Navigating Noise: A Study of How Noise Influences Generalisation and Calibration of Neural Networks

Anonymous authors
Paper under double-blind review

Abstract

Enhancing the generalisation abilities of neural networks (NNs) through integrating noise such as MixUp or Dropout during training has emerged as a powerful and adaptable technique. Despite the proven efficacy of noise in NN training, there is no consensus regarding which noise sources, types and placements yield maximal benefits in generalisation and confidence calibration. This study thoroughly explores diverse noise modalities to evaluate their impacts on NN's generalisation and calibration under in-distribution or out-of-distribution settings, paired with experiments investigating the metric landscapes of the learnt representations, across a spectrum of NN architectures, tasks, and datasets. Our study shows that AugMix and weak augmentation exhibit cross-task effectiveness in computer vision, emphasising the need to tailor noise to specific domains. Our findings emphasise the efficacy of combining noises and successful hyperparameter transfer within a single domain but the difficulties in transferring the benefits to other domains. Furthermore, the study underscores the complexity of simultaneously optimising for both generalisation and calibration, emphasising the need for practitioners to carefully consider noise combinations and hyperparameter tuning for optimal performance in specific tasks and datasets.

1 Introduction

Neural networks (NNs) have demonstrated remarkable capabilities across various tasks, yet they often grapple with overfitting to training data, resulting in suboptimal generalisation performance on unseen samples (Srivastava et al., 2014; Bishop, 1995; Sietsma & Dow, 1991). Addressing this issue, conventional techniques such as weight decay (Krogh & Hertz, 1991) and early stopping (Prechelt, 2002) have been employed to regularise NN training. Alongside these methods, the introduction of noise during the NN's training has emerged as a potent strategy to enhance generalisation (Sietsma & Dow, 1991; Neelakantan et al., 2017; Camuto, 2021; Kukačka et al., 2017). Diverging from weight decay and early stopping that modulate the model's search within the hypothesis space, noise injections embrace randomness during training, fostering exploration of a broader array of representations (He et al., 2019). The appeal of noise injections extends further due to their versatile applicability across diverse tasks, datasets, and NN architectures. These attributes establish noise injections as a convenient approach for enhancing NN's algorithmic performance.

Various noise injection methodologies have been proposed, encompassing **activation** techniques such as Dropout (Srivastava et al., 2014; Gal & Ghahramani, 2016) and Gaussian Dropout (Kingma et al., 2015), **weight** noises such as DropConnect (Wan et al., 2013) or additive Gaussian noise (Blundell et al., 2015), **target** methods such as label smoothing (Szegedy et al., 2016), **input-target** strategies exemplified by MixUp (Zhang et al., 2018), **input** modifications such as AugMix (Hendrycks et al., 2020) or the standard horizontal flipping and center cropping (Krizhevsky et al., 2009), **model** approaches including weight perturbation (Ash & Adams, 2020), and **gradient** perturbations involving Gaussian noise (Neelakantan et al., 2017). Despite the diversity of these techniques, comprehensive and fair comparisons are scarce, leaving a gap in understanding which approach is helpful for specific datasets, tasks and models. This study aspires to bridge this gap by presenting:

1. The first systematic empirical investigation into the impact of noise injections on NN generalisation and calibration across diverse datasets, tasks and NN architectures. Our exploration extends to evaluation

under in-distribution (ID) and out-of-distribution (OOD) scenarios and their transferability across architectures and datasets.
2. A methodological framework for simultaneously combining various noise injection approaches.
3. Visualisation of the learnt representation landscape across noises, jointly comparing calibration and generalisation performance.

The findings show that AugMix, weak augmentation, and Gaussian noise prove effective across diverse tasks, emphasising their versatility. Task-specific nuances in noise effectiveness, such as AugMix's superiority in computer vision (CV) and Output Diversified Sampling (ODS) in natural language processing (NLP), highlight the need for tailored approaches. Combining noises, careful hyperparameter tuning, and task-specific considerations are crucial for optimising NN's performance.

2 Related Work

The concept of noise injections in this study refers to the deliberate introduction of perturbations into different aspects of NN training – including **input data**, **targets**, **activations**, **weights**, **gradients**, and **model** parameters. These perturbations aim to enhance the generalisation performance of NNs without presupposing any particular data or model characteristics, focusing solely on the underlying task, be it classification or regression. Under these conditions, we review several different noise injection strategies.

Input Noise: Pioneering work by Sietsma & Dow (1991) demonstrated the benefits of training with added input Gaussian noise, while Bishop (1995) established its linkage to regularisation in the least squares problems. Variants of MixUp have exhibited efficacy in augmenting both generalisation and calibration (Zhang et al., 2018; Müller et al., 2019; Guo et al., 2019; Yao et al., 2022; Guo et al., 2017). MixUp linearly blends two samples and their label classification, with CMixUp expanding this approach to regression problems. AugMix, domain-specific to CV, extends the concept by applying a sequence of image processing operations to the input, bolstering robustness in OOD settings. From the adversarial robustness domain, ODS augments inputs with random noise to diversify the inputs (Tashiro et al., 2020). **Target Noise:** Employed extensively, label smoothing (Pereyra et al., 2017), MixUp and CMixUp through target interpolation (Zhang et al., 2018; Yao et al., 2022) emerge as key target noise strategies. Label smoothing replaces one-hot targets with softened counterparts, effectively improving NN's robustness in classification (Müller et al., 2019). **Activation Noise:** Widespread activation noise includes Dropout or Gaussian noise injections. Dropout (Srivastava et al., 2014; Noh et al., 2017) randomly deactivates activations through randomly sampled 0-1 noise, while Gaussian noise injections add noise to activations (Kingma et al., 2015; DeVries & Taylor, 2017). Bayesian NNs (Gal & Ghahramani, 2016) incorporate these injections during training and evaluation, in contrast to our work's focus solely on their application in training. **Weight Noise:** Unlike Dropout, DropConnect (Wan et al., 2013) randomly deactivates weights or connections between neurons, while Gaussian noise injections add noise to weights (Blundell et al., 2015). Note that we do not model the variance of the Gaussian noise through learnable parameters, as in (Blundell et al., 2015), but rather fix it through a searchable hyperparameter. We do this to ensure a fair comparison with other noise injection approaches, such as Dropout, which do not have learnable parameters and would require changing the model architecture to accommodate them. **Gradient Noise:** Annealed Gaussian noise added to gradients during training has demonstrated its efficacy in improving NN generalisation Neelakantan et al. (2017); Welling & Teh (2011); Zhou et al. (2019); Chaudhari & Soatto (2015); Wu et al. (2020). **Model Noise:** A recent contribution, Gaussian noise injection through periodic weight shrinking and perturbation Ash & Adams (2020), improves retraining generalisation.

