



Original software publication

## Variational Neural Networks implementation in Pytorch and JAX

 Illia Oleksiienko <sup>a,\*</sup>, Dat Thanh Tran <sup>b</sup>, Alexandros Iosifidis <sup>a</sup>
<sup>a</sup> Department of Electrical and Computer Engineering, Aarhus University, Denmark

<sup>b</sup> Department of Computing Sciences, Tampere University, Finland


### ARTICLE INFO

#### Keywords:

 Bayesian Neural Networks  
 Bayesian deep learning  
 Uncertainty estimation  
 PyTorch  
 JAX

### ABSTRACT

Bayesian Neural Networks consider a distribution over the network's weights, which provides a tool to estimate the uncertainty of a neural network by sampling different models for each input. Variational Neural Networks (VNNs) consider a probability distribution over each layer's outputs and generate parameters for it with the corresponding sub-layers. We provide two Python implementations of VNNs with PyTorch and JAX machine learning libraries that ensure reproducibility of the experimental results and allow implementing uncertainty estimation methods easily in other projects.

### Code metadata

 Current code version  
 Permanent link to code/repository used for this code version  
 Permanent link to Reproducible Capsule  
 Legal Code License  
 Code versioning system used  
 Software code languages, tools, and services used  
 Compilation requirements, operating environments & dependencies

 v1.0.0  
<https://github.com/SoftwareImpacts/SIMPAC-2022-198>  
<https://codeocean.com/capsule/2963476/tree/v1>  
 Apache-2.0 license  
 Git  
 Python, R,  $\LaTeX$   
 Pytorch/JAX (Tensorflow), plotnine, pandas, numpy, neural-tangents, scipy, scikit-learn, dm-haiku, dm-acme, chex, optax, fire, tqdm, tensorboard, tensorboardX, torchvision, absl-py, tensorflow

 If available Link to developer documentation/manual  
 Support email for questions

[io@ece.au.dk](mailto:io@ece.au.dk)

## 1. Introduction

Uncertainty estimation in neural networks provides an ability to detect possible errors in predictions that can be caused by the lack of training data or out-of-distribution samples. These failures may occur silently in regular neural networks and cause danger in critical tasks, such as medical image analysis or autonomous driving. To estimate the network's uncertainty in its predictions, several approaches have been proposed which are mostly following the Bayesian Neural Network (BNN) framework [1–3]. BNNs introduce a probability distribution over the neural network's weights. For each input, they sample a set of different models from the weights' distribution which describe it from different points of view. The deviation in the prediction of these different models is an estimate of the uncertainty of the network for the specific input.

BNNs usually come with increased computational cost and memory requirements compared to neural networks providing point estimations for their inputs. The choice of the weights' probability distribution results in different computational costs, but it also affects the statistical quality of the network. This is due to that more complex and parameterized distributions allow for greater freedom, but require more computations to train and execute the network, while simplified distributions may not represent the problem at hand. Existing methods explore a variety of distributions, including the Gaussian [4,5], Bernoulli [6], and Categorical [7] distributions.

Variational Neural Networks (VNNs) [8] introduce a new type of uncertainty estimation for neural networks by considering a distribution over each layer's outputs and generate the distribution's parameters by processing inputs with corresponding sub-layers. To keep a low

The code (and data) in this article has been certified as Reproducible by Code Ocean: (<https://codeocean.com/>). More information on the Reproducibility Badge Initiative is available at <https://www.elsevier.com/physical-sciences-and-engineering/computer-science/journals>.

\* Corresponding author.

E-mail addresses: [io@ece.au.dk](mailto:io@ece.au.dk) (I. Oleksiienko), [thanh.tran@tuni.fi](mailto:thanh.tran@tuni.fi) (D.T. Tran), [ai@ece.au.dk](mailto:ai@ece.au.dk) (A. Iosifidis).

