

---

# The Peril of Popular Deep Learning Uncertainty Estimation Methods

---

Yehao Liu\*  
EPFL

Matteo Pagliardini\*  
EPFL

Tatjana Chavdarova  
UC Berkeley

Sebastian U. Stich  
CISPA

## Abstract

Uncertainty estimation (UE) techniques—such as the Gaussian process (GP), Bayesian neural networks (BNN), Monte Carlo dropout (MCDropout)—aim to improve the interpretability of machine learning models by assigning an estimated uncertainty value to each of their prediction outputs. However, since too *high* uncertainty estimates can have fatal consequences in practice, this paper analyzes the above techniques. Firstly, we show that GP methods always yield *high* uncertainty estimates on out of distribution (OOD) data. Secondly, we show on a 2D toy example that both BNNs and MCDropout do *not* give high uncertainty estimates on OOD samples. Finally, we show empirically that this pitfall of BNNs and MCDropout holds on real world datasets as well. Our insights (i) raise awareness for the more cautious use of currently popular UE methods in Deep Learning, (ii) encourage the development of UE methods that approximate GP-based methods—instead of BNNs and MCDropout, and (iii) our empirical setups can be used for verifying the OOD performances of any other UE method. The source code is available at <https://github.com/epfml/uncertainty-estimation>.

## 1 Introduction

While the complexity of deep learning models allows them to outperform humans on a growing number of tasks, it often renders them *uninterpretable* to human users. This limits the use of these methods in applications where decisions can have significant consequences, such as for instance in the fields of health care, finance or autonomous driving. Uncertainty estimation (UE) methods aim to improve model interpretability, by associating an estimate of its uncertainty to each output [Kim et al., 2016]. As such, uncertainty estimates are expected to be high in value whenever we give as input a sample that is out of the distribution (OOD) of the dataset that the model was trained with.

Gaussian processes (GP) are considered as the gold standard for UE [Rasmussen and Williams, 2005, van Amersfoort et al., 2021]. Similarly, Bayesian Neural Networks (BNNs) yield mathematically grounded UE methods that extends standard neural network (NN) training to posterior inference. However GPs and BNNs do not scale well with the dimension of the input data and the parameters space, respectively, and are often infeasible for real-world problems. Thus, one of the most popular methods in deep learning is Monte Carlo Dropout (MCDropout), which method applies Dropout [Srivastava et al., 2014] at inference time to compute uncertainty estimates. An active branch of work is focused on finding better and more efficient approximation methods for UE. In order for the community to move forward in the right direction, it is important to verify that the these UE methods work well and to understand their differences and flaws.

**Overview of contributions.** We (i) show on a simplistic 2D example that surprisingly *both* BNNs and MCDropout methods might estimate low uncertainty for OOD data, while GP can detect OOD

---

\*Equal contribution.

Contact: {yehao.liu, matteo.pagliardini}@epfl.ch, tatjana.chavdarova@berkeley.edu, stich@cispa.de.

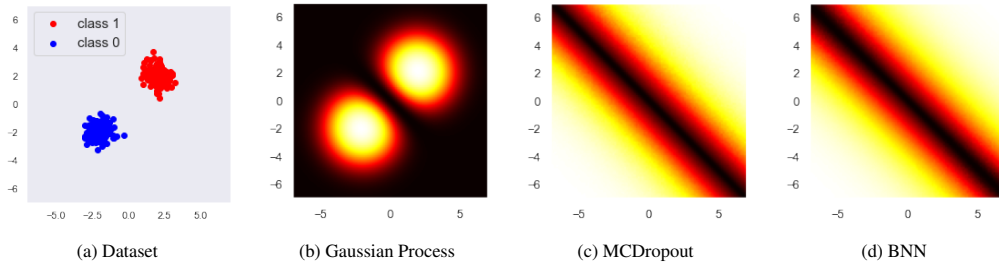


Figure 1: Uncertainty estimates by GP and MCDropout on a 2D dataset, depicted in (a). We display the uncertainty estimates of GP (b), MCDropout (c), and BNN trained with HMC (d), for points of the 2D input space (the darker, the higher the uncertainty). See § 3.1 for discussion and App. B for implementation details.

samples. We (ii) we empirically demonstrate that the above perils of BNNs and MCDropout occur in real-world situations such as on MNIST—for OOD data generated by interpolation between two real samples—and ResNet-18 trained on CIFAR-10—for samples from CIFAR-100 that were not used during training. Finally, we (iii) argue analytically on a simple example why GP methods succeed.

