

NCTV: Neural Clamping Toolkit and Visualization for Neural Network Calibration

Lei Hsiung^{1,3}, Yung-Chen Tang^{1,2}, Pin-Yu Chen³, Tsung-Yi Ho^{1,4}

¹National Tsing Hua University

²MediaTek Inc.

³IBM Research

⁴The Chinese University of Hong Kong

{hsiong, yctang}@m109.nthu.edu.tw, pin-yu.chen@ibm.com, tyho@cse.cuhk.edu.hk

Abstract

With the advancement of deep learning technology, neural networks have demonstrated their excellent ability to provide accurate predictions in many tasks. However, a lack of consideration for *neural network calibration* will not gain trust from humans, even for high-accuracy models. In this regard, the gap between the confidence of the model’s predictions and the actual correctness likelihood must be bridged to derive a well-calibrated model. In this paper, we introduce the Neural Clamping Toolkit, the first open-source framework designed to help developers employ state-of-the-art model-agnostic calibrated models. Furthermore, we provide animations and interactive sections in the demonstration to familiarize researchers with calibration in neural networks. A Colab tutorial on utilizing our toolkit is also introduced.

1 Introduction

With the increasing number of tasks that deep neural networks can handle, including medical diagnosis, image classification, natural language processing, etc., it is indispensable to increase the human trustworthiness of AI models. For example, if an AI model classifies a pathological image as malignant, a radiologist might need to know what it is based on and how likely that the prediction is correct. A confidence level is, therefore, an essential basis for physicians and radiologists to perform disease diagnosis or tumor analysis on medical images (Jiang et al. 2012; Esteva et al. 2017). Because the risk will be magnified in safety-related scenarios, it is also crucial to provide accurate prediction confidence.

However, most neural networks are not required to possess accurate confidence values when trained or landing, resulting in the potential risk of bias between confidence values and accuracy. To address this concern, one of the approaches is through *neural network calibration*, that is, making the confidence of model prediction align with its true correctness likelihood (Guo et al. 2017).

In the existing literature, two calibration approaches have been proposed: the *in-processing* way involves training or fine-tuning the model (Tian et al. 2021; Liang et al. 2020), and the *post-processing* way mainly focuses on processing or remapping the output of the pre-trained model logits (Guo et al. 2017; Esteva et al. 2017; Kull et al. 2019;

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

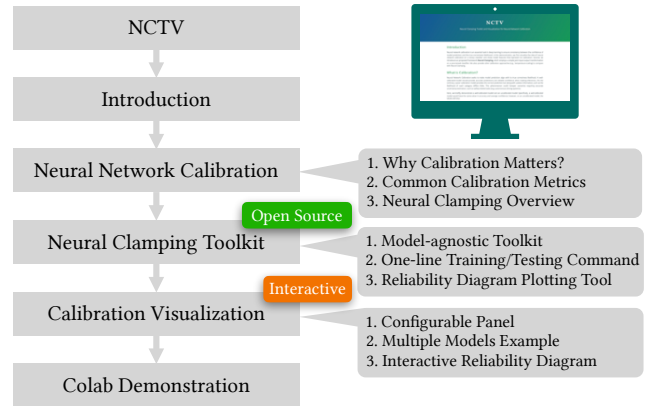


Figure 1: NCTV Overview. Browse on: hsiong.cc/NCTV

Gupta et al. 2020). However, these methods are either time-consuming and computationally expensive or suffer from a lack of effectiveness. Therefore, (Tang, Chen, and Ho 2022) proposed *Neural Clamping*, a novel post-processing calibration framework for neural networks, the first approach that utilizes a joint input-output transformation for model-agnostic calibration. At the model input and output, respectively, Neural Clamping appends a universal perturbation and temperature scaling for all classes. It provides a novel framework for post-processing calibration and includes temperature scaling as a special case. In addition, they also develop theoretical analyses and experiments to prove that this method can effectively reduce calibration error.

Extending on this idea, in this paper, we present *NCTV: Neural Clamping Toolkit and Visualization for Neural Network Calibration*, the first open-source framework and web-based demonstration to familiarize researchers and users with neural network calibration. As shown in Figure 1, NCTV also includes an interactive section that allows the user to observe real-time changes in the reliability diagram by manipulating the calibration tool in the configurable panel. We have also created a step-by-step notebook tutorial on Google Colab to guide users using our toolkit on their own models and visualize the performance before and after calibration in reliability diagrams.

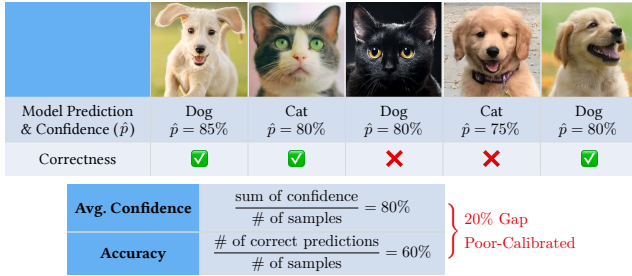


Figure 2: The example of poor-calibrated neural network.

2 NCTV: The Overview

2.1 Neural Network Calibration

Given a K -way neural network classifier $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^K$, and $x \in \mathcal{X}$ is an data sample with ground-truth $y \in \{0, \dots, K\}$. Let \hat{y} and \hat{p} denote the model f_θ predicts the most likely class of x and its confidence, we called that f_θ is *calibrated* if $\mathbb{P}(y = \hat{y} | p = \hat{p}) = \hat{p}$. Therefore, a poor-calibrated neural network might have poor alignment between the model’s predictions and its confidence levels. As shown in Figure 2, the model is poorly calibrated because there is a wide gap between its accuracy and average confidence level.

