Overview over uncertainty estimation in neural networks

Maximilian Pintz | Fraunhofer IAIS | 19.10.2022

Deep Learning

Deep neural networks

- Versatile systems for modeling complex patterns in data
- Learn from given data examples (training data) with the ability to generalize to new (unseen) data
- Successful across a large bandwidth of applications

Goal: Extract patterns from data records

[1] Input (e.g. street scene recording)

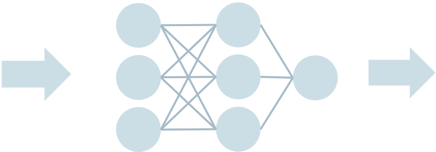

Deep neural network

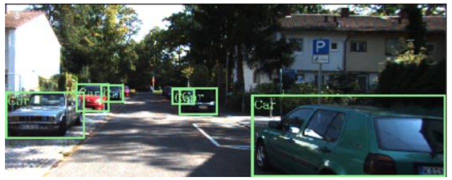

Output pattern (e.g. position of cars)

[2] 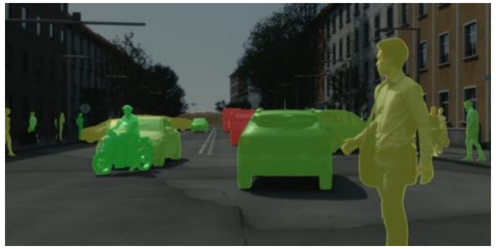

[3] 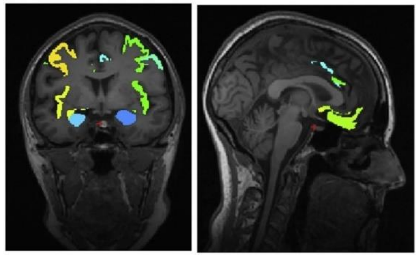

[4] 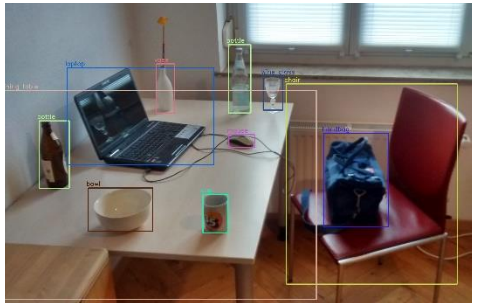

References: [1] Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2013). Vision meets robotics: The kitti dataset. The International Journal of Robotics Research, 32(11), 1231-1237. Image taken from the 2d object detection dataset. [2] Magnus Wrenninge and Jonas Unger. “Synscapes: A Photorealistic Synthetic Dataset for Street Scene Parsing”. arXiv:1810.08705. Image source: „Synscapes Examples“, 7DLabs, URL: https://7dlabs.com/synscapes-examples [3] „Objects detected with OpenCV's Deep Neural Network module (dnn) by using a YOLOv3 model trained on COCO dataset capable to detect objects of 80 common classes.” by MThelier, CC BY-SA 4.0 [3] Alexander Selvikvag Lundervold and Arvid Lundervold. “An overview of deep learning in medical imaging focusing on MRI”. In: Zeitschrift für Medizinische Physik 29.2 (2019), pp.102–127.

Motivation

How reliable are deep neural networks?

Deep neural networks are increasingly employed in safety-critical systems (e.g. autonomous vehicles, medical diagnosis)

Failure of these systems potentially has severe consequences

⇒ Risk mitigation and an assessment of their reliability are needed

Deep neural networks are impacted by **uncertainties**

- Uncertainty: limited knowledge about the task or random factors that impact model performance
- Assessing and handling uncertainties is crucial for building reliable systems

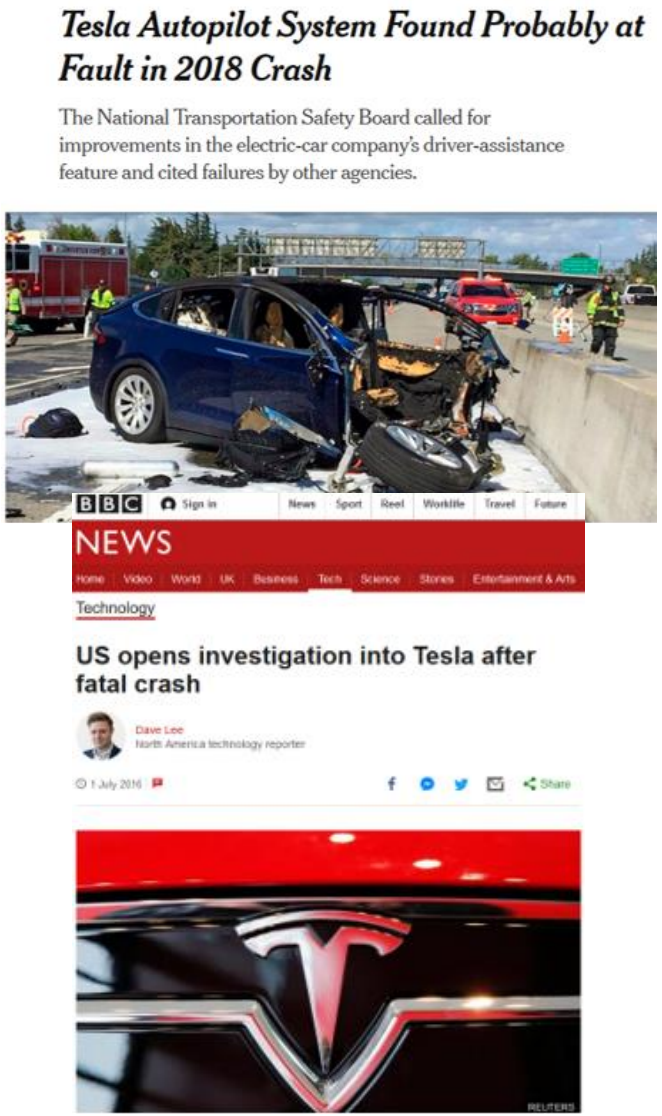

Tesla Autopilot System Found Probably at Fault in 2018 Crash

The National Transportation Safety Board called for improvements in the electric-car company's driver-assistance feature and cited failures by other agencies.

BBC NEWS

Technology

US opens investigation into Tesla after fatal crash

Dave Lee
North America technology reporter

Sources of uncertainty
Lack of training data

[1]

Uncertainty: Have we supplied our model with enough training data?

Impact: Lack of training data can worsen model performance, especially performance on uncommon inputs

References: [1] Braun et al. „EuroCity Persons: A Novel Benchmark for Person Detection in Traffic Scenes.“ In IEEE Transactions on Pattern Analysis and Machine Intelligence, 2019

Sources of uncertainty
Data uncertainty

Uncertainty: Is the data the model learns from clear without ambiguity?

Impact: Uncertainty in the input can propagate to the output

Examples:
- Coin toss
- Sensor noise
- Weather conditions, occluded sensors, …
- Annotation error / ambiguity



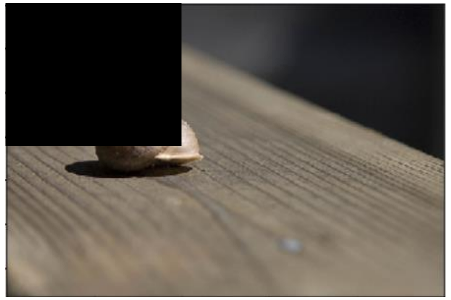


[1]

References: [1] Geiger et al. (2013). Vision meets robotics: The kitti dataset. The International Journal of Robotics Research, 32(11), 1231-1237. Image taken and modified from the 2d object detection dataset.