## 2 Background and Related Works

Gaussian processes [GPs, Rasmussen and Williams, 2005] are a non-parametric distance aware output function. GPs model the similarity between data points with a kernel function, and use the Bayes rule to model a *distribution over functions* by maximizing the marginal likelihood [Rasmussen and Williams, 2005]. As such, GPs require access to the full dataset at inference time, and although there exist approximations, this family of methods does *not* scale well with the dimension of the data.

Bayesian Neural Networks (BNNs) are stochastic neural networks trained using a Bayesian approach. While standard neural network training performs a maximum likelihood estimation (MLE) of the parameters of the network  $\omega \in \Omega$ , training BNNs extends to estimating the posterior distribution. While the formulation is straightforward (see App. A), due to the integration with respect to the whole parameter space  $\Omega$  this computation is intractable for DNNs. In this work, we train the BNNs with two common training methods: Mean-Field Variational Inference (MFVI) Blei et al. [2017], and Hamiltonian Monte Carlo (HMC) Neal [2012]—the latter expected to be more accurate.

Lastly, Monte Carlo Dropout (MCDropout) applies Dropout at inference time allows for computationally efficient UE. Gal and Ghahramani [2016] show that *Dropout* [Srivastava et al., 2014] approximates Bayesian inference of a GP. Given  $M$  models  $\{\mathcal{C}_\omega^m\}_{m=1}^M$  sampled with MC Dropout, where each outputs non-scaled values—“logits”, we define  $\hat{\mathbf{y}} \in \mathbb{R}^C$  as the average prediction:  $\hat{\mathbf{y}} = \frac{1}{M} \sum_{m=1}^M \text{Softmax}(\mathcal{C}_\omega^m(\mathbf{x}))$ . Given  $\hat{\mathbf{y}}$ , there are several ways to estimate the model’s uncertainty, see [Gal, 2016, §3.3]. We use the *entropy* of the output distribution (over the classes) to quantify the uncertainty estimate of a given sample  $\mathbf{x}$ :  $\mathcal{H}(\mathbf{x}, \omega) = -\sum_{c \in C} \hat{y}_c \log \hat{y}_c$ .

**Relevant Works.** Closest to ours are [van Amersfoort et al., 2020] and the recent [van Amersfoort et al., 2021] where the authors point out on a 2D toy example that the uncertainty estimates of Deep Ensembles [Lakshminarayanan et al., 2017] and Deep Kernel Learning [Wilson et al., 2015] models, respectively, are *low* on points that are “far away” from the regions where the training data are located. Foong et al. [2019] point out limitations in the expressiveness of the predictive uncertainty estimate given by mean-field variational inference, in particular that the method fails to give calibrated uncertainty estimates in between separated regions of observations.

## 3 The Failure Mode of Popular Uncertainty Estimation Methods

We now first illustrate the behaviour of GP and MCDropout on a set of examples—see details on the experimental setups in App. B, and then study the GP behaviour theoretically.

