
Impact of Noise on Calibration and Generalisation of Neural Networks

Martin Ferienc^{*1} Ondrej Bohdal^{*2} Timothy Hospedales²³ Miguel Rodrigues¹

Abstract

Noise injection and data augmentation strategies have been effective for enhancing the generalisation and robustness of neural networks (NNs). Certain types of noise such as label smoothing and MixUp have also been shown to improve calibration. Since noise can be added in various stages of the NN’s training, it motivates the question of when and where the noise is the most effective. We study a variety of noise types to determine how much they improve calibration and generalisation, and under what conditions. More specifically we evaluate various noise-injection strategies in both in-distribution (ID) and out-of-distribution (OOD) scenarios. The findings highlight that activation noise was the most transferable and effective in improving generalisation, while input augmentation noise was prominent in improving calibration on OOD but not necessarily ID data.

1. Introduction

Noise injection methods have emerged as a promising approach to enhance the generalisation of neural networks (NNs) (Srivastava et al., 2014; Neelakantan et al., 2017). Given the importance of noise for Bayesian NNs (BNNs) (Gal & Ghahramani, 2016; Blundell et al., 2015; Welling & Teh, 2011), we hypothesise that noise injections during training of standard NNs can also positively impact their calibration. Calibration refers to the alignment of prediction’s accuracy to their confidence (Guo et al., 2017).

Examples of noise injection approaches include dropout (Srivastava et al., 2014; Gal & Ghahramani, 2016), label smoothing (Szegedy et al., 2016), MixUp (Zhang et al., 2018), Gaussian noise (Blundell et al., 2015), shrinking and per-

turbing NN weights (Ash & Adams, 2020), and gradient noise (Neelakantan et al., 2017). By introducing noise during the training, these methods encourage active exploration of the parameter space and can be applied to various components of the network, including the input, targets, activations, gradients and the model itself. In this paper, we aim to provide a fair comparison of noise injection methods during training and investigate their impact on both calibration and generalisation of NNs in a computer vision classification setting. We ensure fairness of the comparison through dedicated hyperparameter optimization per noise type and we examine the transferability of found hyperparameters from one dataset or architecture to another. To robustly evaluate both generalisation and calibration we consider testing the methods on both test in-distribution (ID) and out-of-distribution (OOD) data.

The key takeaways from our work are: 1) Activation noise, especially dropout (Srivastava et al., 2014), improves generalisation and marginally also calibration across architectures and datasets. 2) Input augmentation, MixUp (Zhang et al., 2018), improves calibration and generalisation on OOD data but not necessarily ID data. 3) Model noise and gradient noise improve generalisation and calibration, but only to a smaller extent than input or activation noise.

2. Related Work

Standard NNs were shown to lack calibration (Guo et al., 2017), motivating the need for approaches focusing on training NNs such that their confidence matches their accuracy. Bayesian NNs (BNNs) (Blundell et al., 2015; Gal & Ghahramani, 2016; Welling & Teh, 2011) and NN ensembles (Lakshminarayanan et al., 2017) are popular approaches for obtaining well-calibrated models, but they are computationally expensive as they require random sampling and multiple forward passes during test time. Alternative methods have been proposed without increasing computational complexity, particularly during training. They include different loss functions (Kumar et al., 2018; Mukhoti et al., 2020; Bohdal et al., 2021) and temperature scaling (Guo et al., 2017). However, these approaches have their own limitations and may not be suitable for all scenarios. On the other hand, most noise injections are applicable to any NN architecture and any task.

^{*}Equal contribution ¹Department of Electronic and Electrical Engineering, University College London ²School of Informatics, University of Edinburgh ³Samsung AI Center, Cambridge. Correspondence to: Martin Ferienc <martin.ferienc.19@ucl.ac.uk>, Ondrej Bohdal <ondrej.bohdal@ed.ac.uk>.