Uncertainty Estimation

Observations:

- Deep neural networks are impacted by several uncertainties, which cause changes in the end result and impact task performance
- Many of these uncertainties are inherent to the process and cannot (practically) be reduced, especially in open-world scenarios

Uncertainty estimation: Quantifying the impact of uncertainties on the end result

- Tries to answer: „How certain is the model of the prediction?“
- Understand better whether model outputs can be trusted

Uncertainty Estimation in Steering Angle Prediction [1]

References: [1] Loquercio, A., Segu, M., & Scaramuzza, D. (2020). A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, *5*(2), 3153-3160.

Applications of uncertainty estimation

Out-of-distribution detection

Identifying „unusual“ inputs, in particular inputs that do not stem from the training data distribution

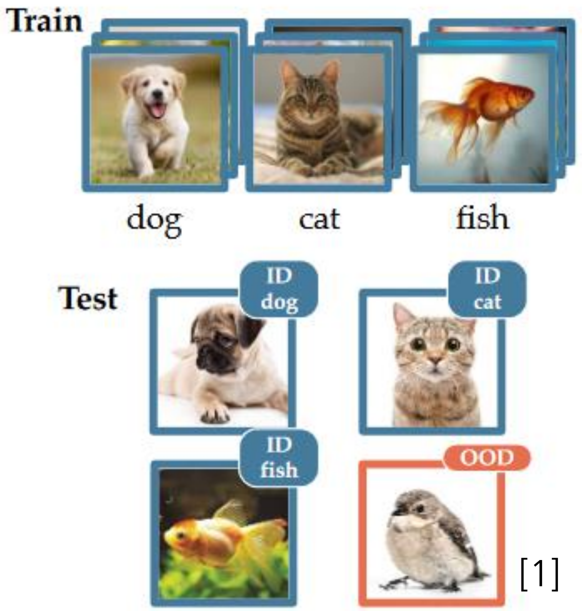

[1]

Active Learning

Selecting uncertain prediction for manual labeling

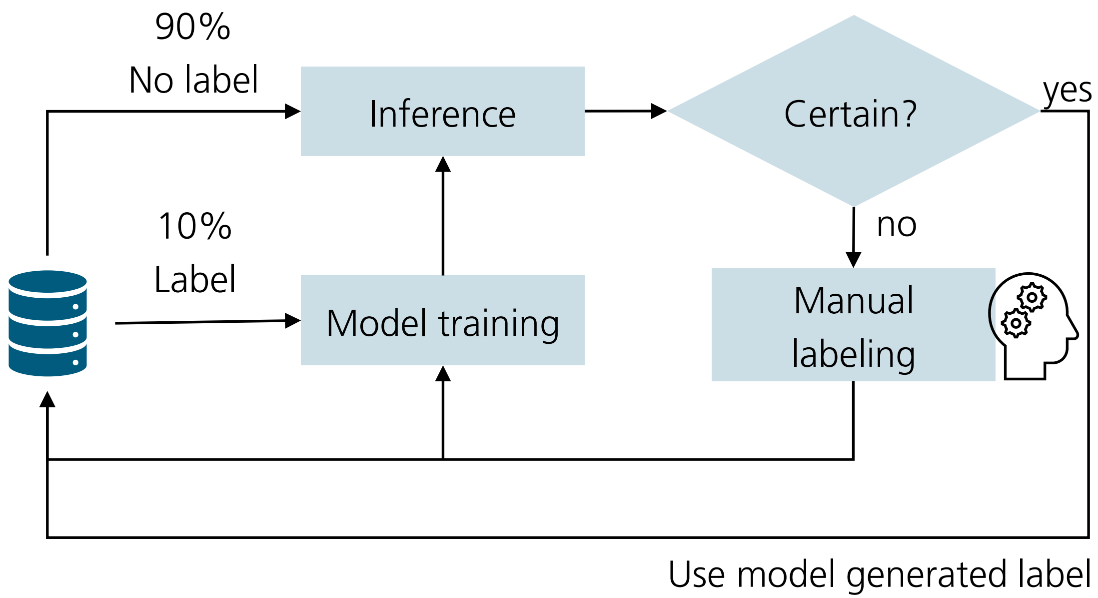

References: [1] Yang et al. „Generalized Out-of-Distribution Detection: A Survey“, arxiv:2110.11334

Applications of uncertainty estimation

Computer Vision

Natural Language Processing


„Bank"

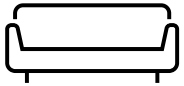
„Bank"

„I am quite sure that A equals B."

Further applications

Reinforcement learning, Robotics, Meta-learning, Few-shot-learning …

Uncertainty estimation for neural networks

Probabilistic Deep Learning

In probabilistic deep learning, the neural network computes a probability distribution instead of a single output

Estimation of the probability distribution $p(y \mid x)$ over model outputs that captures (predictive) uncertainty

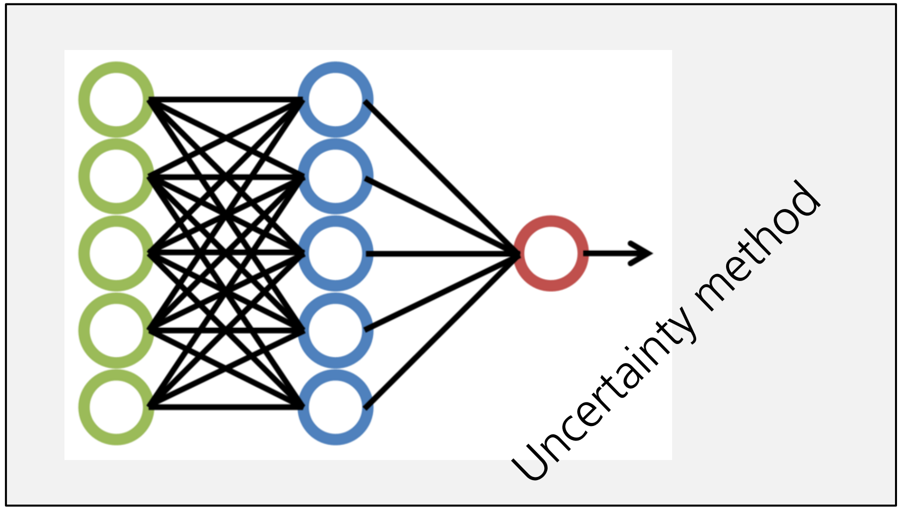

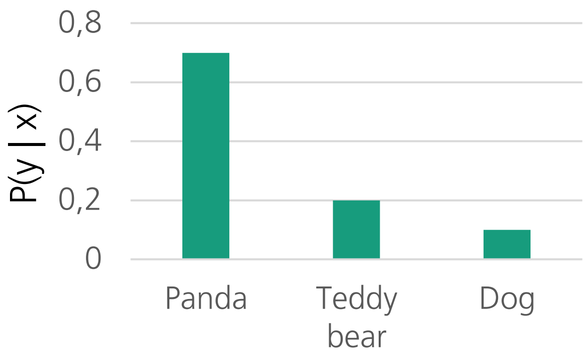

Fraunhofer IAIS

Types of uncertainty
Aleatoric and epistemic

Aleatoric uncertainty

Inherent (non-reducible) randomness in the learning process [1], in particular the data