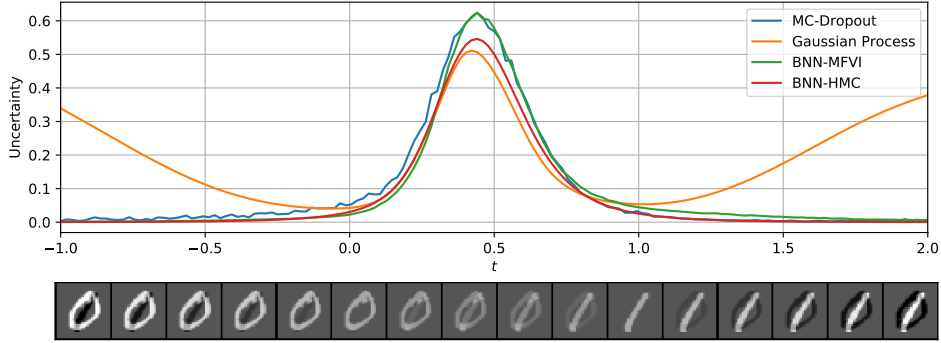


Figure 2: Estimated uncertainty by GP, BNN and MCDropout. We train a GP model, a NN classifier with dropout, and two BNNs on MNIST handwritten digits of 0 and 1’s and display the uncertainty estimation on linear interpolated images of the form  $t \cdot \mathbf{x}^1 + (1 - t) \cdot \mathbf{x}^0$  for  $t \in [-1, 2]$  (grayscale values adjusted for better display in the illustration). The BNNs were trained with two different training methods: Mean-Field Variational Inference (MFVI), and Hamiltonian Monte Carlo (HMC). All methods give high uncertainty for convex combinations of training data ( $t \in [0, 1]$ ), highest uncertainty for  $t \approx 0.5$ . GP detects OOD data for  $t \notin [0, 1]$  while BNN/MCDropout estimates both have low uncertainty. This demonstrates that BNN and MCDropout fail to detect this kind of OOD data.

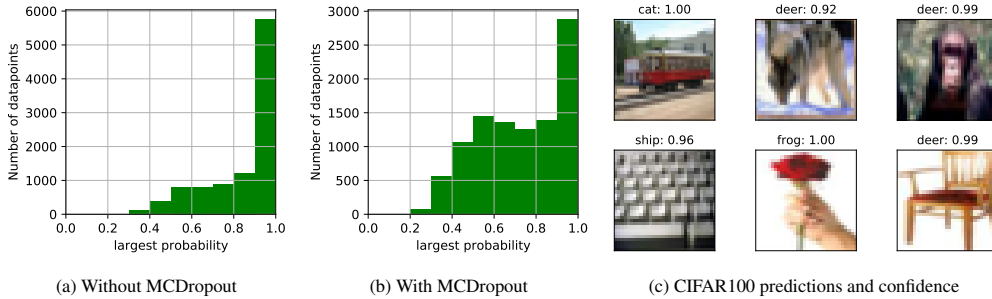


Figure 3: MCDropout struggles detecting OOD samples: we trained a ResNet-18 on CIFAR-10 and used it to classify images from CIFAR-100. We observe in (b) that many CIFAR-100 datapoints, which are outside of the training distribution, are nonetheless inducing a confident prediction, despite using MCDropout. We show in (c) some CIFAR-100 images along with their class and confidence, obtained sampling 20 models via MCDropout.

### 3.1 Numerical Illustrations

**Motivating Two-Dimensional Example (Figure 1).** We consider a binary classification problem with two-dimensional bi-modal input data, generated from  $\mathcal{N}(-2, -2)$  and  $\mathcal{N}(2, 2)$  with label 1 and 0, resp. A Bayes optimal classifier is given by  $\text{sign}(x - y)$  for  $(x, y)$ -coordinates. We train a GP model and a two-layer perceptron on the training data. Both resulting models can perfectly classify the training data. We observe in Fig. 1 that GP gives high uncertainty estimates for input data far away from the two modes, while BNNs and MCDropout both assign high uncertainty only near to the decision boundary ( $x = -y$ ) and fail to identify many OOD regions of the input space.

**MNIST data (Figure 2).** We train GP and MCDropout on the digits of class 0 and 1 of the MNIST [LeCun et al., 2010] dataset. We first randomly select two data points of each class,  $\mathbf{x}^0, \mathbf{x}^1$ . We then test the models on artificial data generated by linear combination of the form  $t \cdot \mathbf{x}^1 + (1 - t) \cdot \mathbf{x}^0, t \in [-1, 2]$ . Although these samples are not necessarily realistically looking images, we expect that the uncertainty estimates are high when  $t$  is close to 0.5, and for  $t \notin [0, 1]$ . However, in Fig. 2 we see that GP successfully detects OOD data for  $t \notin [0, 1]$  while MCDropout fails. However, we observed that MCDropout is able to detect other kinds of OOD data, such as e.g. the MNIST digits 2–9 that were not used for training the model (See Tab. 1 in App. C.) To follow this experimental protocol for an arbitrary method, we provide some helper code, see App. B.4 for details.

**CIFAR data (Figure 3).** Lastly, we train ResNet-18 [He et al., 2016a] with dropout on CIFAR-10 [Krizhevsky, 2012] and measure the uncertainty with MCDropout on samples from CIFAR-100. In Fig. 3 we observe that although MCDropout improves over model without MCDropout enabled, it yet fails to detect these OOD samples.