For **input noise** injection, commonly used are MixUp and Output Diversified Sampling (ODS) methods. MixUp (Zhang et al., 2018) linearly interpolates between two samples and their labels, while ODS (Tashiro et al., 2020) augments the input to diversify predictions and was used in the context of adversarial examples but not calibration. MixUp has been shown to improve calibration and generalisation (Zhang et al., 2022), but its transferability between datasets and architectures has not been explored. Additionally, we investigate naive Gaussian and uniform noise injection, which adds Gaussian or uniform noise to the input during training. In terms of **target noise** injection, label smoothing (Pereyra et al., 2017) and MixUp (Zhang et al., 2018) label interpolation are frequently used. Label smoothing replaces hard targets with soft targets and has already been shown to improve calibration, but not on OOD data (Müller et al., 2019). **Activation noise** injections include Dropout, Gaussian and uniform noise injections. Dropout (Srivastava et al., 2014) randomly sets activations to zero. Gaussian noise injection (Blundell et al., 2015; Camuto et al., 2020; Alemi et al., 2017; Yu et al., 2021) adds Gaussian noise to the activations, while uniform noise injection adds uniform noise. In BNNs, these injections are applied both during training and evaluation, whereas in this work we only apply noise during training. Furthermore, **gradient noise** has been shown to improve generalisation through adding annealed Gaussian noise to the gradients during training (Neelakantan et al., 2017; Welling & Teh, 2011). However, it was not benchmarked on calibration, especially without ensembling weights at different training time-steps. Finally for **model noise** injection, recently Gaussian noise injection via shrinking and perturbing weights (Ash & Adams, 2020) at a given epoch frequency was shown to improve retraining generalisation, but calibration on ID or OOD data was not considered.

To the best of our knowledge, the noise injections have been studied 1) separately (Zhang et al., 2022; Müller et al., 2019), 2) orthogonally for generalisation and calibration on ID or OOD data, and 3) without a unified hyperparameter (HP) optimization protocol. This research aims to start the conversation into a comprehensive analysis of the noise injection methods and their relationship to generalisation and calibration, across datasets and NN architectures, providing valuable insights into their effectiveness and practicality.

3. Methodology

This study focuses on training a NN with noise perturbations to investigate their impact on NN’s accuracy and calibration, identifying which perturbations are helpful and when. The noises are divided between **input**, **target**, **activation**, **gradient** and **model**, and their deployment during training is outlined in Algorithm 1 via blue lines. The probability

Algorithm 1 Training of Neural Network with Noise

Require: Training dataset $D = \{(x^b, y^b)\}_{b=1}^B$, B batches, learning rate η , number of epochs T , weights θ , operation $g(\cdot, \theta)$, hidden states h^b , hidden depth D , activation $f(\cdot)$, probability of applying noise to a batch p

- 1: Initialize θ randomly
- 2: **for** $t = 1$ to T **do**
- 3: **for** $b = 1$ to B **do**
- 4: Randomly select (x^b, y^b) from D
- 5: Sample $e \sim U(0, 1)$ {**If** $e < p$ }
- 6: **Input noise:** Modify x^b
- 7: **Target noise:** Modify y^b
- 8: **for** $i = 1$ to D **do**
- 9: $h_i^b = g(h_{i-1}^b, \theta)$ {Where $h_0^b = x^b$ }
- 10: **Activation noise:** Modify h_i^b before activation
- 11: $h_i^b = f(h_i^b)$
- 12: **end for**
- 13: Compute predicted output $\hat{y}^i = g(h_D^b, \theta)$
- 14: Compute loss $\mathcal{L}(\hat{y}^i, y^i)$ and gradients $\nabla_{\theta} \mathcal{L}$
- 15: **Gradient noise:** Modify $\nabla_{\theta} \mathcal{L}$
- 16: Update weights: $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}$
- 17: **end for**
- 18: **if** $t \bmod \text{frequency} = 0$ **and** $t < 0.75T$ **then**
- 19: **Model noise:** Modify θ
- 20: **end if**
- 21: **end for**
- 22: **return** θ

of applying each noise to a batch out of B batches is determined by the HP $p \in [0, 1]$, except model noise, which was applied with a selected frequency during the T training epochs. The noises have associated HPs and tuning ranges.

Input noise: The input noise consisted of 2 naive variants and 2 variants which tapped into predictions or the targets to compute the noise. The two naive variants consisted of adding Gaussian or uniformly sampled noise $n \sim U(-\sigma, \sigma)$; $n \sim \mathcal{N}(0, \sigma)$ added to inputs x with standard deviation $\sigma \in [1e^{-4}, 1e^{-1}]$. We considered ODS (Tashiro et al., 2020) with respect to $\epsilon \in [1e^{-4}, 1e^{-1}]$ and temperature $T \in [0.5, 5.0]$, and MixUp (Zhang et al., 2018) with $\alpha \in [0, 1]$ which also modified the targets accordingly. **Target noise:** In addition to MixUp we considered a static noise introduced to the labels y in the form of label smoothing (Müller et al., 2019) with the smoothing factor $l \in [0, 0.25]$. **Activation noise:** The hidden states prior to applying the activation function, $\{h_i^b\}_{i=0}^D$, where D is the depth of the net, could be disturbed by 3 types of activation noise: additive Gaussian or Uniform $n \sim U(-\sigma, \sigma)$; $n \sim \mathcal{N}(0, \sigma)$ with $\sigma \in [1e^{-4}, 1e^{-1}]$ as a tunable HP or multiplicative Dropout (Srivastava et al., 2014) that incorporates a dropout rate $d \in [0, 1]$. The activation noise was used prior to an activation $f(\cdot)$ for all linear