In previous work, the impact of noise per injection type was studied. Poole et al. (2014) show that injecting noise at different layers of autoencoders implements various regularisation techniques and can improve feature learning and classification performance. Cohen et al. (2019) show that smoothing classifiers with Gaussian noise naturally induces robustness in the L2 norm. Wei et al. (2020) disentangle and analytically characterise the explicit regularisation effect from modifying the expected training objective and the implicit regularisation effect from the stochasticity of Dropout noise in NNs. Camuto (2021); Camuto et al. (2020) show that training NNs with Gaussian noise injections on inputs and activations regularises them to learn lower frequency functions, improves generalisation and calibration on unseen data but also confers robustness to perturbation. Jang et al. (2021) show that training NNs on noisy images can improve their robustness and match human behavioural and neural responses. Lastly, Kukačka et al. (2017) provided a taxonomy of regularisation in NNs, covering multiple noise-based approaches. Past work has studied noise injection techniques in isolation, mainly focused on generalisation alone, lacked comprehensive hyperparameter optimisation, and rarely evaluated the robustness of distribution shift. For example, only MixUp, AugMix and label smoothing have been studied in terms of calibration (Guo et al., 2017; Müller et al., 2019; Guo et al., 2019; Yao et al., 2022).

While promising, these methods require further unified analysis to determine their relationships across datasets, tasks and architectures. Our work addresses these gaps by *1.)* studying the impact across datasets, tasks and architectures; *2.)* benchmarking the impact of noise injections' hyperparameters on transferability between datasets and architectures; *3.)* studying confidence-calibration in addition to generalisation; *4.)* performing a comprehensive hyperparameter search with fair comparisons; *5.)* evaluating robustness to distribution shift; *6.)* providing a methodological framework for combining and tuning various noise injection approaches across categories; and lastly *7.)* visualising the learnt representation or learning landscape across noise injections in 1D or 2D (Goodfellow et al., 2014; Li et al., 2018) across both generalisation and calibration.

3 Methodology

We establish a structured methodology to investigate noise injections' effects on NNs. The noise types are divided into **input**, **input-target**, **target**, **activation**, **weight**, **gradient** and **model**, and we enable their conditional deployment through probabilities $\{p^i_{noise}\}^S_{i=1}$ in the range $0 \leq p^i_{noise} \leq 1$, where S denotes the number of noises. The training allows simultaneous consideration of S noise types, each associated with specific hyperparameters $\{\delta^i\}^S_{i=1}$. The hyperparameter ranges for all noise types are outlined in the Appendix. The probabilities $\{p^i_{noise}\}^S_{i=1}$ allow us to tune the frequency of applying each noise type, while the hyperparameters $\{\delta^i\}^S_{i=1}$ will enable us to adjust the magnitude of each noise type. This enables us to tune both the magnitude and frequency of noise injections, unlike, for example, Dropout (Srivastava et al., 2014), which only allows the tuning of the magnitude, and it is applied every batch.

Algorithm 1 provides a comprehensive overview of the training process, executed throughout E epochs with L batches processed per epoch. For every batch, input and target data (x_b, y_b) are randomly drawn from the training dataset $\mathcal{D} = \{(x_b, y_b)\}^L_{b=1}$. For each noise in S, we sample a uniform random variable $\epsilon \sim U(0, 1)$, and if $\epsilon < p^i_{noise}$, we enable noise i with hyperparameters δ^i for the current batch b. The enabled noises

Algorithm 1 Training of a Neural Network with Noise

Require: Training dataset $\mathcal{D} = \{(x_b, y_b)\}^L_{b=1}$, L batches, number of epochs E, network depth D, weights $W = \{W^d\}^D_{d=1}$, hidden states z^d_b, activations $\{\phi^d(\cdot)\}^D_{d=1}$, weighted operations $\{f^d(\cdot, W^d)\}^D_{d=1}$, S noise types, probabilities of applying noise to a batch $\{p^i_{noise}\}^S_{i=1}$, Noise hyperparameters (HPs) $\{\delta^i\}^S_{i=1}$.

Initialise W randomly
for $e = 1$ to E **do**
  for $b = 1$ to L **do**
    Randomly select a batch (x_b, y_b) from $\mathcal{D}$
    for $i = 1$ to S **do**
      Sample $\epsilon \sim U(0, 1)$
      if $\epsilon < p^i_{noise}$ **then**
        Enable noise i with hyperparameters δ^i for b
      end if
    end for
    Input noise: Modify x_b
    $z^0_b = x_b$
    Target noise: Modify y_b
    for $d = 1$ to D **do**
      Weight noise: Modify W^d
      Compute hidden state $z^d_b = f^d(h^{d-1}_b, W^d)$
      Activation noise: Modify z^d_b; if $d < D$
      $z^d_b = \phi^d(z^d_b)$
    end for
    Assign predictions $\hat{y}_b = z^d_b$
    Compute loss $\mathcal{L}(\hat{y}^i, y^i)$ and gradients $\nabla_W \mathcal{L}$
    Gradient noise: Modify $\nabla_W L$
    Update weights W
  end for
  if $t \mod \text{frequency} = 0$ **and** $e < 0.75E$ **then**
    Model noise: Modify W
  end if
end for

are applied in the order: *1.)* input, input-target, *2.)* target, *3.)* weights, *4.)* activations, *5.)* gradients and *6.)* model, denoted by green lines. The green lines are implemented as for-loops and if-statements, which iterate over the noise types and apply the noise if the noise is enabled. If multiple noises are in the same category, the user specifies their order. For example, suppose MixUp, label smoothing and Dropout are all enabled. In that case, the MixUp is applied to the input-target followed by label smoothing applied to the target, and Dropout is applied to the activations for the current batch b. Our approach accounts for networks of depth D, denoted by $\{f^d(\cdot, W^d)\}_{d=1}^D$, involving weights together with biases $W = \{W^d\}_{d=1}^D$ and activations $\{\phi^d(\cdot)\}_{d=1}^D$ to produce hidden states $\{z_b^d\}_{d=1}^D$. z_0^b corresponds to the input x_b, while z_b^D represents the output prediction $\hat{y}_b$.