Examples:

- Annotation error / ambiguity
- Sensor noise
- Corruptions

Typically modeled via a probability distribution on model outputs p(y | x)

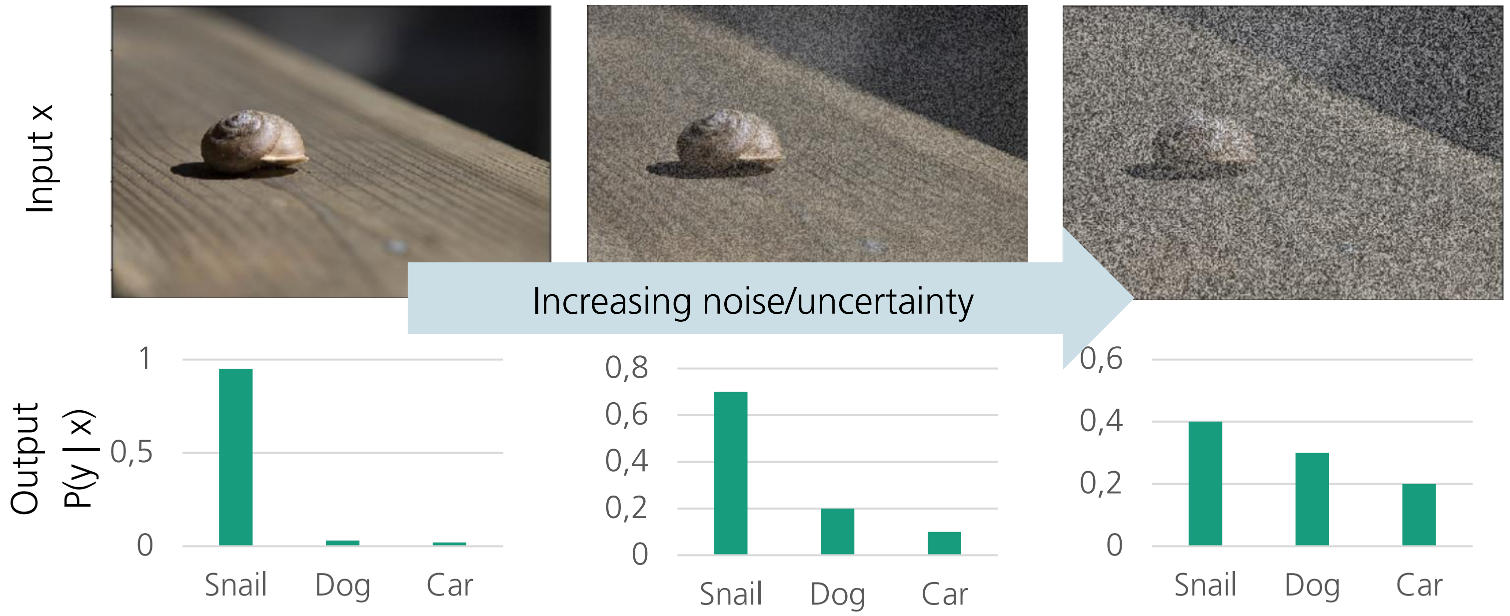

References: [1] Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, *110*(3), 457-506. [2] Kohl. et al. „A Probabilistic U-Net for Segmentation of Ambiguous Images“, NIPS 2018

Types of uncertainty
Aleatoric and epistemic

Epistemic uncertainty

Lack of knowledge about the perfect model („Bayes-predictor")

Examples:

- Model mismatch
- Lack of training data
- Hyperparameter tuning
- Data shuffling, augmentation, splits

Typically modeled via distributions over models (most commonly model parameters p(w))

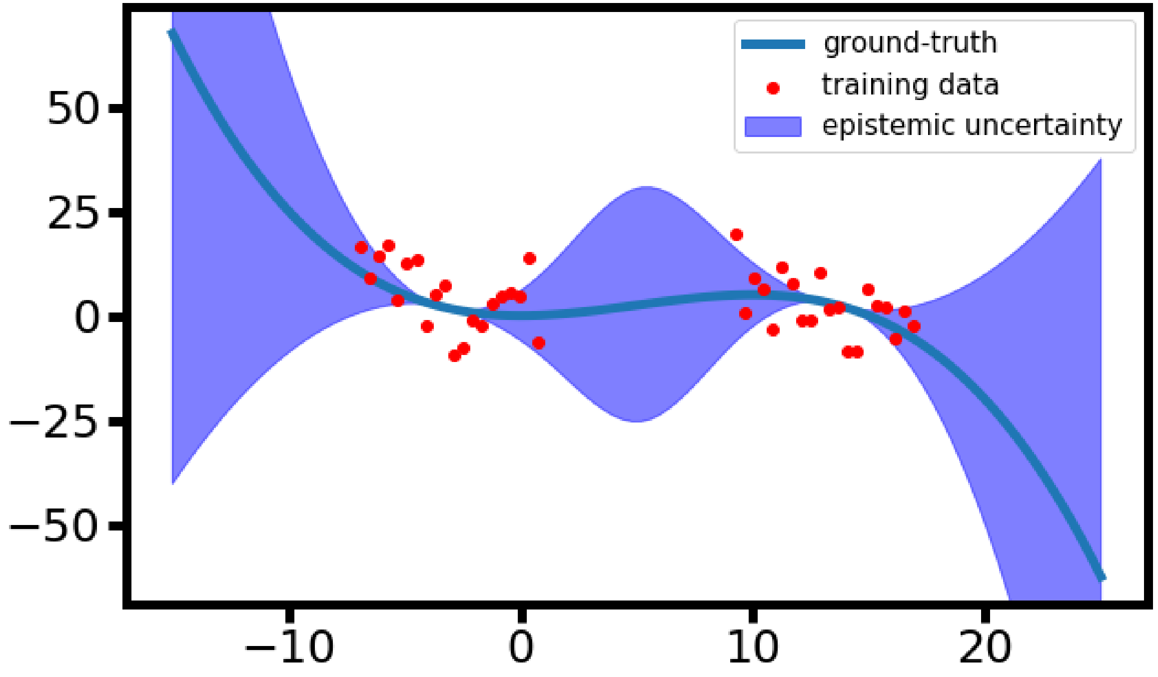

Evaluation of uncertainty estimates

Desiderata for „good“ uncertainty estimates

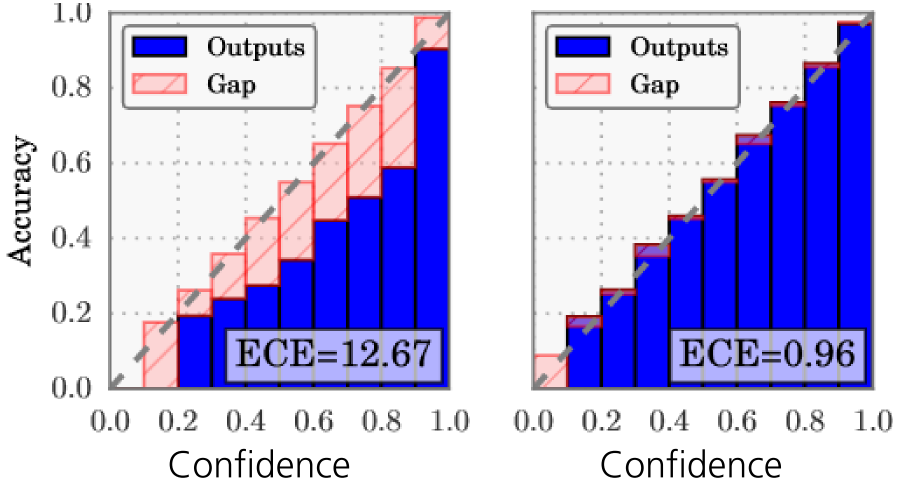

[2]