or convolutional operations $g(\cdot, \theta)$ but not in the output layer. **Gradient noise:** The noise applied to the gradients $\nabla_{\theta} \mathcal{L}$ followed (Neelakantan et al., 2017) with the step size $\eta \in [0, 1]$ and the annealing factor $\gamma \in [0, 1]$. **Model noise:** Lastly the model noise follows the idea of shrinking and perturbing the weights θ (Ash & Adams, 2020) with a shrink factor $\mu \in [0.0, 1.0]$ and standard deviation $\sigma \in [0.0, 1e^{-3}]$ with frequency of perturbing every frequency $\in [0, 80]$ epochs, except the last 25% of training epochs.

4. Experiments

Settings We first tune the learning rate and L2 regularisation of a no-noise network which are reused when tuning the HPs of each noise injection method on three different combinations: ResNet-18 paired with CIFAR-10 or CIFAR-100 and a fully connected (FC) network paired with SVHN. The tuning is performed using 1/4 of the training budget over the course of one day, using model-based Tree-structured Parzen Estimator method (Bergstra et al., 2011). With these settings we are able to evaluate about 40 configurations selected using Bayesian Optimization. Our protocol allows us to optimize the performance of each noise injection method and provide fair comparison. Full experimental details are in Appendix A, including a summary of the identified HPs.

To assess the effectiveness of the noise injection methods, we measure their performance using three metrics: Error [\downarrow , %], Expected Calibration Error (ECE) (Guo et al., 2017) [\downarrow , %], and Negative Log-Likelihood (NLL) [\downarrow] that we report in Appendix B. These metrics provide insights into the accuracy and its match with the confidence of the NNs’ predictions. We evaluate the performance on both the ID test set and an augmented OOD set that includes an average over visual corruptions across 19 categories and 5 severities (Hendrycks & Dietterich, 2019). These corruptions include, for example, adding snow or fog to the image, changing the brightness or saturation of the image or blurring the image. We conduct experiments on a series of deployment scenarios where 1) the tuned HPs are directly used on the tuned dataset and architecture, 2) the architecture is changed but the HPs are kept, 3) the HPs come from a different source dataset. The results presented are averaged across 3 seeds and the best results are in **bold**.

4.1. Analysis

Tuned Hyperparameters In this scenario, we evaluate the performance of the noise injection methods when the HPs are tuned specifically for the dataset and architecture. The results for these experiments are in Tables 1 and 2, and they show that activation and input augmentation noises are prominent in improving the accuracy and calibration of the networks across the datasets. Dropout was the most effective for improving ID generalisation in CIFAR-10 and

CIFAR-100, while MixUp was the most effective for SVHN. Uniform activation worked the best for improving ID calibration in CIFAR-10 and CIFAR-100, whereas ODS was the best in SVHN. The strong result obtained by ODS on SVHN shows that adversarial attack techniques may be useful also for other uses-cases, including calibration. Interestingly, some of the improvements carried to OOD data, for example, where the error on SVHN or CIFAR-100 was the lowest with MixUp or dropout. However, when considering calibration on OOD data, MixUp was dominant for CIFAR-10 and CIFAR-100. On average, dropout improved generalisation and MixUp improved calibration when considering both ID and OOD data. The naive Gaussian and uniform input noise perturbations did not bring significant improvements.

Architecture Transfer In this scenario, we assess the performance of the noise injection methods when the HPs are transferred to a different architecture while keeping the dataset constant. We conduct experiments using SVHN with ResNet-18 with HPs tuned on an FC network. Furthermore, we use HPs tuned for ResNet-18 for both CIFAR-10 and CIFAR-100 and we change the architecture to WideResNet-18. The results are presented in Tables 3 and 4. Considering the performance on ID data, we see that dropout reduced error across architectures and also improved calibration. Contrary to the improvements seen on SVHN when using FC, MixUp did not reduce the error when using ResNet-18 and it even recorded worse performance on OOD data than no noise at all. Switching focus to OOD data, model perturbation moderately improved calibration for CIFAR-100, while activations had a negative impact and led to worse calibration. Even though WideResNet-18 and ResNet-18 are relatively similar, transferring hyperparameters for example for MixUp in CIFAR-100, did not prove efficient as seen in calibration on OOD data which became worse than not using any noise at all. In summary, activation noises, most notably dropout, performed well on improving generalisation on both ID and OOD data and moderately on calibration on ID data. However, no method was able to consistently improve calibration on OOD data after the architecture was changed.