### 3.2 GP Analysis on OOD samples

The following result shows that for a test data sample that is not correlated with the training data—measured with the kernel—the estimated uncertainty is *high*, the proof is omitted and given in App. D.

**Theorem (informal) 1.** *Consider a  $n$ -dimensional binary GP classification model with kernel  $K(\mathbf{x}, \mathbf{x}')$ . Let  $\mathbf{k}_*$  denote the similarity vector  $(\mathbf{k}_*)_i = K(\mathbf{x}_i, \mathbf{x}_*)$  between a test sample  $\mathbf{x}_*$  and training data  $\mathbf{x}_i, i = 1, \dots, n$ . If  $\|\mathbf{k}_*\|$  is small—i.e. if the similarity between the test sample and the training data points is low—then the GP prediction for  $\mathbf{x}_*$  will have high uncertainty.*

This theorem just argues one direction and in particular does not show that the uncertainty is low only for in-distribution data. Moreover, if the kernel overestimates the similarity of the data, OOD samples cannot be detected. We leave the extension of this statement to higher dimensions and more general models to future work.

## 4 Discussion & Conclusion

We found that the GP model performed well in OOD detection in all the tasks and experiments we studied. We analyze GPs analytically and argue that access to a kernel function that measures similarity to the training data can explain this behavior. In contrast, BNNs and standard NN might only memorize training data implicitly and struggle to detect some OOD data points.

While an active line of work focuses on finding alternative approximations of BNNs [Ober and Aitchison, 2021, Mbuva et al., 2021], in this work we point out that BNNs and MCDropout are unable to detect certain OOD samples. Our results advocate the need for efficient UE techniques that approximate GP-based methods and that MCDropout’s predictions should be taken with a grain of salt—although this is the most efficient and frequently used method in practice. We believe that our contribution will lead the community to collaborate on new benchmarks and tests of UE methods, in order to develop methods with theoretically guaranteed properties. This is necessary to develop future-proof UE technologies that can be used for safety-critical applications in the years to come.

## Acknowledgements

The authors would like to thank Prof. M. Jaggi (EPFL) for his support. TC is supported by the Swiss National Science Foundation (SNSF), grant P2ELP2\_199740.

## References

- R. F. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani. Predictive inference with the jackknife+, 2020.
- M. J. Betancourt and M. Girolami. Hamiltonian Monte Carlo for hierarchical models. *arXiv preprint arXiv:1312.0906*, 2013.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- P. Esposito. BLiTz—Bayesian Layers in Torch Zoo (a Bayesian Deep Learning library for Torch). <https://github.com/piEsposito/blitz-bayesian-deep-learning/>, 2020.
- A. Y. K. Foong, Y. Li, J. M. Hernández-Lobato, and R. E. Turner. ‘In-Between’ Uncertainty in Bayesian Neural Networks. In *ICML 2019 Workshop on Uncertainty and Robustness in Deep Learning*, 2019.
- Y. Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.

- Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016a.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016b.
- L. V. Jospin, W. Buntine, F. Boussaid, H. Laga, and M. Bennamoun. Hands-on bayesian neural networks – a tutorial for deep learning users. *arXiv preprint arXiv:2007.06823*, 2021.
- A. Kendall and Y. Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Advances in Neural Information Processing Systems*, page 5580–5590. Curran Associates Inc., 2017.
- B. Kim, R. Khanna, and O. O. Koyejo. Examples are not enough, learn to criticize! Criticism for Interpretability. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2014.
- A. Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 05 2012.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 2017.
- Y. LeCun, C. Cortes, and C. Burges. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- R. Mbuwva, W. T. Mongwe, and T. Marwala. Separable Shadow Hamiltonian Hybrid Monte Carlo for Bayesian Neural Network Inference in wind speed forecasting. *Energy and AI*, 6:100108, 2021.
- R. M. Neal. MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, 2012.
- S. W. Ober and L. Aitchison. Global inducing point variational posteriors for Bayesian Neural Networks and Deep Gaussian processes. In *International Conference on Machine Learning*, pages 8248–8259. PMLR, 2021.
- I. Osband. Risk versus uncertainty in deep learning: Bayes, bootstrap and the dangers of dropout. *Workshop on Bayesian Deep Learning*, 2016.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037, 2019.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- J. van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal. Simple and scalable epistemic uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, 2020.
- J. van Amersfoort, L. Smith, A. Jesson, O. Key, and Y. Gal. On feature collapse and deep kernel learning for single forward pass uncertainty. *arXiv preprint arXiv:2102.11409*, 2021.
- C. Williams and D. Barber. Bayesian classification with Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, 1998.
- A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning, 2015.