For **input** noise, we explore AugMix, ODS, weak augmentation: random cropping and horizontal flipping, and additive Gaussian noise injections (Hendrycks et al., 2020; Tashiro et al., 2020; Sietsma & Dow, 1991). For **input-target** we explore MixUp and CMixUp (Zhang et al., 2018; Yao et al., 2022). For **target** noise, we consider label smoothing, and the target noise also inherently involves MixUp and CMixUp (Zhang et al., 2018; Yao et al., 2022; Müller et al., 2019). The **activation** noise examines Dropout and additive Gaussian noise (Srivastava et al., 2014; Kingma et al., 2015) prior to activations for all linear or convolutional layers, except the last layer. For **weight** noise, we consider Gaussian noise added to the weights (Blundell et al., 2015) or DropConnect (Wan et al., 2013) for all linear or convolutional layers, except the last layer. We consider **gradient** Gaussian noise added to all gradients of the loss function (Neelakantan et al., 2017). After the update of the weights, the **model** noise is applied to the weights, for which we consider shrinking the weights and adding Gaussian noise (Ash & Adams, 2020), but not in the last 25% of the training epochs. Out of these noises, label smoothing, MixUp and ODS are exclusive to classification, and CMixUp is applicable only in regression. AugMix and weak augmentation are exclusive to the CV data. The other noises are broadly applicable across tasks.

4 Experiments

Next, we present the concrete datasets, tasks and architectures used in our experiments, followed by experiments on ID data in Section 4.1, OOD data in Section 4.2, combined noises in Section 4.3, transferability in Section 4.4 and lastly the metric landscape visualisations in Section 4.5.

Tasks, Architectures and Datasets: We consider various setups, including computer vision (CV) classification and regression, tabular data classification and regression, and natural language processing (NLP) classification. For CV classification we include datasets such as CIFAR-10, CIFAR-100 (Krizhevsky et al., 2009), SVHN (Netzer et al., 2011), and TinyImageNet (Le & Yang, 2015), along with neural architectures such as a fully-connected (FC) net and ResNet (He et al., 2016). For CV regression, we introduce a rotated version of CIFAR-100 to predict the rotation angle, and we also use the WikiFace dataset, where the aim is to predict the age based on the image of the face. We use the ResNet model in both cases. In the realm of tabular data classification and regression, we use an FC network and evaluate noises on diverse datasets, including Wine, Toxicity, Abalone, Students, Adult for classification and Concrete, Energy, Boston, Wine, Yacht for regression (Asuncion & Newman, 2007). We explore NLP classification using the NewsGroup dataset (Lang, 1995) paired with global pooling convolutional NN (Kim, 2014) and a transformer (Vaswani et al., 2017). The Appendix details the datasets, architectures, and gives the complete numerical results.

Metrics: To assess the effectiveness of the noise injection methods in classification, we measure their performance using three metrics: Error ($\downarrow$, %), Expected Calibration Error (ECE) (Guo et al., 2017) ($\downarrow$, %) with 10 bins and the categorical Negative Log-Likelihood (NLL) ($\downarrow$). For regression, we use the Mean Squared Error (MSE) ($\downarrow$) and the Gaussian NLL ($\downarrow$). We test the generalisation of the models by evaluating their performance on the ID test set. For CV classification and regression, we test the robustness of the models by assessing their performance on an OOD test set by applying corruptions (Hendrycks & Dietterich, 2019) to the ID test set. These corruptions include, for example, adding snow or fog to the image, changing the brightness or saturation of the image or blurring the image across 5 intensities. We created the OOD test set for tabular data by adding or multiplying the inputs with Gaussian or Uniform noise or by zeroing some of the input features with Bernoulli noise, similarly across 5 intensities. To summarise the results, we collect the results for each approach for each dataset and metric and rank them relative to the no noise baseline.

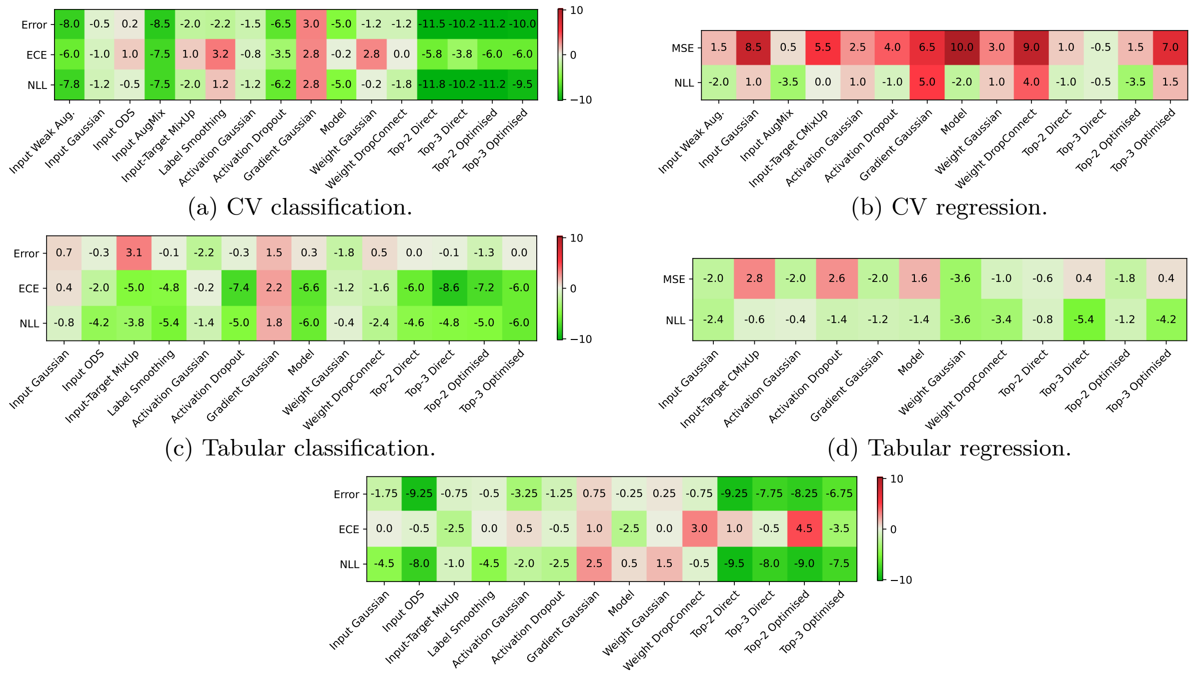

(a) CV classification.

(b) CV regression.

(c) Tabular classification.

(d) Tabular regression.

(e) NLP classification.

Figure 1: In-domain evaluation of the differences in rankings compared to not using any noise.

For example, -1 means that the approach is one rank better than the no noise baseline, and 1 means that the approach is one rank worse than the no noise baseline. We then average the ranks across the datasets for each task and metric.

Hyperparameter Optimisation: 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. By tuning the learning rate and L2 regularisation, we wanted to simulate a realistic scenario where the practitioner seeks to add noise to their existing model and does not want to jointly tune the model's hyperparameters and the noise injection method. The tuning was performed with 1 seed, and the winning hyperparameters were retrained 3 times with different seeds. 10% of the training data was used as the validation set to select the best model, with validation NLL used as the selection objective to combine both generalisation and calibration. The tuning is performed using model-based Tree-structured Parzen Estimator method (Bergstra et al., 2011) with successive halving pruning strategy (Jamieson & Talwalkar, 2016). We evaluate 50 trials for each setting, which allows us to manage the trade-off between compute costs and a reasonable number of trials.