Desired properties of „good“ uncertainty estimates

Alignment of predicted probability distributions with given data

- Alignment per data point (Proper Scoring Rules)
- Alignment over data subsets (Calibration)

Application-specific properties

- Distinguishing between in- and out-of-distribution data points, between true & false positives, … (Entropy, AUC, …)
- Performance in active learning tasks, …

Qualitative properties

- Increasing uncertainty under data-shift, …

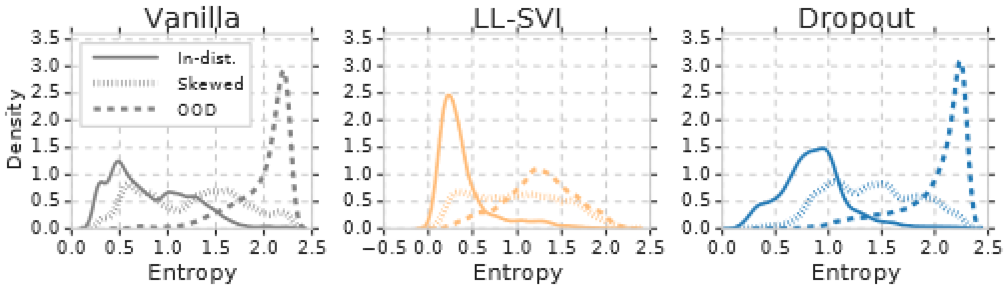

[1]

References: [1] Ovadia et al. (2019). Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, *32*. [2] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017, July). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321-1330). PMLR.

Uncertainty estimation methods for neural networks

Probabilistic Deep Learning

1. Direct parametrization
2. Bayesian Inference
3. Ensembling
4. Recalibration
5. (Test-time augmentation)

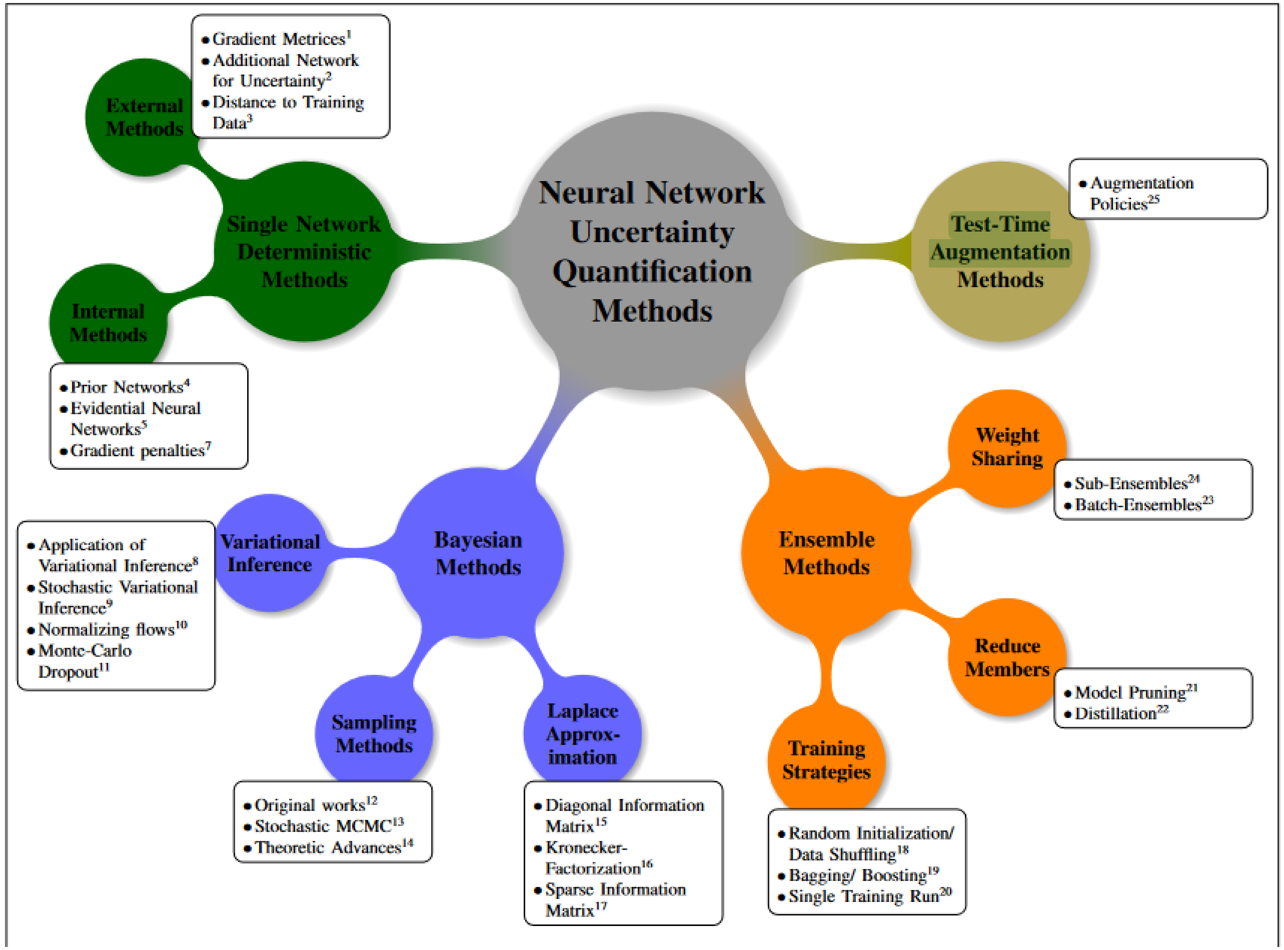

References: [1] Gawlikowski et al. (2021). A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*.