## A Overview of Uncertainty Estimation Methods

Below, we first describe in more detail BNNs—see also [Kendall and Gal, 2017] for a full review of UE methods, and then list relevant works to ours.

### A.1 Bayesian Neural Networks

**Bayesian Neural Networks (BNNs).** Given  $N$  training datapoints  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ , training BNNs extends to posterior inference by estimating the posterior distribution:

$$p(\boldsymbol{\omega}|\mathcal{D}) = \frac{p(\mathcal{D}|\boldsymbol{\omega})p(\boldsymbol{\omega})}{p(\mathcal{D})}, \quad (1)$$

where  $p(\boldsymbol{\omega})$  denotes the prior distribution on a parameter vector  $\boldsymbol{\omega} \in \Omega$ . Given a new sample  $\mathbf{x}^*, y^*$ , the predictive distribution is then:

$$p(y^*|\mathbf{x}^*, \mathcal{D}) = \int_{\Omega} p(y^*|\mathbf{x}^*, \boldsymbol{\omega})p(\boldsymbol{\omega}|\mathcal{D})d\boldsymbol{\omega}. \quad (2)$$

In § 2 we point out that BNNs are often computationally expensive, what arises due to the above integration with respect to the whole parameter space  $\Omega$ .

**Training BNNs.** Two methods exist to train BNNs:

- *Markov Chain Monte Carlo (MCMC) algorithm*—samples the posterior directly, however requires to cache a collection of samples  $\{\boldsymbol{\omega}_k\}_{k=1}^K$ .
- *Variational inference approach (“Variational Bayes”)*—learns a variational distribution  $q(\boldsymbol{\omega})$  to approximate the exact posterior.

Regarding the former, one way to use posterior uncertainty is to sample a set of values  $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_K$  from a posterior  $p(\boldsymbol{\omega}|\mathcal{D})$ , and then average their predictive distributions:

$$p(y^*|\mathbf{x}^*, \mathcal{D}) \approx \sum_{k=1}^K p(y^*|\mathbf{x}^*, \boldsymbol{\omega}_k).$$

In the context of BNNs most popular choice to replace the sampling from the posterior is *Hamiltonian MC (HMC)* [see e.g., Neal, 2012, Betancourt and Girolami, 2013], also known as *hybrid Monte Carlo*, which method uses the derivatives of the density function being sampled to generate efficient transitions spanning the posterior. HMC uses an approximate Hamiltonian dynamics simulation based on numerical integration which is then corrected by performing a Metropolis acceptance step. Unfortunately, HMC does not scale to large datasets, because it is inherently a batch algorithm—it requires visiting the entire training set for every update.

As the former MCMC approach requires revisiting each data point for each update, the latter approximate approach of Variational inference scales better, and thus this approach gained a lot of popularity in the context of BNNs. Variational Bayes approximates the complicated posterior distribution with a “simpler” variational approximation  $q(\boldsymbol{\omega})$ , for example a Gaussian posterior with a diagonal covariance, (i.e. fully factorized Gaussian), and in that case each parameter of the model has its own mean and variance. Analogously to variational autoencoders [VAEs, Kingma and Welling, 2014], we define a variational lower bound:

$$\log p(\mathcal{D}) \geq \mathcal{F}(q) \triangleq \underbrace{\mathbb{E}_{q(\boldsymbol{\omega})}[\log p(\mathcal{D}|\boldsymbol{\omega})]}_{\text{Likelihood term}} - \underbrace{\mathbb{D}_{KL}(q(\boldsymbol{\omega})||p(\boldsymbol{\omega}))}_{\text{KL term}},$$

where  $\mathbb{D}_{KL}$  denotes the Kullback–Leibler (KL) divergence, and the KL term encourages  $q$  to match the prior. Unlike VAEs,  $p(\mathcal{D})$  is fixed, and we are *only* maximizing  $\mathcal{F}(q)$  with respect to the variational posterior  $q$  (i.e. a mean and standard deviation for each weight). Same as for VAEs, the gap equals the KL divergence from the true posterior:  $\mathcal{F}(q) = \log p(\mathcal{D}) - \mathbb{D}_{KL}(q(\boldsymbol{\omega})||p(\boldsymbol{\omega}|\mathcal{D}))$ . Hence, maximizing  $\mathcal{F}(q)$  is equivalent to approximating the posterior.