Dataset Transfer Under these settings, we investigate the transferability of hyperparameters by evaluating the noise injection methods on the same architectures but using different datasets. Specifically, we evaluate SVHN with ResNet-18 and HPs from CIFAR-100/ResNet-18, CIFAR-10 with ResNet-18 and CIFAR-100/ResNet-18 HPs, and CIFAR-100 with ResNet-18 but with CIFAR-10/ResNet-18 HPs. The results are shown in Tables 5 and 6. For all SVHN, CIFAR-10, CIFAR-100, the most significant error improvements across ID or OOD data were achieved using dropout and Gaussian noise. Interestingly, the activation Gaussian noise was able to improve calibration on both ID and OOD data on CIFAR-100, but not on the other datasets. MixUp

Table 1. Error [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with tuned hyperparameters.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	17.1±0.1	20.8±0.0	10.8±0.2	33.9±0.5	37.9±0.5	59.5±0.2
INPUT GAUSSIAN	17.4±0.1	21.0±0.0	11.1±0.1	34.1±0.1	37.9±0.6	59.5±0.3
INPUT UNIFORM	17.2±0.2	20.9±0.1	11.2±0.0	33.9±0.4	37.9±0.3	59.5±0.3
INPUT ODS	15.2±0.2	18.7±0.2	10.9±0.3	31.6±0.4	38.0±0.4	58.2±0.3
INPUT-TARGET MIXUP	13.9±0.1	17.5±0.1	9.8±0.0	28.0±0.2	38.4±0.2	58.3±0.2
LABEL SMOOTHING	16.5±0.2	20.1±0.1	11.5±0.1	33.8±0.2	38.6±0.6	60.5±0.5
ACTIVATION GAUSSIAN	17.2±0.2	20.8±0.2	9.0±0.2	31.4±0.5	34.9±0.1	58.1±0.2
ACTIVATION UNIFORM	17.0±0.2	20.7±0.2	8.9±0.1	31.8±0.5	34.6±0.3	57.8±0.2
ACTIVATION DROPOUT	14.6±0.0	18.0±0.0	8.7±0.3	30.9±0.6	31.4±0.3	56.6±0.3
GRADIENT	17.2±0.1	21.0±0.1	11.3±0.3	33.9±0.2	38.1±0.3	59.9±0.2
MODEL	16.4±0.1	20.1±0.1	10.8±0.1	33.6±0.4	38.1±0.3	59.8±0.0

Table 3. Error [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with changed architecture.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	5.2±0.1	9.4±0.1	14.9±0.1	38.2±0.2	37.4±0.5	61.4±0.4
INPUT GAUSSIAN	5.2±0.1	9.4±0.0	14.7±0.3	38.0±0.2	37.4±0.4	61.3±0.3
INPUT UNIFORM	5.2±0.1	9.5±0.1	14.7±0.4	38.3±0.2	37.3±0.4	61.4±0.3
INPUT ODS	19.8±7.7	25.6±8.1	14.5±0.1	36.1±0.1	36.9±0.6	60.2±0.4
INPUT-TARGET MIXUP	5.3±0.1	13.1±0.5	11.8±0.3	31.5±0.3	36.5±0.5	59.5±0.6
LABEL SMOOTHING	5.1±0.1	9.3±0.2	15.2±0.3	37.7±0.1	36.8±0.3	60.9±0.1
ACTIVATION GAUSSIAN	5.3±0.1	9.5±0.1	12.0±0.1	35.7±0.5	43.2±0.8	66.0±0.7
ACTIVATION UNIFORM	5.2±0.1	9.5±0.1	12.9±0.2	37.9±0.7	36.9±0.3	61.0±0.3
ACTIVATION DROPOUT	4.2±0.1	8.0±0.1	11.1±0.2	33.2±0.5	32.1±0.3	58.3±0.6
GRADIENT	5.5±0.0	10.0±0.2	15.0±0.1	37.9±0.4	37.3±0.3	61.2±0.5
MODEL	4.4±0.0	8.5±0.0	14.7±0.4	38.1±0.1	36.5±0.2	61.1±0.0