Uncertainty estimation methods for neural networks

Direct parametrization

Concept:

- Neural network computes parameters of a probability distribution in one forward pass (no sampling)

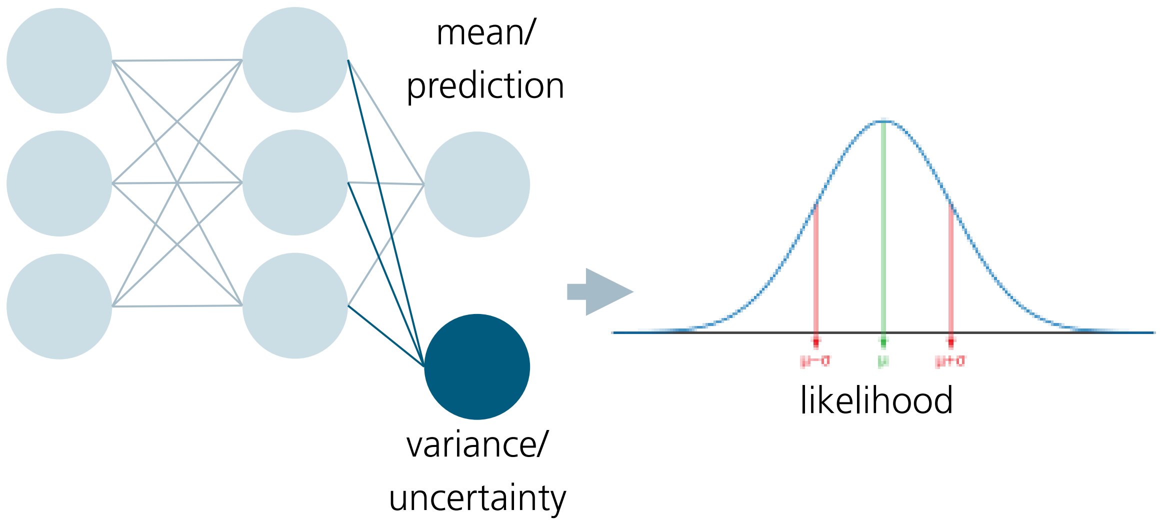

Simple variant: „Likelihood parametrization“

- Adding additional neurons to the output layer for computing distribution parameters (e.g. mean and variance)
- Model outputs the „likelihood“ (a distribution over model outputs), which describes aleatoric uncertainty
- Training: Maximize the likelihood over training data

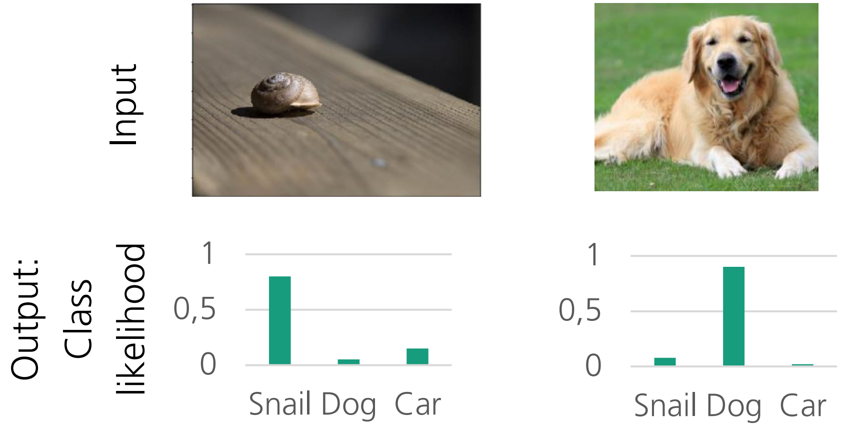

Uncertainty estimation methods for neural networks

Direct parametrization

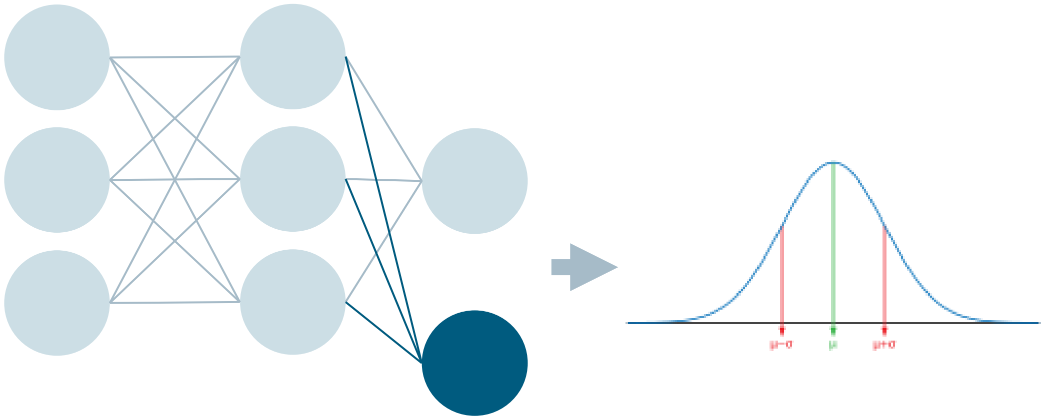

Likelihood parametrization

Pro:

- Low compute overhead at training and inference
- Uncertainty estimates with a single forward pass, no sampling required
- Easy to combine with other uncertainty methods
- Typically easy to implement (depends on the complexity of the distributions)

Cons:

- Only models aleatoric uncertainty (in particular does not capture out-of-distribution uncertainty)
- Requires change of architecture/loss (although typically just the output layer)
- Parametric assumption

Uncertainty estimation methods for neural networks

Direct parametrization

Current methods

Joint estimation of aleatoric and epistemic uncertainty with a single forward pass …
… but increased complexity in terms of implementation and compute overhead

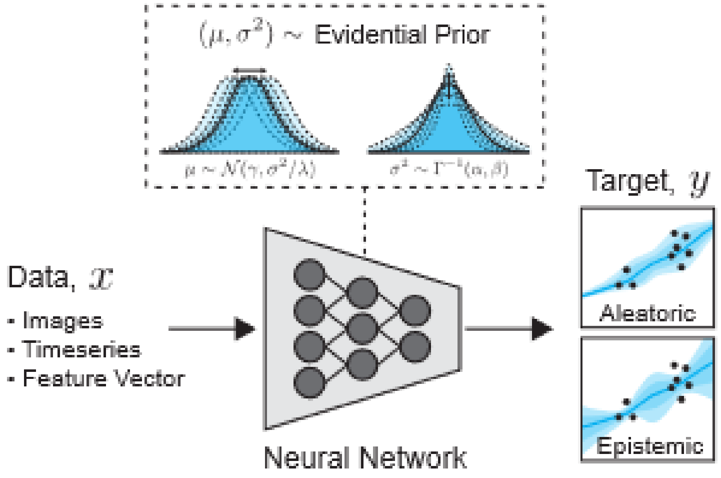

Evidential deep learning [1]

