Zürcher Hochschule für Angewandte Wissenschaften
zh aw School of Engineering

Are you serious? Probabilistic Modelling in Deep Neural Networks

Elvis Murina, Vasily Tolkachev, Beate Sick, Oliver Dürr
Zurich University of Applied Sciences, School of Engineering, Winterthur, Switzerland

Introduction and Motivation

The output a deep neural network in a classification setting is commonly interpreted as class probabilities and no uncertainty measures are given. In this case study we demonstrate the need for credibility intervals and show how to obtain them.

Principle Question

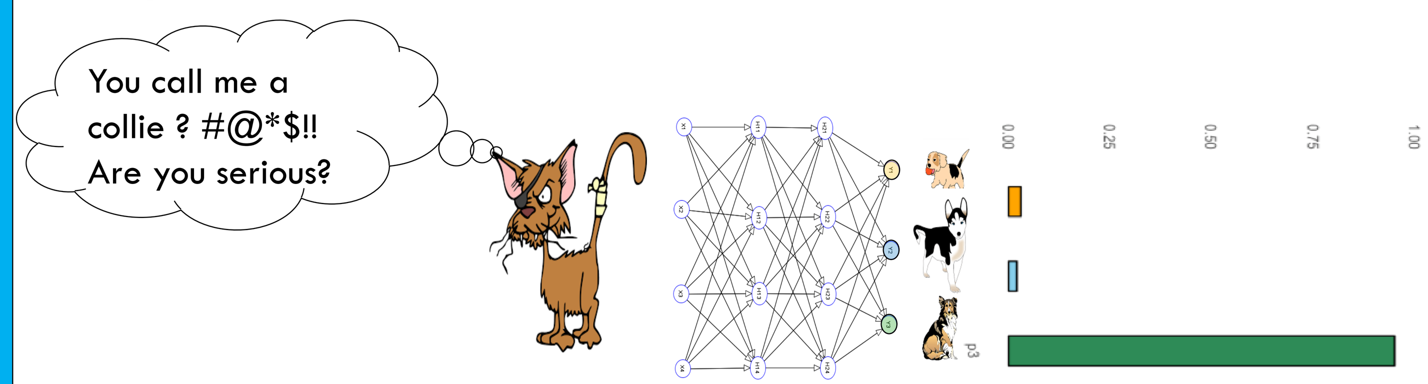

How sure is the classifier?

A case study
We use an open source data set [1] for single cell classification [2], and remove one class (mitochondria) during training.

Classes

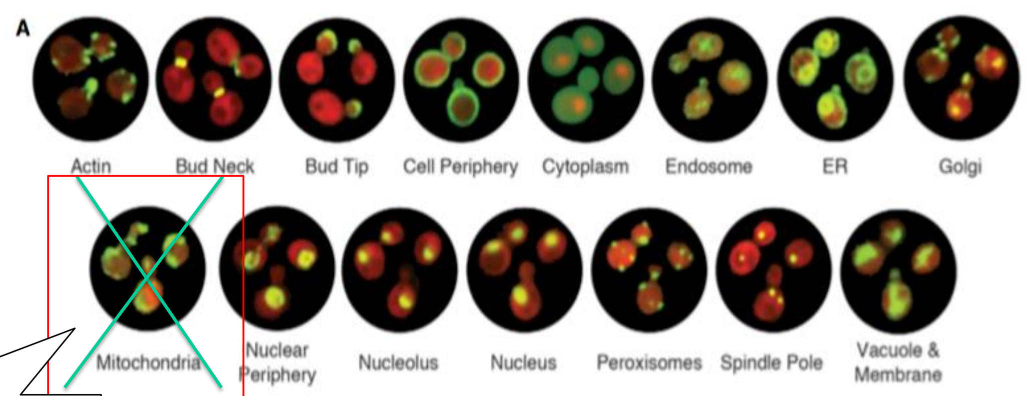

Not used for training

Predictions for unknown class

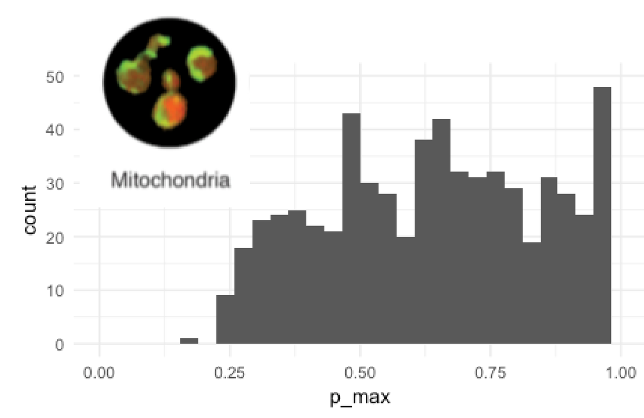

Observation:

When using the trained classifier to predict an image of a new class, it is forced to classify it as one of the classes used in training. Remarkably, quite high probabilities for obviously wrong classes are predicted.
To identify such cases an additional measure of confidence is needed.