Table 5. Error [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with changed dataset.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	5.3±0.2	9.5±0.1	11.6±0.3	33.6±0.3	38.2±0.2	59.8±0.1
INPUT GAUSSIAN	5.2±0.0	9.5±0.1	11.6±0.2	33.4±0.2	38.2±0.2	59.7±0.2
INPUT UNIFORM	5.1±0.1	9.4±0.1	11.2±0.3	34.0±0.4	38.2±0.2	59.9±0.0
INPUT ODS	5.2±0.2	9.3±0.1	11.5±0.1	32.4±0.5	37.5±0.1	56.8±0.2
INPUT-TARGET MIXUP	5.1±0.1	11.6±0.3	10.9±0.1	28.2±0.5	38.0±0.3	58.7±0.1
LABEL SMOOTHING	5.2±0.1	9.4±0.1	11.3±0.3	34.2±0.3	37.3±0.0	58.7±0.1
ACTIVATION GAUSSIAN	5.2±0.2	9.3±0.2	10.4±0.3	31.4±0.3	35.5±0.2	57.8±0.1
ACTIVATION UNIFORM	5.2±0.1	9.4±0.0	13.6±0.4	37.0±0.6	35.6±0.3	57.8±0.0
ACTIVATION DROPOUT	4.1±0.0	7.9±0.0	9.4±0.1	30.0±0.9	32.1±0.3	56.6±0.4
GRADIENT	6.1±0.2	10.9±0.2	11.8±0.3	33.4±0.2	38.2±0.2	61.0±0.2
MODEL	4.4±0.1	8.6±0.1	12.0±0.2	33.8±0.2	38.2±0.2	59.6±0.1

has demonstrated varying results, for example on SVHN or CIFAR-10 the calibration on ID data was worse than not using any noise at all, while in CIFAR-100 there was a marginal improvement. Nevertheless on OOD data MixUp was able to improve calibration across all datasets.

Summary The effectiveness of noise injections varies with the dataset and architecture. Nevertheless, especially in the tuned regime, certain settings of different noises improved both generalisation and calibration. Activation noise injections demonstrated promising results for error reduction across ID data, while input augmentations seemed to be the most effective for OOD data. Dropout was the most effective in improving error on ID or OOD data, and it proved to be transferable across architectures and datasets. MixUp was the best in improving the performance on OOD data in terms of calibration and accuracy, but not necessarily on the ID data. Interestingly, hidden in its mediocrity, model noise was able to marginally improve accuracy and calibration across majority of considered scenarios. Additional evaluation in terms of NLL, in Appendix B, has shown MixUp,

Table 2. ECE [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with tuned hyperparameters.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	15.5±0.1	18.3±0.0	5.4±0.1	12.7±0.2	16.9±0.5	15.2±0.3
INPUT GAUSSIAN	15.8±0.1	18.5±0.1	5.3±0.2	12.8±0.4	16.9±0.5	15.3±0.1
INPUT UNIFORM	15.6±0.2	18.4±0.2	5.5±0.1	12.8±0.4	17.3±0.1	15.3±0.1
INPUT ODS	4.6±0.1	5.6±0.1	5.6±0.1	11.8±0.3	16.0±0.5	15.0±0.1
INPUT-TARGET MIXUP	9.4±0.1	9.9±0.1	11.0±0.2	11.3±0.2	16.0±0.2	13.8±0.3
LABEL SMOOTHING	15.0±0.1	14.1±0.1	10.1±0.1	11.5±0.1	28.8±1.1	20.3±0.8
ACTIVATION GAUSSIAN	15.5±0.2	18.3±0.1	4.2±0.2	14.9±0.7	11.6±0.5	23.3±0.8
ACTIVATION UNIFORM	15.4±0.2	18.3±0.2	3.8±0.1	15.8±0.5	9.9±0.2	20.6±0.2
ACTIVATION DROPOUT	6.5±0.2	7.8±0.1	5.3±0.3	13.4±0.5	12.2±0.3	25.8±0.0
GRADIENT	15.6±0.1	18.5±0.1	5.4±0.1	13.0±0.1	16.4±0.5	14.8±0.1
MODEL	13.7±0.2	16.3±0.1	5.3±0.1	12.7±0.5	15.4±0.2	14.4±0.2