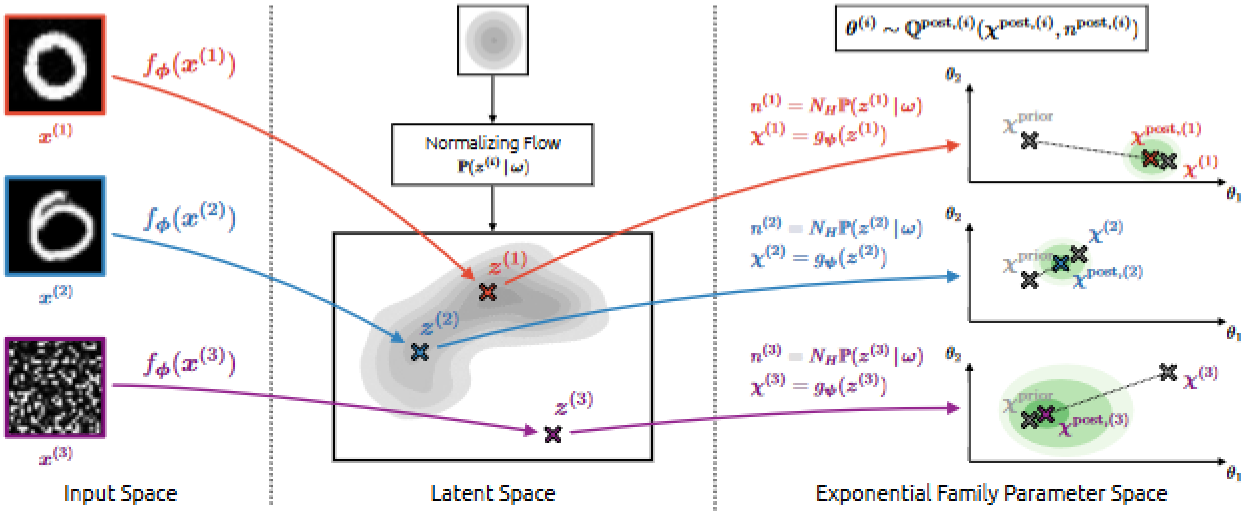

PostNets [2]

References: [1] Amini et al. (2020). Deep evidential regression. Advances in Neural Information Processing Systems, 33, 14927-14937. [2] Charpentier et al. (2021, September). Natural Posterior Network: Deep Bayesian Predictive Uncertainty for Exponential Family Distributions. In *International Conference on Learning Representations*.

Uncertainty estimation methods for neural networks

Bayesian Inference for Deep Learning

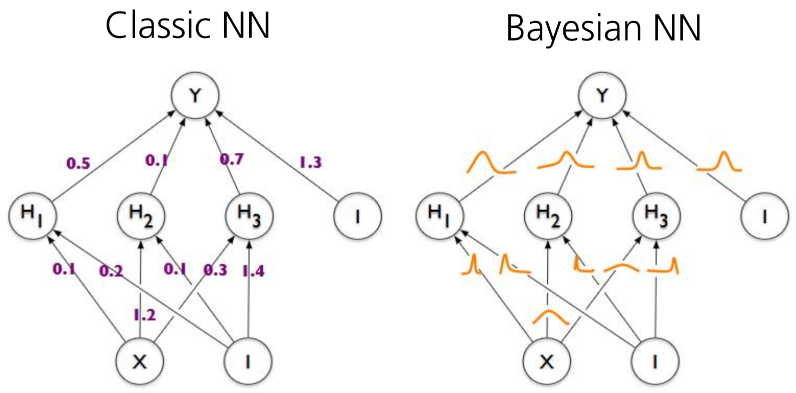

Bayesian inference

Concept:

- Theoretical framework for incorporating epistemic uncertainty
- Extension of likelihood parametrization
- Derives two central distributions
  - Posterior distribution $p(w \mid D)$ for epistemic uncertainty: distribution on model parameters w given training data D
  - Predictive posterior: $p(y \mid x, D)$ to jointly capture epistemic and aleatoric uncertainty: distribution on model outputs y given an unseen input x, and training data D

Thomas Bayes

For neural networks, it is infeasible to compute the exact posterior and predictive posterior

=> Approximate methods are needed (Discussed in the field „approximate Bayesian inference“)

 19.10.2022

Uncertainty estimation methods for neural networks
MC Dropout

Dropout mechanism [1]

Concept:

- Commonly used method for mitigating „overfitting“, the problem that a neural network learns the training data by heart instead of the underlying patterns
- Randomly deactivate neurons during training
- Training an ensemble of multiple, simpler sub-networks, instead of a complex single network
- Leave dropout inactive during inference

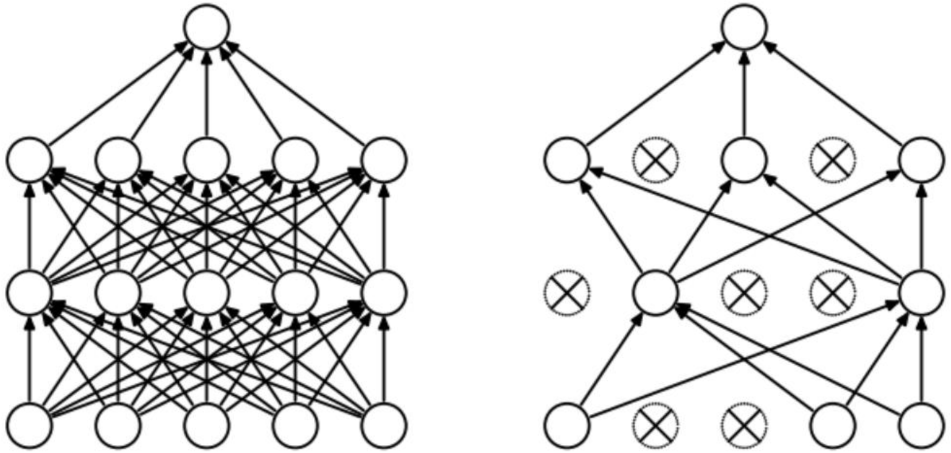

(a) Standard Neural Net (b) After applying dropout.

References: [1] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15, 1 (January 2014), 1929–1958.

Uncertainty estimation methods for neural networks

MC Dropout

MC Dropout [1]

Concept:

- Leave dropout activated during inference
- Collect multiple predictions from different sub-networks for fixed input
- Compute sample statistics (mean, variance) for uncertainty estimation
- Approximation of Bayesian inference

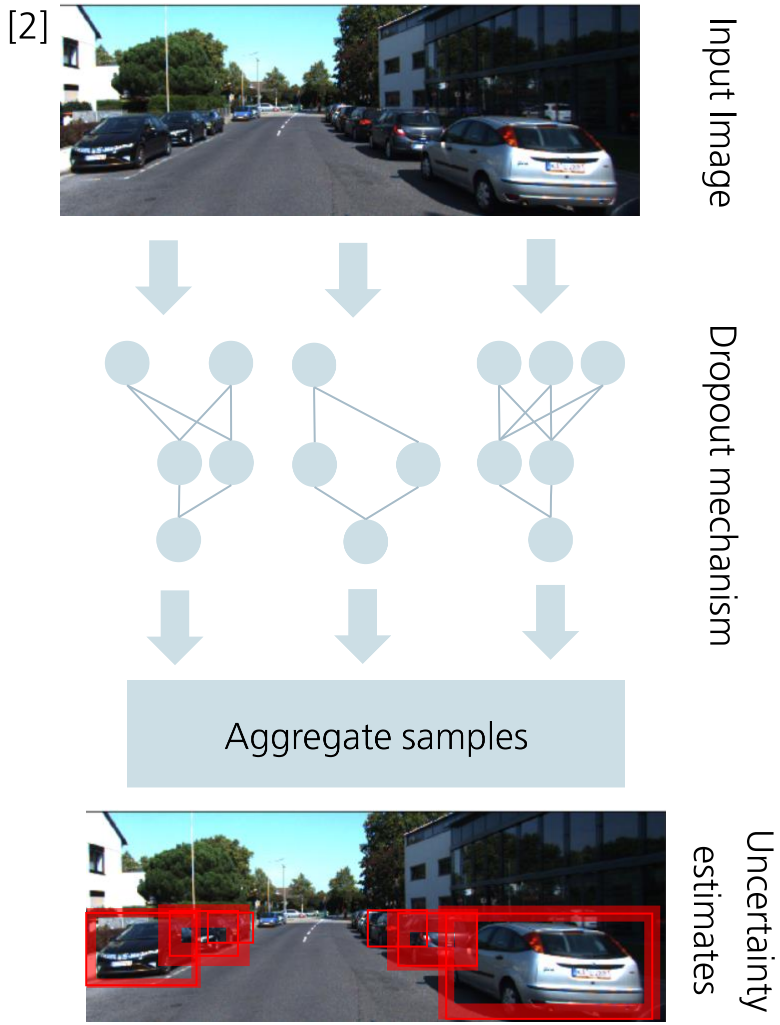

References: [1] Gal, Y., & Ghahramani, Z. (2016, June). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning* (pp. 1050-1059). PMLR. [2] Geiger et al. (2013). Vision meets robotics: The kitti dataset. The International Journal of Robotics Research, 32(11), 1231-1237. Image taken and modified from the 2d object detection dataset.