4.1 In-Domain Evaluation

In Figure 1, we show the in-domain (ID) performance of NNs trained with various noise injection methods across CV classification and regression, tabular data classification and regression, and NLP classification. Overall, we observe that the noise injection methods significantly improve the generalisation and calibration in many cases, but different noise types are needed for various tasks. In CV classification, almost all noises improve the error rate, with many simultaneously improving calibration. The most beneficial noises are AugMix, weak augmentation and Dropout. MixUp and label smoothing are a surprise to a certain extent as they improved generalisation but not calibration. In CV regression, improving generalisation was challenging, with no improvement. However, NLL has been improved by several noises, with the best balance given by AugMix, weak augmentation and Dropout. These results suggest that image augmentation broadly benefits CV, confirming expectations.

In tabular data classification, several noises have improved the error rate to a smaller extent or kept it at a similar level. In contrast, almost all noises have improved ECE and NLL. The improvements were particularly

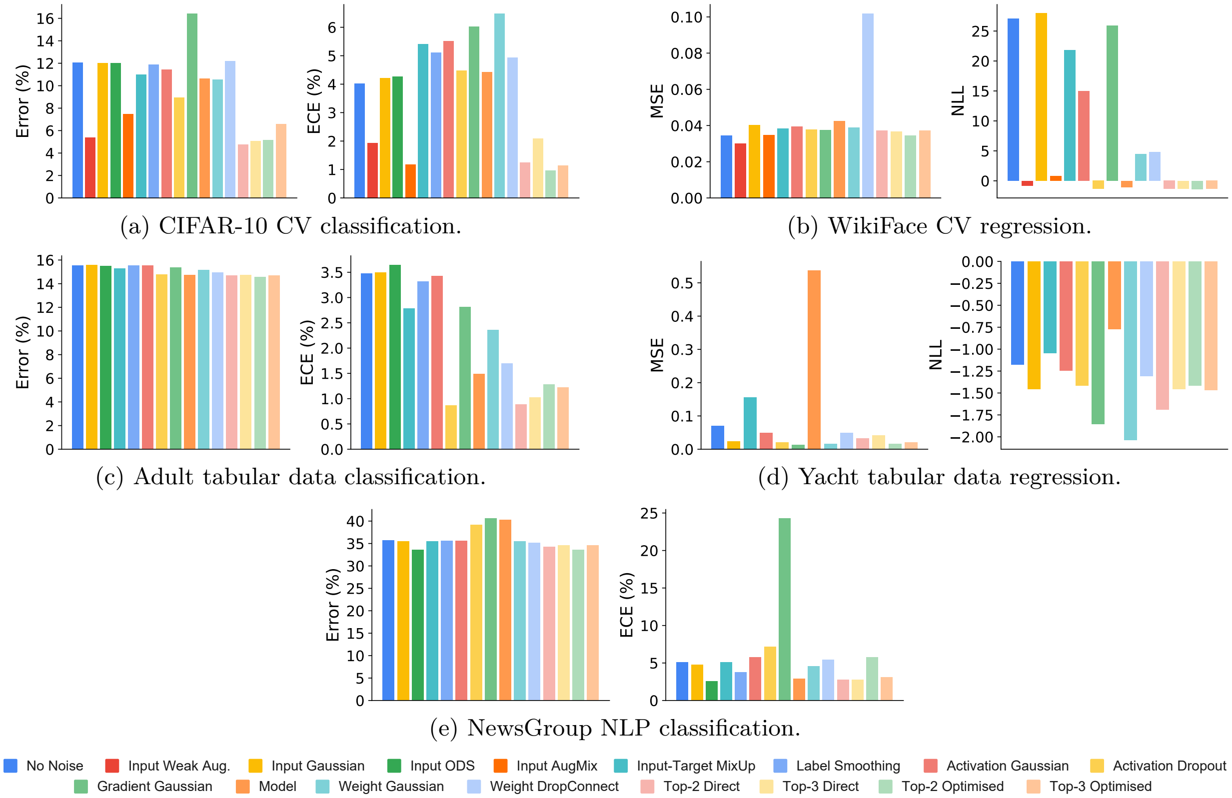

(a) CIFAR-10 CV classification.

(b) WikiFace CV regression.

(c) Adult tabular data classification.

(d) Yacht tabular data regression.

(e) NewsGroup NLP classification.

Figure 2: Detailed in-domain performance of NNs trained with various noises across the five tasks.

impactful in several cases, with model noise, Dropout, and label smoothing being the best. While ODS is designed to improve adversarial robustness, it improved ECE and NLL while keeping similar error rates. All noises improve NLL for tabular regression, and some significantly improve MSE. Gaussian noise applied to the weights or the inputs and DropConnect are the most useful types of noise and improve both metrics. In NLP classification, almost all noises improve error, with some also improving calibration simultaneously. The best noises are ODS, Gaussian added to the input, and label smoothing, which differs from what was best for CV. These noises significantly lowered error and NLL, while MixUp and model noise were particularly useful for reducing ECE. ODS was beneficial for improving error and marginally calibration, which can be a surprise as this technique was not previously considered for improving generalisation or calibration.

In Figure 2, we show detailed results for selecting representative datasets across the 5 tasks. We see the improvements in error can be large for CIFAR-10, for example, halving it in some of the best cases – weak augmentation and AugMix, with Dropout also leading to a few percentage point improvements. The situation is similar for ECE, where weak augmentation and AugMix make the ECE one-half or one-third. Many errors are slightly better, with certain noises making the calibration worse, e.g. MixUp, label smoothing or Gaussian noise added to the activations. For WikiFace, there are more minor improvements in error from weak augmentation and AugMix with overall similar MSE across different noises. Still, the differences in calibration as measured using NLL can be considerable, with most noises improving the NLL significantly.

Moving the focus to tabular data, all noises in the Adult classification dataset improve the error marginally. In contrast, many improve ECE significantly, with the best ones being Dropout, model noise and DropConnect. Most noises have significantly improved MSE for the Yacht regression dataset, but CMixUp and model noise led to significant increases. The best ones have been gradient Gaussian and Gaussian noise added to the weights. NLL has been improved in several cases, including gradient Gaussian and weight Gaussian, demonstrating solid MSE improvements. The errors stay similar for NLP classification on NewsGroup using the global pooling CNN model. ODS leads to the best improvement and a few generalising worse, specifically, Dropout, gradient Gaussian and model noise. ODS and label smoothing have improved ECE.