Dropout at test time

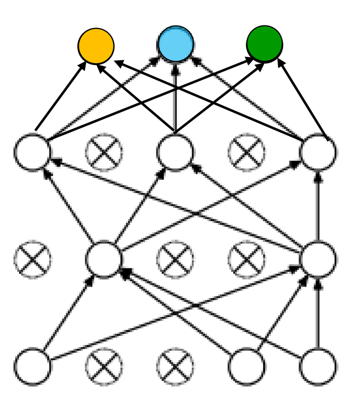

Dropout during training is commonly used as regularization method [3]. It was shown that this can be interpreted as Bayesian learning of a distribution for the weights. [4]

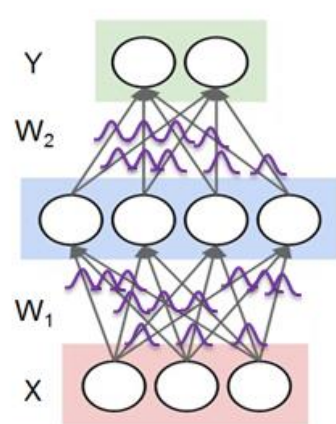

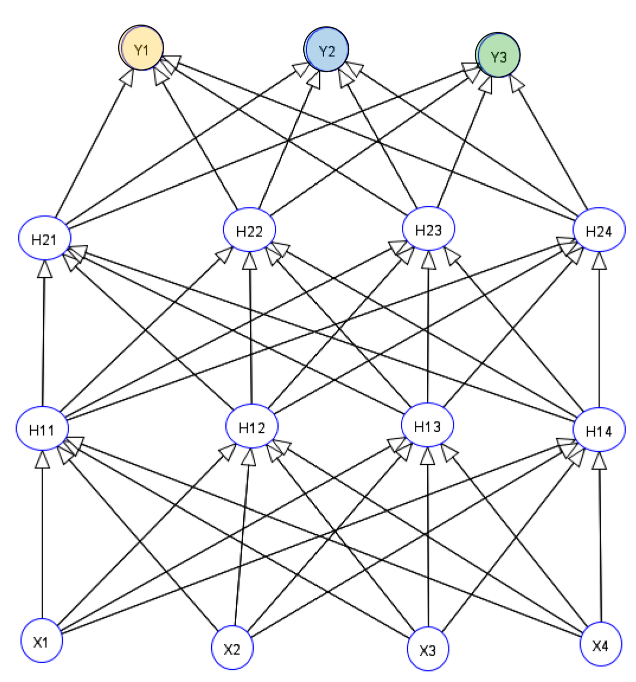

Dropout during testing samples from the weight distribution, yielding multiple predictions for each image. [4]