Uncertainty estimation methods for neural networks

MC Dropout

Pro:

- Widespread method
- Applicable to a broad range of architectures
- Easy to implement
- Low computational overhead during training
- Captures epistemic uncertainty

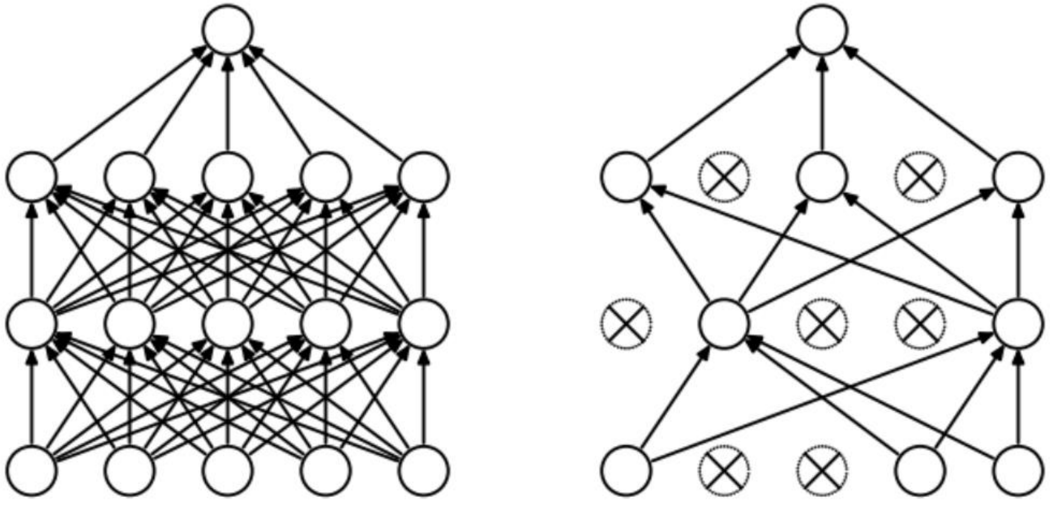
(a) Standard Neural Net (b) After applying dropout. [1]

Con:

- Does not model input-dependent aleatoric uncertainty
  - But can be extended, e.g. by combining with ensembling or likelihood parametrizations
- Requires multiple forward passes during inference
- Approximation errors

References: [1] Srivastava et al. 2014. Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15, 1 (January 2014), 1929–1958.

Uncertainty estimation methods for neural networks
Ensembling

Deep Ensembles [1]

Concept: Combining multiple (uncertainty-estimating) models into a stronger one (ensembling)

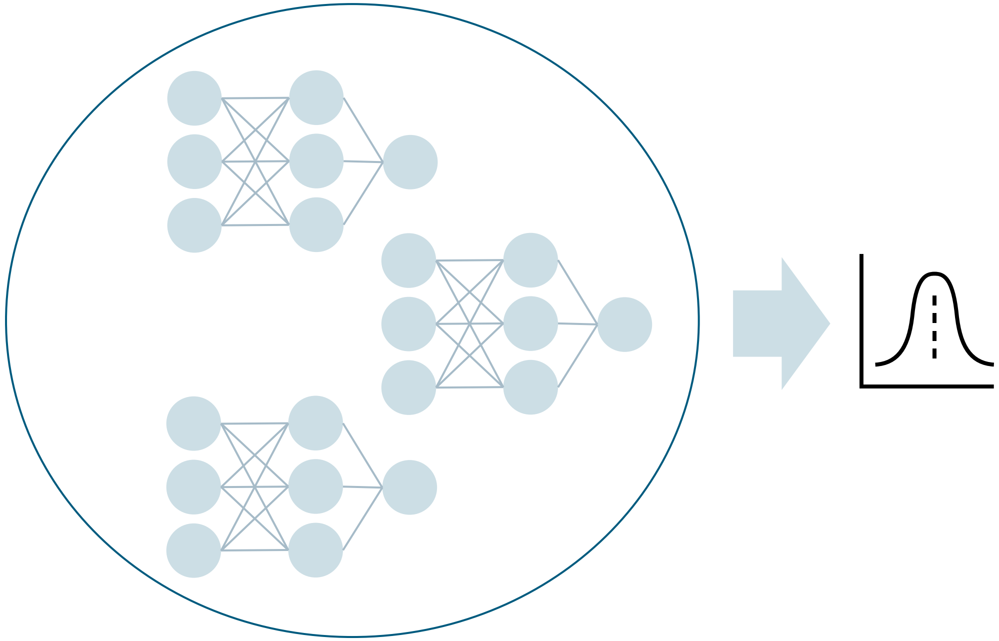

Pro:

- Easy to implement
- Captures epistemic and aleatoric uncertainty
- Easier attribution of uncertainties to concrete sources
- Easy to apply to different architectures and use cases (no change in architecture needed)

Con: High compute-overhead at training and inference (n models)

References: [1] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in neural information processing systems, 30.

Uncertainty estimation methods for neural networks

Ensembling

Current methods

Increased diversity of the ensembles

- Ensembles of models with different hyperparameters („Hyper-deep ensembles“)
- Ensembles of models with different architectures (combination with NAS)

Efficient Ensembles (e.g. dropout, sparse networks, …)

[1]

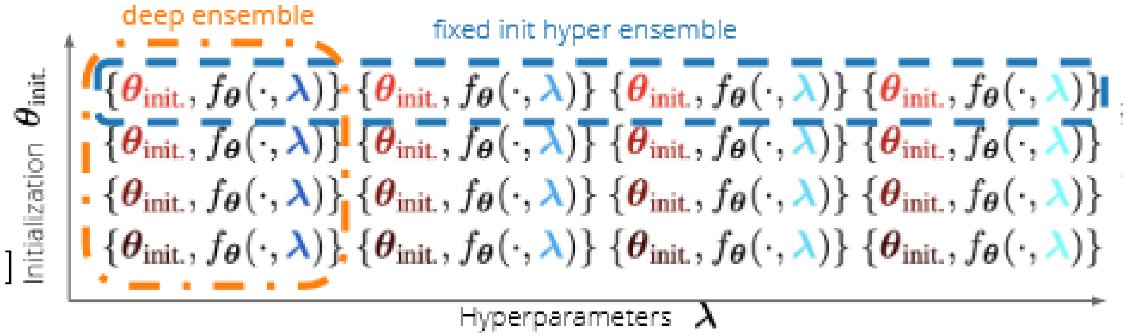

[2]