Table 4. ECE [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with changed architecture.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	4.1±0.1	7.2±0.0	7.3±0.1	11.9±0.3	18.2±0.3	15.5±0.1
INPUT GAUSSIAN	4.1±0.0	7.2±0.0	7.1±0.1	11.8±0.3	17.9±0.2	15.4±0.3
INPUT UNIFORM	4.1±0.0	7.2±0.1	7.1±0.3	12.0±0.5	18.3±0.3	15.3±0.3
INPUT ODS	6.7±1.5	9.1±2.2	7.1±0.1	11.4±0.0	17.8±0.3	15.4±0.5
INPUT-TARGET MIXUP	10.2±0.5	9.9±0.5	12.6±0.3	11.7±0.1	18.8±0.4	18.1±0.5
LABEL SMOOTHING	21.2±0.0	20.2±0.0	12.0±0.3	11.6±0.1	29.0±0.3	19.5±0.3
ACTIVATION GAUSSIAN	4.2±0.0	7.2±0.1	5.6±0.1	11.7±0.4	10.5±0.8	22.2±1.1
ACTIVATION UNIFORM	4.2±0.1	7.2±0.1	6.5±0.1	12.3±0.8	8.6±0.4	19.1±1.2
ACTIVATION DROPOUT	3.6±0.0	6.6±0.1	7.0±0.2	11.1±0.3	9.2±0.3	22.3±0.6
GRADIENT	4.2±0.1	7.5±0.1	7.3±0.1	11.3±0.4	17.7±0.7	15.3±0.3
MODEL	3.4±0.0	6.1±0.1	6.9±0.2	11.8±0.4	15.7±0.4	14.6±0.4

Table 6. ECE [\downarrow , %] comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with changed dataset.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	4.2±0.2	7.2±0.1	4.4±0.0	14.9±1.0	17.1±0.2	14.9±0.3
INPUT GAUSSIAN	4.1±0.1	7.2±0.1	4.6±0.2	15.1±1.2	16.7±0.1	15.2±0.1
INPUT UNIFORM	4.1±0.1	7.2±0.0	5.5±0.1	12.5±0.4	17.2±0.2	15.1±0.1
INPUT ODS	4.2±0.1	7.1±0.1	4.7±0.2	14.8±1.4	16.3±0.0	14.8±0.1
INPUT-TARGET MIXUP	6.6±0.6	6.8±0.5	6.5±0.1	8.9±0.4	16.9±0.5	14.4±0.4
LABEL SMOOTHING	15.9±0.1	15.2±0.1	19.0±0.3	14.4±0.3	25.7±0.1	19.7±0.3
ACTIVATION GAUSSIAN	4.2±0.1	7.2±0.1	5.5±0.1	18.5±0.5	9.3±0.3	12.9±0.4
ACTIVATION UNIFORM	4.2±0.1	7.1±0.1	4.6±0.3	12.6±0.4	8.6±0.1	13.0±0.3
ACTIVATION DROPOUT	3.6±0.0	6.5±0.0	5.6±0.1	19.3±0.7	9.9±0.2	20.8±0.7
GRADIENT	4.7±0.1	8.2±0.2	4.7±0.2	14.9±0.6	12.8±0.3	13.3±0.2
MODEL	3.3±0.1	5.7±0.1	4.8±0.1	15.1±0.5	16.8±0.2	15.0±0.2

dropout and model perturbations to be the most effective.

5. Discussion and Conclusions

Our study explored the impact of various noise injection strategies on the calibration and generalisation of NNs. We identified activation noise to be the most effective and transferable for improving generalisation, while input augmentation noise notably enhanced calibration on OOD data. Despite sole focus on computer vision classification, our research underlines the potential of noise injection as a beneficial tool for boosting NN generalisation and calibration.

Acknowledgements

Martin Ferianc was sponsored through a scholarship from the Institute of Communications and Connected Systems at UCL. Ondrej Bohdal was supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh.

References

- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *ICLR*, 2017.
- Ash, J. and Adams, R. P. On warm-starting neural network training. In *NeurIPS*, 2020.
- Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. Algorithms for hyper-parameter optimization. In *NIPS*, 2011.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *ICML*, 2015.
- Bohdal, O., Yang, Y., and Hospedales, T. Meta-calibration: Learning of model calibration using differentiable expected calibration error. In *ICML UDL Workshop*, 2021.
- Camuto, A., Willetts, M., Simsekli, U., Roberts, S. J., and Holmes, C. C. Explicit regularisation in gaussian noise injections. In *NeurIPS*, 2020.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *ICML*, 2017.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019.
- Kumar, A., Sarawagi, S., and Jain, U. Trainable calibration measures for neural networks from kernel mean embeddings. In *ICML*, 2018.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*, 2017.
- Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.
- Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P., and Dokania, P. Calibrating deep neural networks using focal loss. In *NeurIPS*, 2020.
- Müller, R., Kornblith, S., and Hinton, G. E. When does label smoothing help? In *NeurIPS*, 2019.
- Neelakantan, A., Vilnis, L., Le, Q. V., Kaiser, L., Kurach, K., Sutskever, I., and Martens, J. Adding gradient noise improves learning for very deep networks. In *OpenReview*, 2017.
- Pereyra, G., Tucker, G., Chorowski, J., Kaiser, Ł., and Hinton, G. Regularizing neural networks by penalizing confident output distributions. In *ICLR Workshop*, 2017.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.
- Tashiro, Y., Song, Y., and Ermon, S. Diversity can be transferred: Output diversification for white-and black-box attacks. In *NeurIPS*, 2020.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011.
- Yu, T., Yang, Y., Li, D., Hospedales, T., and Xiang, T. Simple and effective stochastic neural networks. In *AAAI*, 2021.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.
- Zhang, L., Deng, Z., Kawaguchi, K., and Zou, J. When and how mixup improves calibration. In *ICML*, 2022.

