1, pp. 1–114 (2012)
35. Shelmanov, A., Tsymbalov, E., Puzyrev, D., Fedyanin, K., Panchenko, A., Panov, M.: How certain is your transformer? In: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, pp. 1833–1840 (2021). <https://doi.org/10.18653/v1/2021.eacl-main.157>
36. Smith, J.S., Nebgen, B., Lubbers, N., Isayev, O., Roitberg, A.E.: Less is more: sampling chemical space with active learning. *J. Chem. Phys.* **148**(24), 241733 (2018)
37. Snoek, J., et al.: Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In: NeurIPS, pp. 13969–13980 (2019)
38. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
39. Tsymbalov, E., Panov, M., Shapeev, A.: Dropout-based active learning for regression. In: International Conference on Analysis of Images, Social Networks and Texts, pp. 247–258 (2018)
40. Vyas, A., Jammalamadaka, N., Zhu, X., Das, D., Kaul, B., Willke, T.L.: Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11212, pp. 560–574. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01237-3_34
41. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms (2017)
42. Zacharov, I., et al.: “Zhores”—petaflops supercomputer for data-driven modeling, machine learning and artificial intelligence installed in Skolkovo Institute of Science and Technology. *Open Eng.* **9**(1), 512–520 (2019)