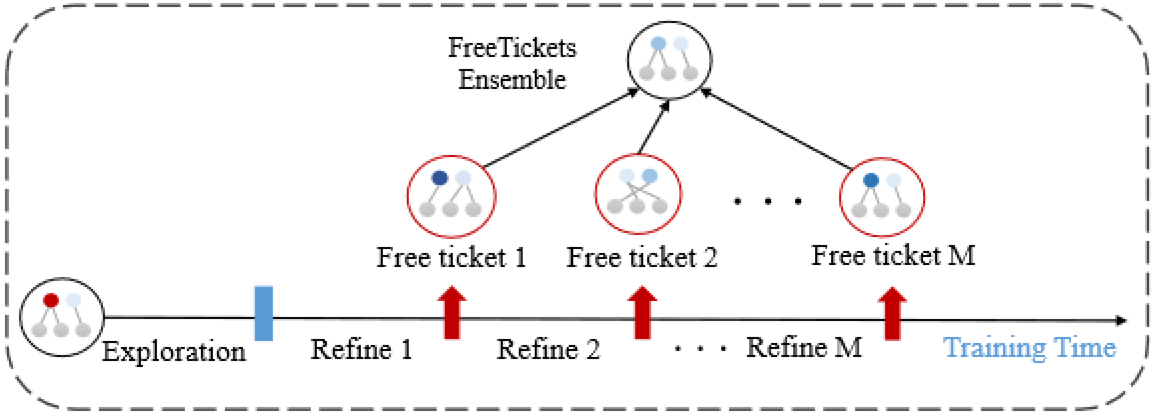

References: [1] Wenzel, F., Snoek, J., Tran, D., & Jenatton, R. (2020). Hyperparameter ensembles for robustness and uncertainty quantification. Advances in Neural Information Processing Systems, 33, 6514-6527. [2] Liu, S., Chen, T., Atashgahi, Z., Chen, X., Sokar, G., Mocanu, E., ... & Mocanu, D. C. (2021). FreeTickets: Accurate, Robust and Efficient Deep Ensemble by Training with Dynamic Sparsity. arXiv preprint arXiv:2106.14568.

Uncertainty estimation methods for neural networks
Recalibration

Concept:

- Map given (uncalibrated) uncertainty estimates (e.g. densities, confidence values) to calibrated ones using a (simple) auxiliary model (Platt/temperature scaling, histogram binning)
- Optimizing a calibration measure on a validation dataset

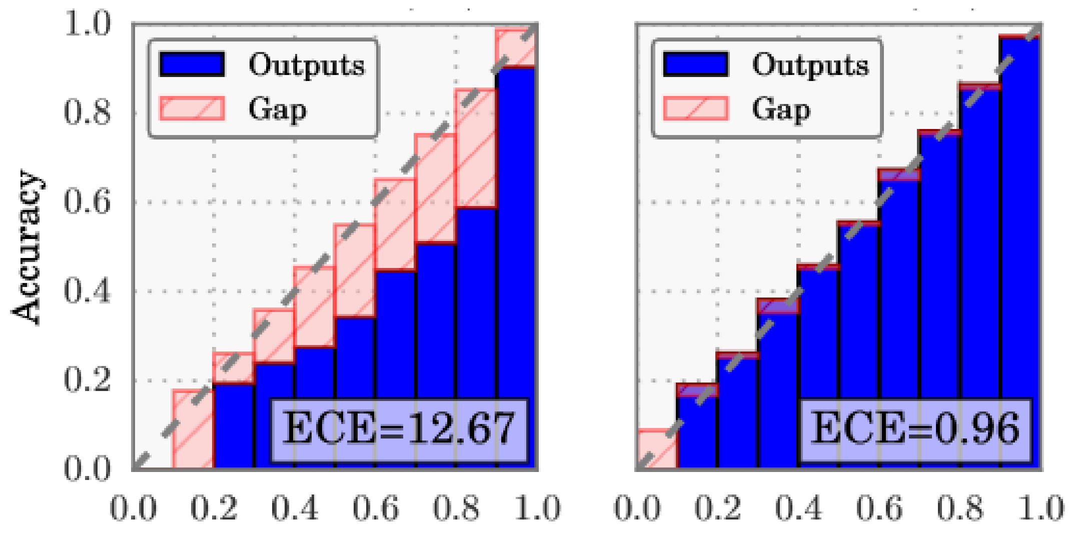

[1]

References: [1] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017, July). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321-1330). PMLR.

Uncertainty estimation methods for neural networks

Recalibration

Pro

- Can be applied after training, black-box access to the model
- Easy-to-implement
- Low computational overhead

Con

- Requires (uncalibrated) uncertainty estimates and validation data
- Calibration may overfit to validation dataset (often bad calibration on out-of-distribution inputs)
- Does not model epistemic uncertainty

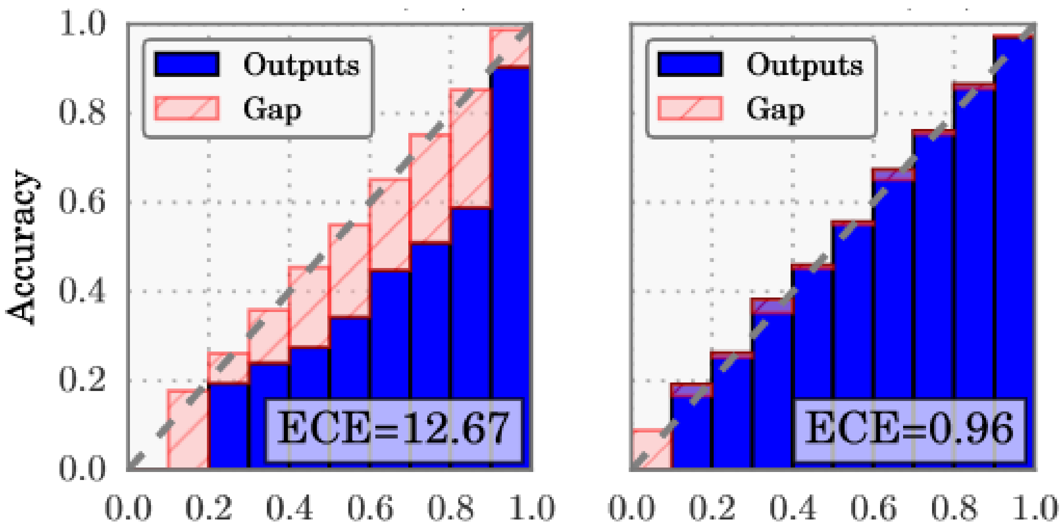

[1]

References: [1] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017, July). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321-1330). PMLR.

Conclusion

- Deep learning is affected by uncertainties
- Uncertainty estimation has several applications and plays an inportant role in reliability assessment
- Uncertainties are modeled via probability distributions
- Two types of uncertainty: aleatoric and epistemic
- Evaluation of uncertainty estimates: proper scoring rules, calibration, out-of-distribution separation, …
- Probabilistic deep learning: different methods for uncertainty estimation in neural networks, each with their own strengths and weaknesses