<https://doi.org/10.1016/j.simpa.2022.100431>

Received 13 September 2022; Accepted 11 October 2022

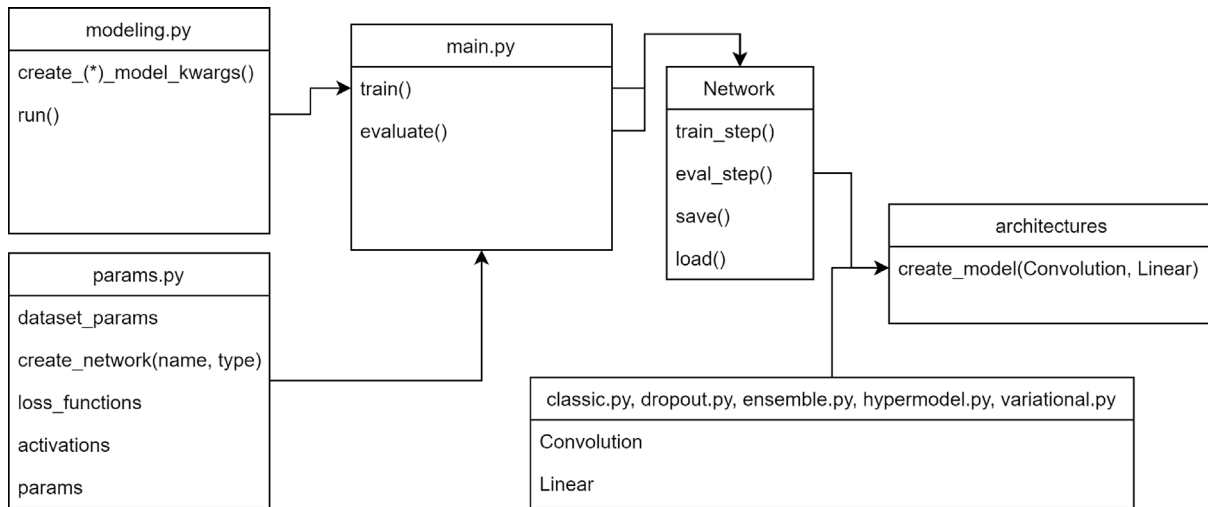


Fig. 1. Structure of the PyTorch implementation of VNNs (R1).

computational cost and memory requirements of VNNs, we consider the Gaussian distribution with its parameters, i.e., mean and variance, generated by the corresponding sub-layers. As shown in [8], VNNs achieve better uncertainty quality compared to Monte Carlo Dropout (MCD) [6] and Bayes By Backprop (BBB) [4], while having similar properties to those methods from the Bayesian Model Averaging perspective.

## 2. Impact overview

The codebase of VNNs is split into two implementations with different purposes. The PyTorch implementation (R1) provides a general-purpose implementation of VNNs using Python 3 [9] and a machine learning framework PyTorch [10] together with code for image classification experiments on MNIST [11] and CIFAR-10 [12] datasets. The JAX implementation (R2) is built upon JAX [13] implementation [14] of Epistemic Neural Networks [15] and provides a JAX version of VNNs together with the uncertainty quality estimation experiments.

The software provides full reproducibility of the research results in [8] and includes code for generating plots and tables therein, which means that validating results and creating new experiments based on VNNs should be an easy task. The structure of the software is modular and allows for fast implementation of additional methods or usage of already existing methods in an existing project by replacing layers of a regular neural network with those provided by the software.

## 3. Functionalities and key features

*The PyTorch implementation (R1).* of VNNs is programmed in a modular manner, where each of the tested uncertainty estimation methods, including VNNs, Bayes By Backpropagation (BBB) [4], Monte Carlo Dropout (MCD) [6], Deep Ensembles [7] and Hypermodels [5], is implemented using a common training interface and isolated from the actual network architecture. The file structure of R1 is represented in Fig. 1 and consists of

- *main.py* is an entry point that implements *train* and *evaluate* functions to train and test models of interest.
- *network.py* implements a *Network* class that encapsulates all the functions need to train, run, load and save the model.
- *architectures* folder contains a set of scripts that implement model architecture classes, such as ResNet, DenseNet, and simpler architectures. These classes derive from the *Network* class and have convolutional and linear layers as generic parameters allowing the creation of models with different uncertainty methods for the same network architecture.

- *classic.py*, *dropout.py*, *ensemble.py*, *hypermodel.py*, *variational.py* implement the corresponding uncertainty methods by creating convolutional and linear layers of the method of interest or a full network generator. By providing these implementations into an architecture creator, a network for the selected uncertainty method and architecture will be generated.
- *params.py* contains data generators for MNIST and CIFAR-10 datasets, string-indexed dictionaries of all network creators, loss functions, activation functions and optimizers.
- *modeling.py* is an alternative entry point and is used to generate sets of experiments for each of the uncertainty methods and run them in parallel.
- *tools/create-grouped-tex-report.py* collects results from the trained models and generates a  $\LaTeX$  table with two best entries for each network type.