Table 7. NLL \downarrow comparison on in-distribution (ID) and out-of-distribution (OOD) test sets with tuned hyperparameters.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	2.1±0.0	2.4±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT GAUSSIAN	2.1±0.0	2.4±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT UNIFORM	2.1±0.0	2.4±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT ODS	0.6±0.0	0.7 ±0.0	0.4±0.0	1.1±0.0	1.9±0.0	2.8±0.0
INPUT-TARGET MIXUP	0.5 ±0.0	0.7 ±0.0	0.4±0.0	0.9 ±0.0	1.8±0.0	2.7 ±0.0
LABEL SMOOTHING	0.7±0.0	0.8±0.0	0.5±0.0	1.1±0.0	2.1±0.0	3.1±0.0
ACTIVATION GAUSSIAN	2.1±0.0	2.4±0.0	0.3 ±0.0	1.1±0.0	1.4±0.0	2.8±0.1
ACTIVATION UNIFORM	2.1±0.0	2.4±0.0	0.3 ±0.0	1.2±0.0	1.4±0.0	2.8±0.0
ACTIVATION DROPOUT	0.6±0.0	0.7 ±0.0	0.3 ±0.0	1.1±0.0	1.3 ±0.0	2.9±0.0
GRADIENT	2.1±0.0	2.5±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
MODEL	1.5±0.0	1.7±0.0	0.4±0.0	1.2±0.0	1.8±0.0	2.9±0.0

A. Experimental Settings Details

We used stochastic gradient descent with momentum 0.9 to train all the networks. The learning rate $\eta = [1e-4, 1e-1]$ together with L2 regularisation $\lambda = [1e-7, 1e-1]$ were initially tuned and then reused for each noise injection method. We used cosine annealing learning rate schedule without restarts (Loshchilov & Hutter, 2017) for all experiments. For each dataset, we only used normalization without any further data augmentations. We used gradient norm clipping of 5.0 to stabilise the training. The batch size was set to 256 for all experiments. 10% of the training data was used as the validation set to select the best model. The tuning was performed with 1 seed and the winning hyperparameters were retrained 3 times with different seeds. The final results are reported as the average of the 3 runs. We used cross-entropy loss for all the experiments. In all cases we trained the networks for 200 epochs.

We used a fully connected network with 4 hidden layers of 150 units followed by ReLU activations and ResNet-18 with [32, 64, 128, 256] channels in 4 stages with [2, 2, 2, 2] blocks with strides [1, 2, 2, 2]. For WideResNet-18 we used the same channel, stride and block configuration as for ResNet-18, but with the Bottleneck block, expansion factor 4, base width 32, base width multiplier 2 and a single group. Both residual architectures use batch normalisation and ReLU activations. We used the default PyTorch weight initialization for all layers. We used 10 bins to measure ECE and a small ϵ of $1e^{-8}$ which was added to the output softmax probabilities to avoid NaNs. The found hyperparameters for each dataset-architecture pair are in Table 10. The descriptions of the hyperparameters as well as the search ranges are provided in Section 3.

B. Negative Log-Likelihood

We provide additional Tables 7, 8 and 9 that report the Negative Log-Likelihood for the different experiments that we have conducted. Comparing the NLL results to the previously collected results on the error and calibration error, there are subtle differences. In Table 7 the NLL for SVHN is

Table 8. NLL \downarrow comparison on in-distribution (ID) and out-of-distribution (OOD) test sets with changed architecture.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	0.3±0.0	0.6±0.0	0.5±0.0	1.2±0.0	1.8±0.0	3.0±0.0
INPUT GAUSSIAN	0.3±0.0	0.6±0.0	0.5±0.0	1.2±0.0	1.8±0.0	2.9±0.0
INPUT UNIFORM	0.3±0.0	0.6±0.0	0.5±0.0	1.3±0.0	1.8±0.0	2.9±0.0
INPUT ODS	0.8±0.3	1.1±0.3	0.5±0.0	1.2±0.0	1.8±0.0	2.9±0.0
INPUT-TARGET MIXUP	0.3±0.0	0.5±0.0	0.5±0.0	1.0 ±0.0	1.7±0.0	2.8 ±0.0
LABEL SMOOTHING	0.4±0.0	0.5±0.0	0.6±0.0	1.2±0.0	2.0±0.0	3.1±0.0
ACTIVATION GAUSSIAN	0.3±0.0	0.6±0.0	0.4 ±0.0	1.2±0.0	1.6±0.0	3.3±0.0
ACTIVATION UNIFORM	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.4±0.0	2.9±0.1
ACTIVATION DROPOUT	0.3±0.0	0.6±0.0	0.4 ±0.0	1.1±0.0	1.2 ±0.0	2.8 ±0.0
GRADIENT	0.3±0.0	0.6±0.0	0.5±0.0	1.2±0.0	1.8±0.0	2.9±0.0
MODEL	0.2 ±0.0	0.4 ±0.0	0.5±0.0	1.2±0.0	1.7±0.0	2.9±0.0