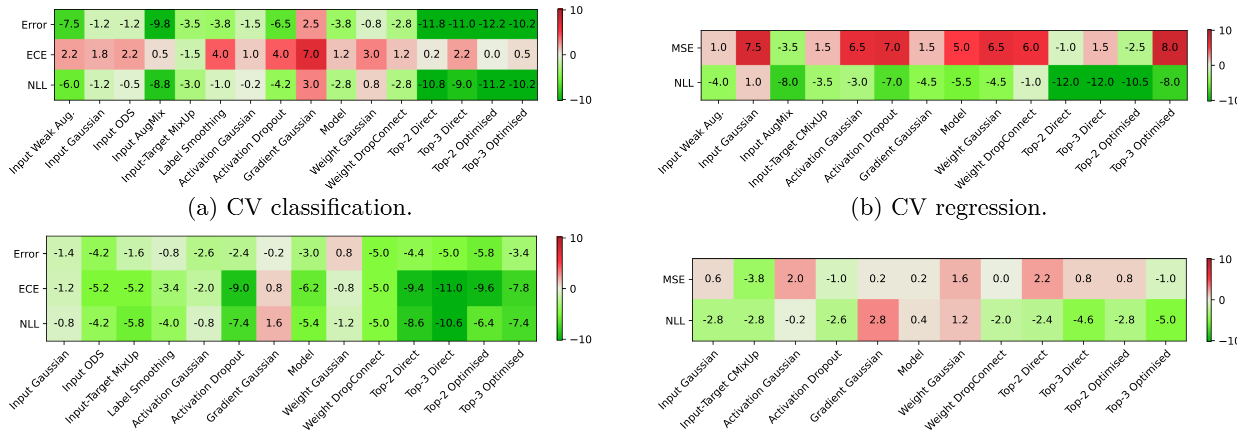

(a) CV classification. (b) CV regression.

(c) Tabular classification. (d) Tabular regression.

Figure 3: OOD evaluation of the differences in rankings compared to not using any noise.

Metric	SVHN	CIFAR-10	CIFAR-100	TinyImageNet	Average
Error	0.912 ± 0.000	0.706 ± 0.000	0.897 ± 0.000	0.912 ± 0.000	0.857
ECE	1.000 ± 0.000	0.618 ± 0.000	0.250 ± 0.177	0.118 ± 0.542	0.496
NLL	0.971 ± 0.000	0.515 ± 0.003	0.868 ± 0.000	0.882 ± 0.000	0.809

(a) CV classification.

Metric	Rotated CIFAR-100	WikiFace	Average
MSE	0.257 ± 0.202	0.581 ± 0.002	0.419
NLL	-0.238 ± 0.239	0.810 ± 0.000	0.286

(b) CV regression.

Metric	Wine	Toxicity	Abalone	Students	Adult	Average
Error	-0.211 ± 0.276	0.761 ± 0.000	0.126 ± 0.518	0.500 ± 0.011	0.867 ± 0.000	0.409
ECE	0.600 ± 0.001	0.905 ± 0.000	0.314 ± 0.114	0.524 ± 0.006	0.905 ± 0.000	0.650
NLL	0.467 ± 0.016	0.981 ± 0.000	0.657 ± 0.000	0.029 ± 0.923	0.867 ± 0.000	0.600

(c) Tabular data classification.

Metric	Energy	Boston	Wine	Yacht	Concrete	Average
MSE	0.000 ± 1.000	-0.077 ± 0.765	-0.513 ± 0.015	-0.282 ± 0.204	-0.128 ± 0.590	-0.200
NLL	-0.333 ± 0.129	0.821 ± 0.000	0.154 ± 0.510	0.128 ± 0.590	0.590 ± 0.004	0.272

(d) Tabular data regression.

Table 1: Kendall Tau correlation between ID and OOD rankings of different noise types for various tasks.

Main Observations: The noises are effective across various tasks and datasets. The shortlist of the most effective methods is AugMix and weak augmentation in CV, model noise, and Gaussian noise added to weights for tabular data and ODS in NLP. Different task types benefit from different types of noise.

4.2 Out-of-Domain Evaluation

We evaluate the performance on the ID test set and an augmented OOD set, including an average over visual corruptions across 19 categories and 5 severities (Hendrycks & Dietterich, 2019). Likewise, we average the performance across 5 categories and 5 severities for tabular data. The summary of the results is in Figure 3, with analysis of correlations between ID and OOD rankings via Kendall Tau score in Table 1. For CV classification, we observe that the generalisation improvements also remain for OOD, but calibration in terms of ECE turns out to be much more challenging to improve. The overall ranking of the best noises remains similar, with AugMix and weak augmentation remaining the best. MixUp rose to prominence thanks to the best OOD calibration and improved errors and NLL. Analysis of Kendall Tau correlation in Table 1a shows that ID and OOD rankings are strongly correlated for error and NLL, while only moderately for ECE. CV regression is similar to classification ranking the best noises, with only AugMix leading to improvements in OOD generalisation. However, calibration is improved by most noises, with AugMix excelling. Only a minor correlation exists between ID and OOD rankings for MSE and NLL metrics. For tabular classification, the noises generally improve all metrics under OOD settings, with Dropout, DropConnect, and model noise being the best. This suggests that model noise and Dropout are among the best noises for both ID and OOD. ID and OOD rankings show a moderate correlation overall. Several noises improve OOD generalisation and calibration for tabular regression, with CMixUp, Dropout and input, and Gaussian added to the activations noises, leading to the best overall improvements. The ID and OOD ranking Kendall Tau correlation is low or negative in this case.

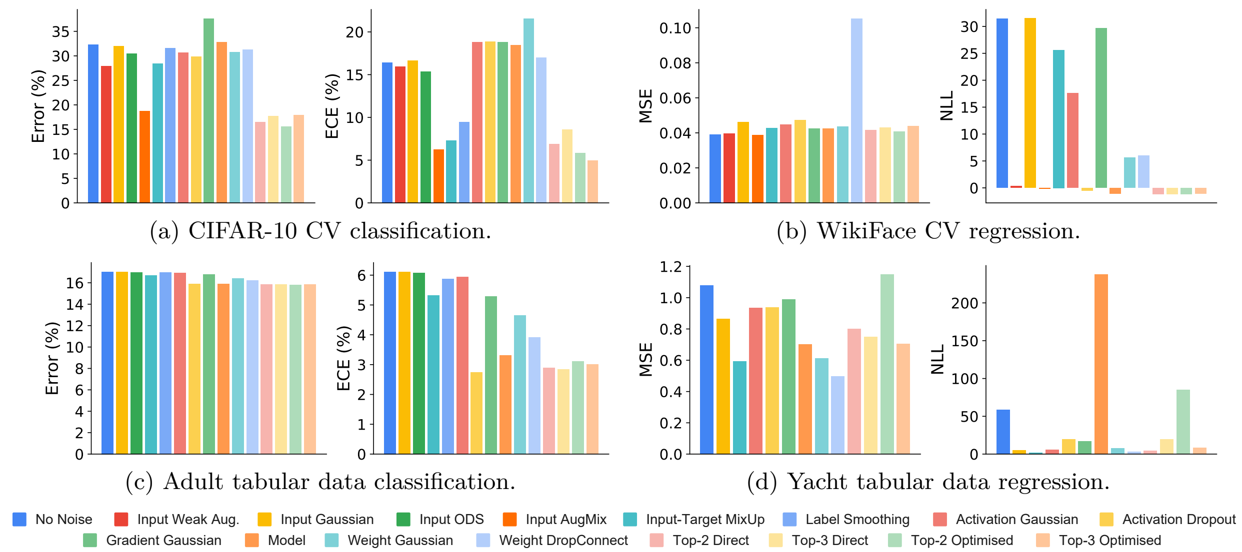

Figure 4: Detailed OOD performance of NNs trained with various noises across the four tasks.