See for example [Jospin et al., 2021] for more detailed discussion on the differences between these two BNN approaches.

## A.2 Additional Related Works

van Amersfoort et al. [2020] combine Deep Kernel Learning [Wilson et al., 2015] framework and GPs by using NNs to learn low dimensional representation where the GP model is jointly trained, resulting in a method called *Deterministic Uncertainty Quantification* (DUQ). More recently, van Amersfoort et al. [2021] showed that the DUQ method based on DKL can map OOD data close to training data samples, referred as “feature collapse”. The authors thus propose the *Deterministic Uncertainty Estimation* (DUE) which in addition to DUQ method ensures that the encoder mapping is bi-Lipschitz. It would be interesting to explore if these methods perform well on similar OOD experiments as those considered in this work.

Other pitfalls of MCDropout have been pointed out in [Osband, 2016] where it is argued that suggest MCDropout approximates the risk, not to uncertainty.

In this work we focused on the most popular UE methods in machine learning, however, it is worth noting that other approaches exist. For example, the *jackknife+* method [Barber et al., 2020] is known to have good coverage guarantees, however it is less popular in deep learning where datasets are typically large, due to its requirement to retrain a model on multiple subsets of the full training set. We leave analyzing the OOD performance of other UE methods for future work.

## B Details on the Implementation

### B.1 Experiments on 2D-Toy data

In this section we list the details of the implementation of Fig. 1. The 2D toy dataset is generated sampling two normal distributions. For class 1, we sample 200 points from  $\mathcal{N}(\boldsymbol{\mu}_1, \sigma)$ , with  $\boldsymbol{\mu}_1 = (2, 2)$  and  $\sigma = \text{diag}((0, 1, 0.1))$ . For class 0, we sample 200 points from  $\mathcal{N}(\boldsymbol{\mu}_0, \sigma)$ , with  $\boldsymbol{\mu}_0 = (-2, -2)$ .

To train our Gaussian Process model, we use the `GaussianProcessClassifier` function<sup>1</sup> from the `sklearn` library Pedregosa et al. [2011], using an RBF kernel. The uncertainty is then calculated as the entropy of the predicted probability distribution, see § 2.

For our MCDropout model, we use an MLP with one hidden layer that contain 300 neurons. The activation function is ReLU and the dropout rate is 0.5. The uncertainty is calculated via average entropy of 100 forward dropout. The network is optimised by Adam with a cross-entropy loss using the PyTorch library [Paszke et al., 2019].

We tuned hyperparameters extensively. Using grid search, we tried different combinations of dropout rates and regularisation coefficients. (Additionally, we also explored increasing the number of neurons up to 10,000, increasing the number of hidden layers up to 4, changing the activation function, trying different values of dropout and regularization. However, none of these variations led to a change of our main observations.)

For BNNs, we use the approximation from “Bayes by Backprop” [Blundell et al., 2015], implemented through the `blitz` Python package [Esposito, 2020]. The network has two hidden layers, with 512 neurons for the first layer, and 128 neurons for the second layer. We use a regularization coefficient of 0.1.

### B.2 Experiments on MNIST

We select the 0 and 1 digits from the MNIST dataset [LeCun et al., 2010], and train the different models using solely these two classes. To train a GP on MNIST, we first train a Multi-layer Perceptron (MLP) of three layers—each of 600, 20 and 2 units, resp., using a dropout rate 0.6. We then use the first two layers as encoder  $\mathcal{E}: \mathbf{x} \mapsto \mathbb{R}^{20}$ ,  $\mathbf{x} \in \mathbb{R}^{28 \times 28}$ , which is kept fixed during the training of the GP model. Thus the GP model is trained given the embeddings of the input images, and the GP implementation is as in B.1.

For MCDropout, we used smaller network of two layers of 500 and 2 units each, and after sampling  $M = 100$  models we use the entropy of the output distribution to compute the uncertainty estimates.