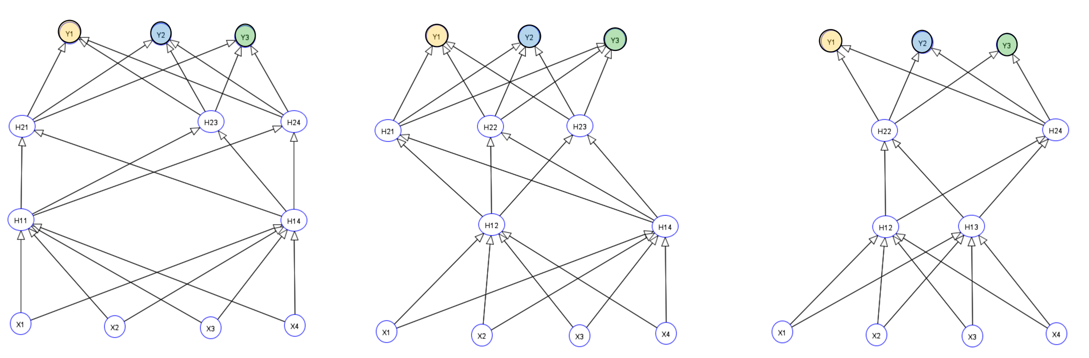

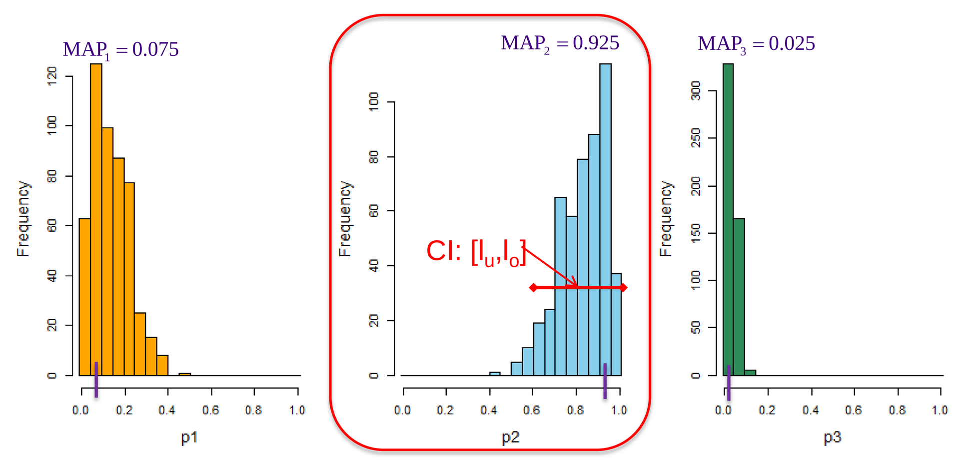

From the predictive distribution of the marginal class probabilities we compute:

- MAP (mode) as estimated class probability
- 66% CI as uncertainty measure

Architecture

The used model architecture, inspired by the VGG16 network [5] has a total of ~$3.3 \cdot 10^6$ parameters. Essential are the dropout layers which are used during training and test phase.