We study selected representative datasets regarding OOD performance in Figure 4. OOD results on CIFAR-10 show that AugMix significantly improves both error and ECE, making ECE one-third of the no noise equivalent. MixUp leads to similarly considerable improvements in ECE and more minor yet significant improvements in error. Several noises, e.g., Dropout and Gaussian noise, added to activations or weights lead to a few percentages worse ECE. On WikiFace, most OOD MSE values are similar, but OOD calibration in NLL is improved significantly for several noises, including AugMix, weak augmentation or Dropout. Improvements in generalisation for the Adult tabular classification dataset are minor, but the improvements in calibration can be significant, for example, Dropout and model noise halving the OOD ECE value. For the Yacht tabular regression dataset, the improvements in generalisation have been more critical, with the same being true for calibration measured in terms of OOD NLL.

Main Observations: We see consistent improvements in OOD generalisation and calibration for tabular data. For CV classification, errors and NLL are improved, but calibration is generally not improved when measured via ECE. CV regression sees improvements in OOD NLL only. The best ID noise types have often remained the best OOD, but overall, the correlations between ID and OOD rankings were lower, especially for tabular data. MixUp, or CMixUp for regression, showed surprising behaviour as it was much more helpful for improving OOD calibration than ID calibration.

4.3 Combination of Noises

Next, we evaluate the combination of noises. We construct them from empirical combinations of the Top 2 or 3 noises from the ID evaluation for each task, based on average rank across respective datasets and metrics. We consider two cases: *1.)* the found hyperparameters of the noises are directly applied, and *2.)* the hyperparameters of the noises are jointly tuned. We utilise the same 50-trial budget to tune the selected noises jointly. The results are already in Figures 1, 2, 3 and 4 and denoted as Top-2 Direct, Top-3 Direct for *1.)*, Top-2 Optimised and Top-3 Optimised for *2.)*. The combinations for Top-2 and Top-3 are in Table 2. To simplify the analysis of how effective the different combinations of noises are, we compute their average rank improvement compared to no noise and report it in Table 3. Notice that when we choose a combination of noises to involve noises from the same category, for example, ODS and input Gaussian are both input noises, these are applied sequentially.

We can draw several observations from Table 3. *1.)* Given our budget for optimising hyperparameters, combining two noises and optimising their hyperparameters was helpful, rather than directly combining their hyperparameters. However, the budget was insufficient for combining three noises, and the found hyperparameters did not perform as well as using fewer noises or directly combining the hyperparameters.

Task	Top-2	Third Method
CV classification	Input AugMix, Input Weak Augmentation	Activation Dropout
NLP classification	Input ODS, Input Gaussian	Target Label Smoothing
Tabular classification	Model, Activation Dropout	Target Label Smoothing
CV regression	Input AugMix, Input Weak Augmentation	Activation Dropout
Tabular regression	Weight Gaussian, Input Gaussian	Weight DropConnect

Table 2: Top task and noise combinations. Underlined methods are from the same type.

Scenario	Top-2 Direct	Top-3 Direct	Top-2 Optimised	Top-3 Optimised
ID	-4.48	-4.61	-4.57	-3.89
OOD	-5.84	-5.87	-6.08	-4.45

Table 3: Average rank improvement over no noise for the different combination strategies.

2.) The combination of three noises performs better if we directly combine their hyperparameters but not when we optimise their hyperparameters, as a significantly larger budget would be needed. *3.)* Considering the relative differences, directly combining the top three noises is a reasonable strategy.

Commenting on the overall performance of the combinations of noises, the combinations are typically better for classification tasks than the individual noises. Still, the opposite may be true for regression. As observed in Figures 1a, 1c and 1e, the combinations are consistently ranked lowest in comparison to using no noise for classification, showing the effectiveness of the combinations. However, Figures 1b and 1d show that regression can benefit from only using one noise at a time. OOD analysis in Figure 3 confirms the benefits of combinations of noises for classification tasks, and it also shows that it can be beneficial for regression, contrary to the ID behaviour. The combinations are generally ranked lower and can improve calibration and generalisation, as seen in lower MSE, NLL, or error and ECE simultaneously.

Main Observations: In general, combining noises is better than individual noises, directly using 3 noises is better than 2 noises, and directly combining noises with their hyperparameters is reasonable. The combination of noises can improve both calibration and generalisation simultaneously.

4.4 Transferability of Hyperparameters Across Datasets and Models

Furthermore, we evaluate the transferability of the hyperparameters across datasets and models. We consider two cases: the transfer of hyperparameters to a new dataset and the transfer of hyperparameters to a new architecture. For the dataset transfer, we consider the following combinations: SVHN to CIFAR-10, CIFAR-10 to CIFAR-100, CIFAR-100 to TinyImageNet, and 3 tabular regression datasets combinations, Concrete to Energy, Boston to Wine, Yacht to Concrete. We consider the following combinations for the architecture transfer: FC to ResNet-18 for SVHN and ResNet-18 to ResNet-34 for CIFAR-10, CIFAR-100 and TinyImageNet. We use a NN with an additional layer for tabular data, i.e., five layers instead of four.

4.4.1 Dataset Transfer

Figures 5a and 5c show the dataset transfer results for ID settings, with OOD settings shown in Figures 8a and 8c in the Appendix. We observe generally good transferability of hyperparameters across datasets for CV classification in ID and OOD settings. In particular, weak augmentation, AugMix and Dropout lead to solid improvements in the ID setting. AugMix also excels in OOD scenarios under dataset transfer, but weak augmentation and Dropout are not as strong in calibration measured using ECE. Certain noises are less transferable, including Gaussian noise added to the input and DropConnect. Hyperparameters for noise in tabular regression are less transferable because of worse generalisation measured using MSE.