<sup>1</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.gaussian\\_process.GaussianProcessClassifier.html](https://scikit-learn.org/stable/modules/generated/sklearn.gaussian_process.GaussianProcessClassifier.html)

Finally, the architecture used for the BNN uncertainty estimation method is an MLP with three hidden layers of 1024, 128 and 2 units each. We used a regularization coefficient of 0.1.

In Fig. 2 we designed the experiment as follows: to measure the uncertainty, we randomly select two data points of each class,  $x^0, x^1$ . We then test the models on artificial data generated by linear combination of the form  $t \cdot x^1 + (1 - t) \cdot x^0, t \in [-1, 2]$ . We also explore the result for  $t$  greater than 1 and less than 0 although it might not make sense in the real world. Then we calculate the uncertainty for different value of  $t \in [-1, 2]$ . We display the values obtained from repeating this procedure 100 times and averaging the values.

For the implementation details, the package and the function are the same as previous section. And all these methods achieve accuracy more than 99.9% on the test dataset.

The Gaussian Process is Deep GP. We at first train a MLP, with two hidden layers and dropout rate 0.6. The first hidden layer has 600 neurons, while the second one only has 20. This stage is feature extraction. For each data, the input space is mapped from  $28 \times 28$  to 20. After this network is well-trained, we use the value of these 20 new features as inputs to train Gaussian Process.

The implementation of MCDropout is more direct. The network has 500 neurons and dropout rate 0.6. The uncertainty is calculated via average entropy of 100 forward dropout. Again, we tried different structure of the network and it doesn't improve at all.

For BNNs. The networks used have two hidden layers, with 1024 neurons in the first one and 128 neurons in the second one. Regularization coefficient is 0.1. The BNN-HMC implementation is done using the `hamiltorch`<sup>2</sup> library, with a prior precision of 5 for each parameter of the model, a trajectory length of 3, and a step size of 0.0005.

### B.3 Experiments on CIFAR-10 [Krizhevsky, 2012]

We use a ResNet-18 architecture [He et al., 2016b] modified to accommodate the MCDropout sampling procedure. The modification consists of adding a dropout layer with dropout probability  $p = 0.2$  after each convolutional layer. We train with an SGD optimizer with a momentum of 0.9 and a weight decay of 0.00001. The learning rate is following a triangular scheduler, first increasing linearly from 0 to 0.2 during the first half of the epochs, and then decreasing linearly back to 0. We train for a total of 30 epochs, reaching a test accuracy of 89%. When using MCDropout, we sample 20 models and average the output distributions. The images of Fig. 3c are taken randomly from the CIFAR100 test images with an MCDropout predicted probability larger than 0.9.

### B.4 How to use the provided source code for OOD evaluation

We provide some code to help reproduce this experimental protocol, see: <https://github.com/epfml/uncertainty-estimation>. The code in the notebook `0-1-interpolation-for-UE-evaluation.ipynb` shows how to build the test digits and how to compute the uncertainty for a simple classifier.

## C Additional Numerical Results

Class (digit)	Uncertainty	Class (digit)	Uncertainty
0	<b>0.0370</b>	5	0.4963
1	<b>0.0200</b>	6	0.4861
2	0.3304	7	0.4949
3	0.4344	8	0.4775
4	0.4689	9	0.4832

Table 1: The uncertainty prediction of MCDropout for the digit 0 to digit 9. With training set only 0 and 1, and measured on all of the test set. The number is the average uncertainty for all of these numbers. We saw that number 2-9 have high uncertainty, while the uncertainty for number 0 and 1 are small.

<sup>2</sup><https://github.com/AdamCobb/hamiltorch>



## D Proof of Theorem 1

Our notation follows that of [Rasmussen and Williams, 2005, pages 41–43], and it is summarized in Table 2.

$k(\mathbf{x}, \mathbf{x}')$	Kernel value between $\mathbf{x}$ and $\mathbf{x}'$
$X$	Training input data
$\mathbf{y}$	Ground-truth labels
$\mathbf{x}_*$	Test data point
$K$	covariance matrix for the (noise free) $\mathbf{f}$ values. i.e $k(X, X)$
$\mathbf{k}_*$	Kernel matrix between test and training input. i.e $k(\mathbf{x}_*, X)$
$\bar{\pi}_*$	The predict probability of new test sample $\mathbf{x}_*$
$\mathbf{f}$	Latent function values
$\hat{\mathbf{f}}$	Calculated by $\mathbb{E}_q[\mathbf{f}   X, \mathbf{y}]$ , the maximum posterior.
$f_*$	Gaussian process (posterior) prediction (random variable)
$W$	The negative Hessian of the $\log p(\mathbf{y}   \hat{\mathbf{f}})$ .