2.2 Neural Clamping Toolkit (NCToolkit)

NCToolkit possesses three characteristics: 1) it is a model-agnostic framework so that it can be directly applied to any pre-trained model, 2) it is developed in a highly modularized framework, so it is convenient for researchers to extend and operate, and 3) it could significantly outperform state-of-the-art post-processing calibration methods (Tang, Chen, and Ho 2022). The overall framework of NCToolkit is shown in Figure 3, which is composed of four main components. We introduce each component in the following.

Custom Configuration: NCToolkit is developed in the PyTorch framework, therefore, we support common model architectures and datasets. We provide the default configuration, including our pre-trained models and pre-trained Neural Clamping parameters for CIFAR-10 and ImageNet Datasets. Users can modify or extend upon specific needs.

Calibration Tool: We combine input perturbation and temperature scaling for Neural Clamping. This integration could offer state-of-the-art post-processing calibration tool. In NCToolkit, user could simply call `.train_NC()` to start calibrating.

Calibration Metric: To properly measure the model calibration, the most common metric is the Expected Calibration Error (ECE), which could be defined as $\mathbb{E}_{(x,y) \sim \mathcal{D}}(|\mathbb{P}(y = \hat{y} | p = \hat{p}) - \hat{p}|)$. Other popular metrics, such as Static Calibration Error (SCE) and Adaptive Calibration Error (ACE), are also available for developers to evaluate calibration performance.

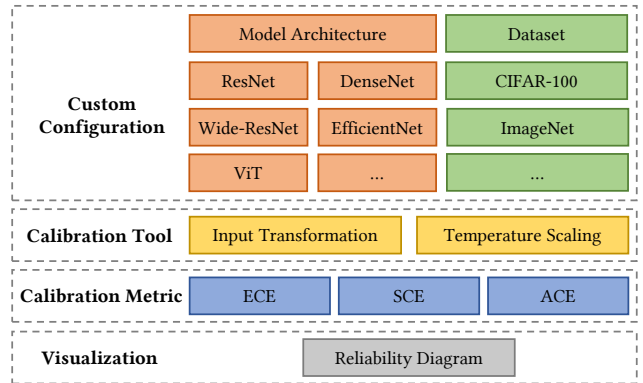


Figure 3: NCToolkit overall framework.




Figure 4: Reliability diagram examples of a neural network with and without applied Neural Clamping for calibration. The bar marked in blue (pink) denotes the actual (expected) sample accuracy, which is a function of confidence. Deviations from the expected bar line indicate calibration errors.

Visualization: We provide a visualization tool to depict calibration performance in *reliability diagrams* (DeGroot and Fienberg 1983; Niculescu-Mizil and Caruana 2005). As shown in Figure 4, the reliability diagram is composed of M bins, where each bin B_m represents the set of samples, whose prediction confidence falls within the interval $I_m = (\frac{m-1}{M}, \frac{m}{M}]$. Specifically, the accuracy of each bin can be defined as:

$$\text{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}(\hat{y}_i = y_i),$$

where $|B_m|$ represents the number of samples in B_m , and \hat{y}_i and y_i are the most likely and ground-truth labels for sample $i \in B_m$. Accordingly, the diagram will vary with the number of bins, and so does the corresponding ECE.

3 Conclusion

This demonstration enables users to gain a deeper understanding of neural network calibration. NCToolkit features a modularized and extensible framework to support developers in calibrating their model. In our demonstration, we also provide Colab tutorials on using our toolkit. Furthermore, the visualization tool is also introduced to allow users to depict the calibration performance of the model.

Acknowledgments

Part of this work was done during Lei Hsiung's visit to IBM Thomas J. Watson Research Center.

References

- DeGroot, M. H.; and Fienberg, S. E. 1983. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2): 12–22.
- Esteva, A.; Kuprel, B.; Novoa, R. A.; Ko, J.; Swetter, S. M.; Blau, H. M.; and Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *nature*, 542(7639): 115–118.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330. PMLR.
- Gupta, K.; Rahimi, A.; Ajanthan, T.; Mensink, T.; Sminchisescu, C.; and Hartley, R. 2020. Calibration of neural networks using splines. *arXiv preprint arXiv:2006.12800*.
- Jiang, X.; Osl, M.; Kim, J.; and Ohno-Machado, L. 2012. Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association*, 19(2): 263–274.
- Kull, M.; Perello Nieto, M.; Kängsepp, M.; Silva Filho, T.; Song, H.; and Flach, P. 2019. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *Advances in neural information processing systems*, 32.
- Liang, G.; Zhang, Y.; Wang, X.; and Jacobs, N. 2020. Improved trainable calibration method for neural networks on medical imaging classification. *arXiv preprint arXiv:2009.04057*.
- Niculescu-Mizil, A.; and Caruana, R. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, 625–632.
- Tang, Y.-C.; Chen, P.-Y.; and Ho, T.-Y. 2022. Neural Clamping: Joint Input Perturbation and Temperature Scaling for Neural Network Calibration. *arXiv preprint arXiv:2209.11604*.
- Tian, J.; Yung, D.; Hsu, Y.-C.; and Kira, Z. 2021. A geometric perspective towards neural calibration via sensitivity decomposition. *Advances in Neural Information Processing Systems*, 34: 26358–26369.