Main Observations: The transfer of hyperparameters from dataset to dataset generally works well for CV classification. However, caution is advised as it is not the case for all noise types. For tabular data regression, tuning of hyperparameters is recommended.

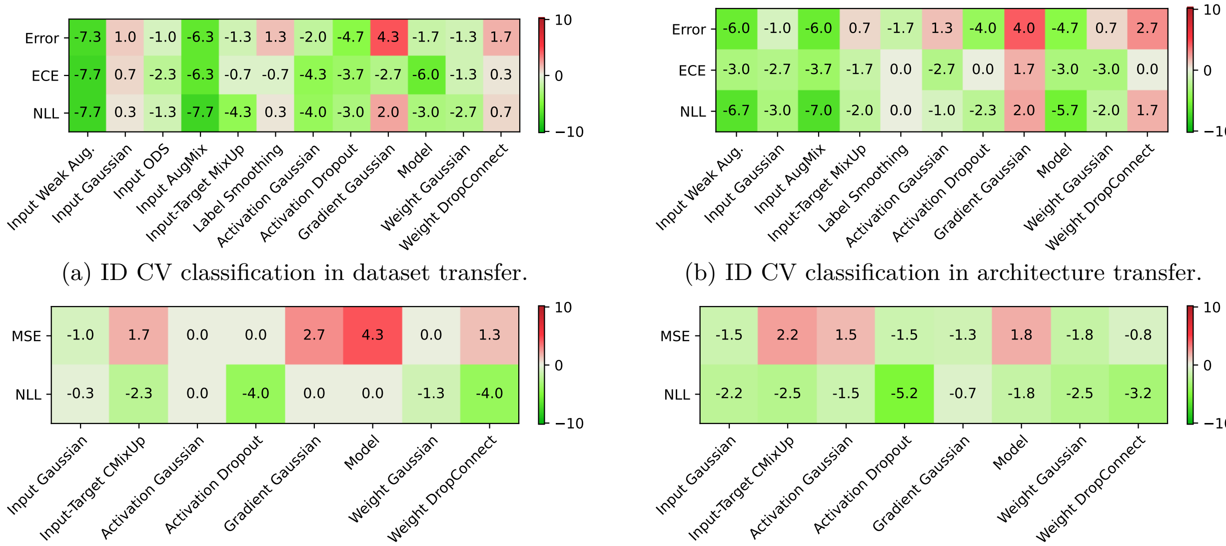

(a) ID CV classification in dataset transfer.

(b) ID CV classification in architecture transfer.

(c) ID tabular regression in dataset transfer.

(d) ID tabular regression in architecture transfer.

Figure 5: Transfer of hyperparameters on in-domain (ID) data.

4.4.2 Architecture Transfer

In Figures 5b and 5d, we show the ID results for the architecture transfer, with Figures 8b and 8d in the Appendix reporting the OOD results. The transferability of noise hyperparameters is lower than across datasets for CV classification, but it is still successful, especially for weak augmentation and AugMix for ID settings. Transfer of hyperparameters for tabular data regression works for certain noise types in the ID setting, including adding Gaussian noise to the input or the weights, which are the top-2 noises for tabular data regression.

Main Observations: Transfer of hyperparameters across architectures appears more challenging than across datasets but can be successful in some instances. Caution is advised, and tuning is recommended.

4.5 Learnt Representation Landscapes

We study the learnt representation landscapes of NNs trained with various noises through the lenses of ID and OOD performance in terms of error, ECE, NLL or MSE. We consider the noises individually, with the ID-found hyperparameters starting from the same weight initialisation for fairness. We visualise linear interpolation modulated through an α parameter between the final, $\alpha = 0$ and initial model, $\alpha = 1$ (Goodfellow et al., 2014). The interpolation empirically investigates the smoothness of the training process. We also visualise the landscape in 2D (Li et al., 2018) by saving the network after each epoch and concatenating the weights. Instead of using random coordinates, we use the first two principal components of the weights as the coordinates. We normalise them based on the magnitude of the original weights, and we project all the weights onto these two components in the vicinity of α and β. The 2D visualisations show us the exploration and exploitation of the training process. In Figures 6 and 7, we compare the metric landscapes of no noise with AugMix and Dropout noises, respectively on CIFAR-10 and WikiFace datasets. We used 20 points for linear interpolation and 100 points for the 2D plots across five selected OOD augmentations and 1000 test data samples for compute efficiency. In red, we show the error or MSE; in green, we offer the NLL or ECE. In the 1D plots, ● and ▲ stand for ID and OOD error or MSE, and × and ■ stand for ID and OOD ECE or NLL. In the 2D plots, the darker combined contours signify worse performance than the lighter parts, and the ★ in blue or black denotes the start or end weights respectively. The Appendix contains the metric landscapes for all other noises, tabular classification – Adult, and regression – Yacht datasets.

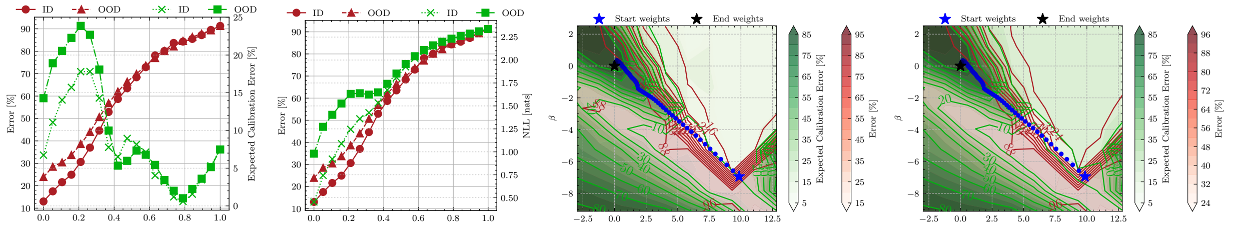

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

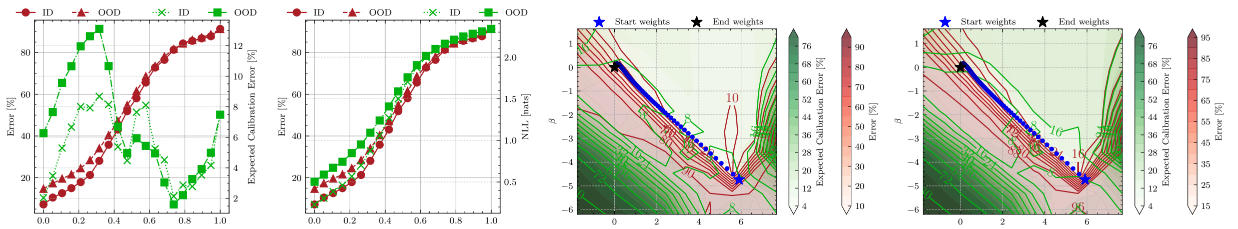

(e) Error and ECE. (f) Error and NLL. (g) Error and ECE on ID. (h) Error and ECE on OOD.