Table 9. NLL \downarrow comparison on in-distribution (ID) and out-of-distribution (OOD) test sets with changed dataset.

NOISE TYPE	SVHN		CIFAR-10		CIFAR-100	
	ID	OOD	ID	OOD	ID	OOD
NO NOISE	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT GAUSSIAN	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT UNIFORM	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0
INPUT ODS	0.3±0.0	0.6±0.0	0.4±0.0	1.1±0.0	1.8±0.0	2.7 ±0.0
INPUT-TARGET MIXUP	0.2 ±0.0	0.4 ±0.0	0.4±0.0	0.9 ±0.0	1.8±0.0	2.7 ±0.0
LABEL SMOOTHING	0.3±0.0	0.5±0.0	0.6±0.0	1.2±0.0	2.0±0.0	3.0±0.0
ACTIVATION GAUSSIAN	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.6±0.0	2.7 ±0.0
ACTIVATION UNIFORM	0.3±0.0	0.6±0.0	0.4±0.0	1.2±0.0	1.5±0.0	2.7 ±0.0
ACTIVATION DROPOUT	0.3±0.0	0.6±0.0	0.3 ±0.0	1.3±0.1	1.2 ±0.0	2.7 ±0.1
GRADIENT	0.4±0.0	0.7±0.0	0.4±0.0	1.2±0.0	1.8±0.0	2.9±0.0
MODEL	0.2 ±0.0	0.4 ±0.0	0.4±0.0	1.2±0.0	1.9±0.0	2.9±0.0

the lowest for MixUp and not ODS for ID and OOD data as suggested by ECE. Furthermore, we can see in Table 8 that when we change the architecture, the model perturbation achieved the lowest NLL on both ID and OOD data for SVHN, while the dropout activation was the best in terms of the error. Dropout retained the lowest NLL in CIFAR-100 and CIFAR-10 on ID data. Looking at Table 9 that reports results with the changed dataset, surprisingly model perturbation was the most dominant in NLL on SVHN, in contrast again to dropout as suggested by the error. Even though for CIFAR-10 the calibration on ID data was the lowest without using any noise, dropout was able to improve the NLL over no noise settings. In summary and contrast to the previous results, in NLL improvements dropout was perceived as the most effective together with MixUp and model perturbations.

Table 10. Found hyperparameters on SVHN, CIFAR-10 and CIFAR-100 datasets.

NOISE TYPE	HYPERPARAMETER	SVHN	CIFAR-10	CIFAR-100
NO NOISE	LEARNING RATE	0.042	0.0082	0.060
	L2	1.3E-07	0.029	0.0031
INPUT GAUSSIAN	σ	0.0029	0.00011	0.00061
	p	0.56	0.83	0.78
INPUT UNIFORM	σ	0.00086	0.0038	0.0013
	p	0.28	0.63	0.61
INPUT ODS	T	1.37	0.85	2.26
	η	0.047	0.0056	0.00080
	p	0.96	0.066	0.67
INPUT-TARGET MIXUP	α	0.99	0.55	0.56
	p	0.91	0.72	0.69
LABEL SMOOTHING	l	0.24	0.11	0.18
	p	0.99	0.65	0.96
ACTIVATION GAUSSIAN	σ	0.0035	0.0064	0.099
	p	1.0	0.54	0.47
ACTIVATION UNIFORM	σ	0.0014	0.014	0.058
	p	0.02	0.70	0.63
ACTIVATION DROPOUT	d	0.16	0.31	0.26
	p	0.78	0.55	0.75
GRADIENT	η	0.020	0.26	0.32
	γ	0.63	0.99	0.99
	p	0.0083	0.19	0.021
MODEL	μ	0.32	0.75	0.57
	σ	9.7E-05	0.00087	0.00016
	FREQUENCY	52	28	12