Results

Predictions for known classes

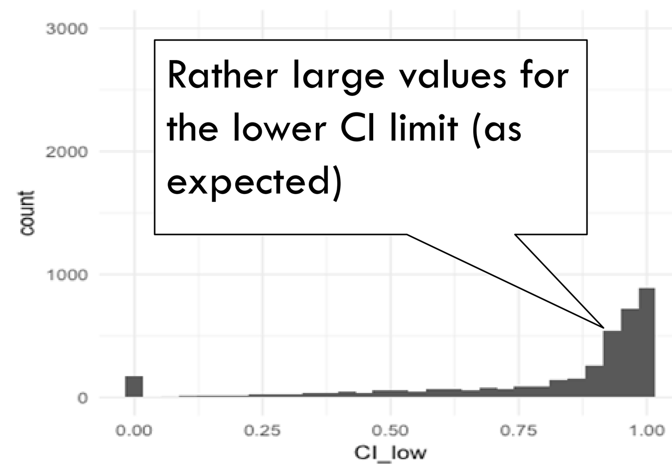

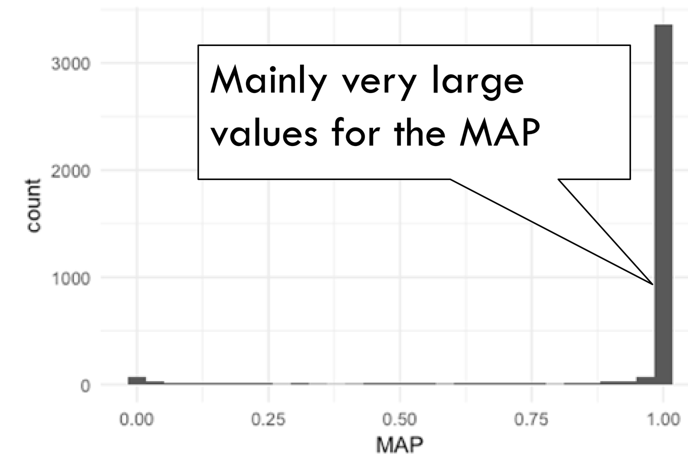

Predictions for unknown class

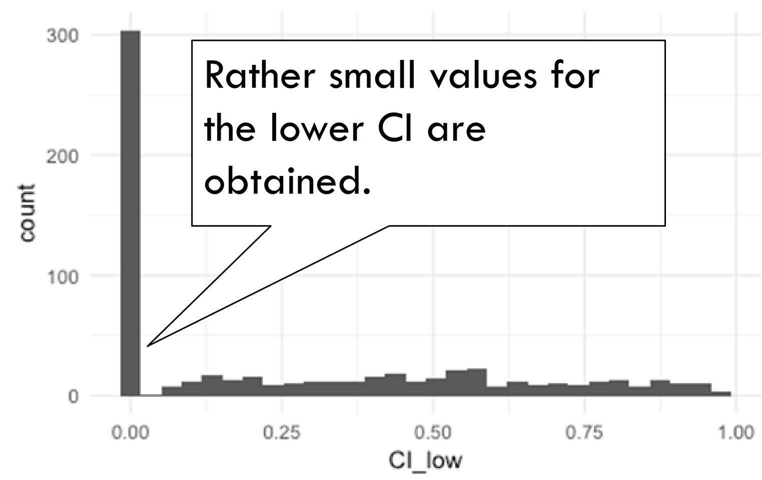

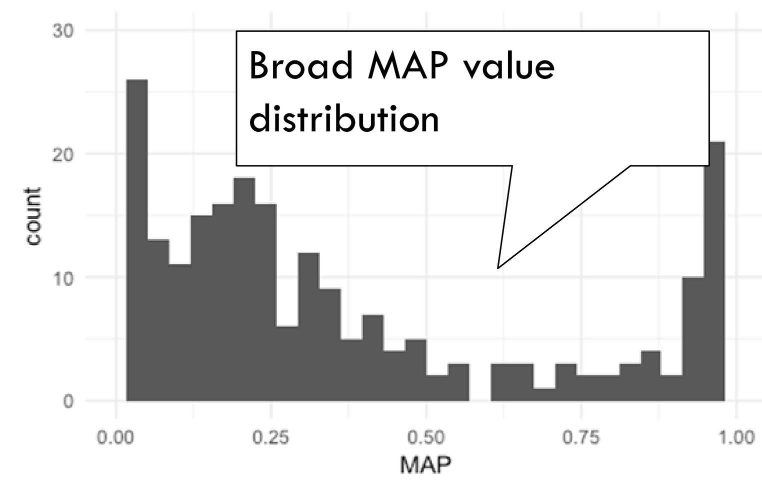

Using the lower CI as filter

After filtering out unsure images indicated by low values of the MAP or the lower CI, we obtain a higher accuracy.

Classical probability predictions (no mc dropout) for filtering yield inferior accuracy for this architecture.

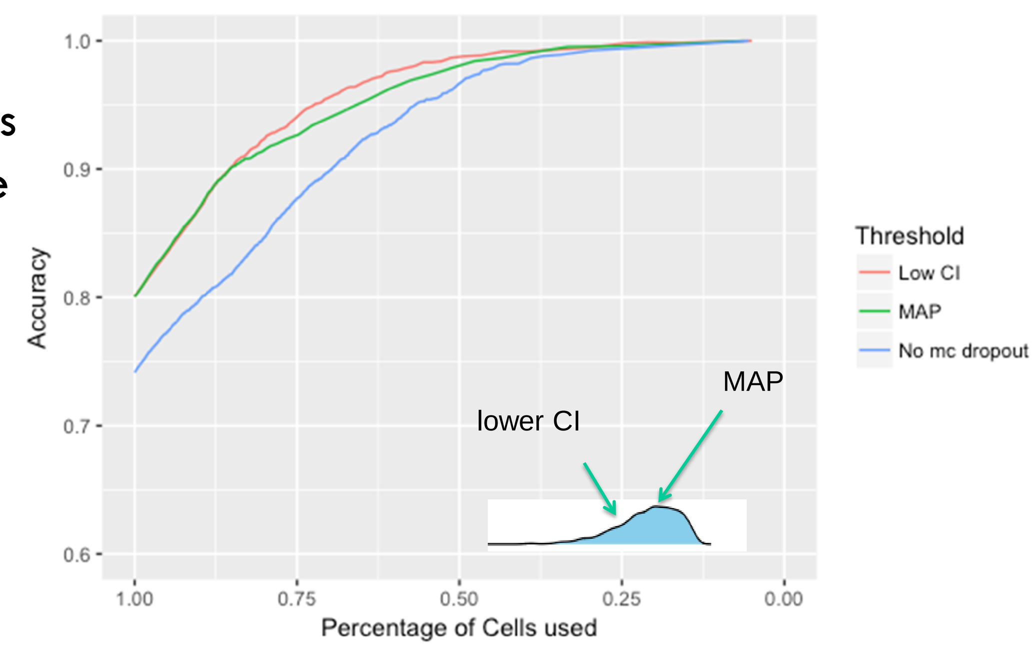

Conclusion / Outlook

- High probability values do not imply high confidence (point estimate)
- Dropout during test time can be used to obtain the predictive distribution
- The MAP of the predictive distribution can be used as a new probability estimate
- The credibility interval can be used to quantify uncertainty
    - Novel classes can be detected by low values for the low CI limit
    - Filtering out uncertain instances increases the accuracy

References

[1] Kraus, Oren Z., et al. "Automated analysis of high-content microscopy data with deep learning." *Molecular systems biology* 13.4 (2017): 924.
[2] Dürr, Oliver and Sick, Beate. "Single-cell phenotype classification using deep convolutional neural networks." *Journal of biomolecular screening* 21.9 (2016): 998-1003.
[3] Srivastava, Nitish, et al. "Dropout: a simple way to prevent neural networks from overfitting." *Journal of machine learning research* 15.1 (2014): 1929-1958.
[4] Gal, Yarin, and Zoubin Ghahramani. "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning." *international conference on machine learning*. (2016).
[5] Simonyan, Karen, and Andrew Zisserman. "Very deep convolutional networks for large-scale image recognition." *arXiv preprint arXiv:1409.1556* (2014).