Table 2: Summary of the notation.

### D.1 Assumption

We study a Gaussian Process  $f$  with  $n$ -D training data  $X$  with kernel  $K$ . We assume the binary classification setting (Gaussian Process Classification with Laplace approximation [Williams and Barber, 1998]).

### D.2 Predictive probability

The predict probability of new test sample  $\mathbf{x}_*$  is approximated via

$$\bar{\pi}_* \simeq \int \sigma(f_*) q(f_* | X, \mathbf{y}, \mathbf{x}_*) df_*, \quad (3)$$

where  $\sigma(z)$  is a link function (logistic or Gaussian CDF),  $q(f_* | X, \mathbf{y}, \mathbf{x}_*)$  is a Gaussian distribution whose mean and variance are given below

### D.3 Parameters for Gaussian distribution above

The parameters for  $q(f_* | X, \mathbf{y}, \mathbf{x}_*)$  are

Mean:

$$\mathbb{E}_q[f_* | X, \mathbf{y}, \mathbf{x}_*] = \mathbf{k}_*^\top K^{-1} \hat{\mathbf{f}} = \mathbf{k}_*^\top \nabla \log p(\mathbf{y} | \hat{\mathbf{f}}).$$

Variance:

$$\begin{aligned} \mathbb{V}_q[f_* | X, \mathbf{y}, \mathbf{x}_*] &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top K^{-1} \mathbf{k}_* + \mathbf{k}_*^\top K^{-1} (K^{-1} + W)^{-1} K^{-1} \mathbf{k}_* \\ &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (K + W^{-1})^{-1} \mathbf{k}_*, \end{aligned}$$

where  $\hat{\mathbf{f}} = \mathbb{E}_q[\mathbf{f} | X, \mathbf{y}]$ .

### D.4 Theorem

**Theorem 1.** Consider the  $n$ -dimensional binary GP classification model described above with kernel  $K(\mathbf{x}, \mathbf{x}')$ . Let  $\mathbf{k}_*$  denote the similarity vector  $(\mathbf{k}_*)_i = K(\mathbf{x}_i, \mathbf{x}_*)$  between a test sample  $\mathbf{x}_*$  and training data  $\mathbf{x}_i$ ,  $i = 1, \dots, n$ . If  $\|\mathbf{k}_*\| \leq \varepsilon$  is small then Gaussian Process Classification with Laplace approximation will assign a probability close to 0.5 for both of two classes, with probability approaching  $\frac{1}{2}$  as  $\varepsilon \rightarrow 0$ .

*Proof.* As can be seen from the derivations above, if  $\|\mathbf{k}_*\| \leq \varepsilon$  is small, then the mean tends to zero as  $\varepsilon \rightarrow 0$ , and variance tends to  $k(\mathbf{x}_*, \mathbf{x}_*)$ . We now calculate the predictive probability with this mean and variance.

$$\begin{aligned}
\bar{\pi}_* &:= \int \sigma(z) \mathcal{N}(z \mid \bar{f}_*, \mathbb{V}[f_*]) dz \\
&= \int \sigma(z) \mathcal{N}(z \mid 0, k(\mathbf{x}_*, \mathbf{x}_*)) dz \\
&= \int (\sigma(z) - \frac{1}{2}) \mathcal{N}(z \mid 0, k(\mathbf{x}_*, \mathbf{x}_*)) dz + \frac{1}{2} \int \mathcal{N}(z \mid 0, k(\mathbf{x}_*, \mathbf{x}_*)) dz \\
&\stackrel{\varepsilon \rightarrow 0}{=} \frac{1}{2}.
\end{aligned} \tag{4}$$

In the third equation, the first term vanishes because it is an odd function. Hence, we showed that when  $\|\mathbf{k}_*\| \leq \varepsilon$  is small, the prediction tends to  $\frac{1}{2}$  for  $\varepsilon \rightarrow 0$ .  $\square$

As we know, the entropy is the largest when two classes both have probability 0.5 in binary classification. This means that if the testing input is far away from training input, then GP would return a high uncertainty for this input.