The presented repository structure allows for an easy expansion of the project:

- Adding a new network type requires implementing either convolutional and linear layers (as shown in *classic.py*, *dropout.py*, *variational.py*), or a network creator that takes an architecture as an input (as shown in *ensemble.py*, *hypermodel.py*).
- Adding a new architecture requires implementing it with convolutional and linear layer classes given as parameters.
- Adding a new layer type requires implementing versions of it for each of the network types, which is usually straightforward, following the already implemented ones.
- Using these methods for other projects can be achieved by copying layer implementations to the existing project with classical neural networks and replacing regular layers in the network with the desired ones.

*The JAX implementation (R2).* is based on the Epistemic Neural Networks repository [14] and provides a code to reproduce uncertainty quality experiments. All main files of R2 are located at *enn/experiments/neurips\_2021* with the following file structure:

- *run\_testbed.py* is an entry point that generates a problem for the selected experiment parameters and trains a set of models for the selected method.
- *agent\_factories.py* contains model creators for each of the uncertainty methods and functions to create networks with different parameters to be trained.
- *enn/networks* folder contains implementations of the different uncertainty methods.

- `tools/create_enn_plots.py` is used to generate plots for each method, groups of experiments and a summary.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgments

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871449 (OpenDR). This publication reflects the authors’ views only. The European Commission is not responsible for any use that may be made of the information it contains.

### References

- [1] D.J.C. Mackay, Probable networks and plausible predictions — a review of practical Bayesian methods for supervised neural networks, *Network: Comput. Neural Syst.* 6 (3) (1995) 469–505.
- [2] A.G. Wilson, P. Izmailov, Bayesian deep learning and a probabilistic perspective of generalization, 2020, [arXiv:2002.08791](https://arxiv.org/abs/2002.08791).
- [3] T. Charnock, L. Perreault-Levasseur, F. Lanusse, Bayesian neural networks, 2020, [arXiv:2006.01490](https://arxiv.org/abs/2006.01490).
- [4] C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural networks, 2015, [arXiv:1505.05424](https://arxiv.org/abs/1505.05424).
- [5] V. Dwaracherla, X. Lu, M. Ibrahimi, I. Osband, Z. Wen, B.V. Roy, Hypermodels for exploration, 2020, [arXiv:2006.07464](https://arxiv.org/abs/2006.07464).
- [6] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, 2016, [arXiv:1506.02142](https://arxiv.org/abs/1506.02142).
- [7] I. Osband, J. Aslanides, A. Cassirer, Randomized prior functions for deep reinforcement learning, 2018, [arXiv:1806.03335](https://arxiv.org/abs/1806.03335).
- [8] I. Oleksiienko, D.T. Tran, A. Iosifidis, Variational neural networks, 2022, [arXiv:2207.01524](https://arxiv.org/abs/2207.01524).
- [9] G. Van Rossum, F.L. Drake, *Python 3 Reference Manual*, CreateSpace, 2009.
- [10] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimselshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in: *Advances in Neural Information Processing Systems*, Vol. 32, Curran Associates, Inc., 2019, pp. 8024–8035.
- [11] L. Deng, The mnist database of handwritten digit images for machine learning research, *IEEE Signal Process. Mag.* 29 (6) (2012) 141–142.
- [12] A. Krizhevsky, Learning multiple layers of features from tiny images, Tech. Rep, 2009.
- [13] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of Python+NumPy programs, 2018, URL [http://github.com/google/jax](https://github.com/google/jax).
- [14] I. Osband, Z. Wen, M. Asghari, M. Ibrahimi, X. Lu, B.V. Roy, Epistemic neural networks implementation, 2021, URL <https://github.com/deepmind/enn>.
- [15] I. Osband, Z. Wen, M. Asghari, M. Ibrahimi, X. Lu, B.V. Roy, Epistemic neural networks, 2021, [arXiv:2107.08924](https://arxiv.org/abs/2107.08924).