Figure 6: No noise (top) and Input AugMix (bottom) on CIFAR-10.

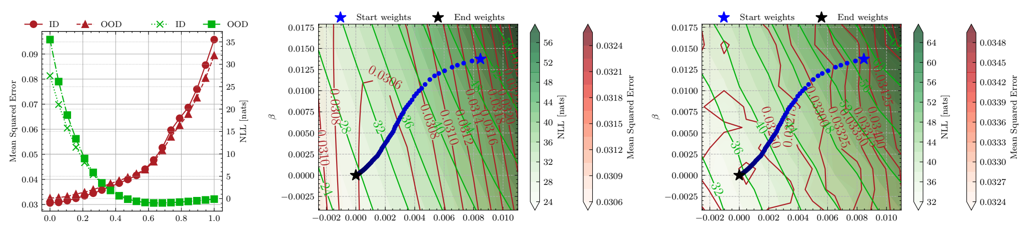

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

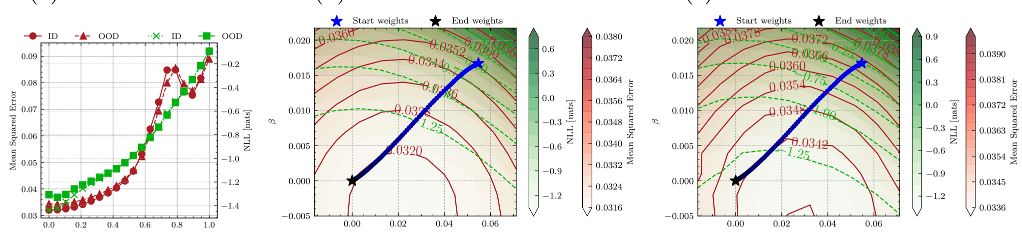

(d) MSE and NLL. (e) MSE and NLL on ID. (f) MSE and NLL on OOD.

Figure 7: No noise (top) and Activation Dropout (bottom) on WikiFace.

Observing Figures 6 and 7, we first notice the ID and OOD results are similar, with the OOD results being slightly worse across all metrics. This includes both the 1D plots and the 2D plots. Second, as seen in Figures 6a and 6b, the curves for error, representing generalisation, and ECE or NLL, representing confidence calibration, do not share the same shapes or curvatures. MSE and NLL curves in Figure 7a are more similar than error and ECE curves. Looking at the 1D plots, for example in Figures 6a 6b and 7a, the error or MSE can be more smoothly interpolated than ECE or NLL. Figure 7a shows models trained without noise can become overconfident, reflected in large NLL and small MSE. Adding noise such as Dropout can fix this, leading to low NLL for the final model in Figure 7d. Looking at the 2D plots in Figures 6c and 7b, the error or MSE valley is wider than the ECE or NLL valley, and they are not aligned. From a detailed comparison between no noise and AugMix or Dropout in Figures 6 and 7, we observe that AugMix and Dropout can smoothen the optimisation in the 1D plots, but not for ECE, and decrease the gap between ID and OOD performance. The 2D plots show that AugMix and Dropout can explore broader metric landscapes than no noise, shown in ranges of α and β in the 2D plots, and marginally align the error or MSE with NLL. Seen in the lightness of the 2D contour plots, the noises navigate lower NLL or ECE landscapes than no noise.

Our general observations considering both CV and tabular datasets show that while noises such as AugMix, weak augmentation, MixUp or activation and weight noises based around Dropout can smoothen the optimisation regarding error or MSE, they rarely smoothen the optimisation regarding ECE. The metric landscapes often look similar to no noise, but the optimisation ends in more profound valleys. Across the datasets and tasks, label smoothing, input additive Gaussian and ODS have minimal effect on the 2D landscapes or 1D interpolation. The model shrink and perturb make the optimisation more "stairs-like", and the metric landscape explored is broader. Together with gradient Gaussian noise, the shrink and perturb noises explore broader metric landscapes than the others. No method drastically changes the metric landscape or the interpolation from the default, but they can make the optimisation smoother or broader.

Main Observations: The metric landscapes for error or MSE and ECE or NLL are different, and the noises can smoothen the optimisation in terms of error or MSE but not necessarily in terms of ECE or NLL. When a model trained without noise is overconfident, adding noise to the training can resolve it and lead to a significantly better-calibrated model at the end of training.

5 Conclusion

Key Takeaways: Noise injection methods can improve NN performance across various tasks and datasets. This is despite the fact L2 regularisation was already tuned to prevent overfitting, indicating noise injection methods can provide additional benefits beyond standard regularisation. The methods did not have the same efficiency across all tasks and datasets, with significant differences in performance between regression and classification. The most effective noise for CV was AugMix, model shrink and perturb and Gaussian noise added to weights for tabular data classification and regression respectively, while ODS worked the best for NLP. Even though ODS was not designed to improve calibration and generalisation, it has shown promising performance in several cases. Combining noises outperformed individual noises in most classification cases, with regression often benefitting from using only one noise at a time. While directly combining hyperparameters of noises is a reasonable strategy, tuning them can still be valuable if a large budget is used. The noises improved both ID and OOD performance, but the ID rankings were sometimes inconsistent with the OOD evaluation. AugMix remained highly ranked for robustness. The visualisation showed noises can smoothen the optimisation in terms of error or MSE but not necessarily in terms of ECE or NLL. It also showed noise can be helpful in mitigating overconfidence. Overall, results indicate practitioners should consider combining noises, e.g. AugMix and dropout, and tuning hyperparameters for their specific problem.

Limitations: To conduct this study, we had to restrict the experiments' scope. Our scope was limited to standard datasets, tasks such as classification and regression, and standard NN architectures. Testing on more complex data and downstream tasks such as object detection, segmentation, or reinforcement learning would reveal more profound impacts of noise injection. Moreover, we also limited the optimisation to SGD with momentum and a cosine learning rate schedule, which were tuned beforehand to make the hyperparameter search tractable. The computational costs of tuning and comparing noises also restricted the scale of experiments. To draw practical conclusions, we evaluated the noise performance from the perspective of minimising the NLL rather than exploring all possible settings. Further work on selecting hyperparameters without exhaustive tuning would make these techniques more accessible.

Future Directions: The strong performance of AugMix highlights the potential for developing specialised, domain-specific noise techniques. For example, tailored domain-specific noise methods could benefit tabular data-based problems and NLP. Future work should also explore specific data-architecture noise interactions, as the transferability of hyperparameters was limited. Inspired by the annealed gradient noise, annealing noise levels over training may also prove helpful, as early noise could encourage robustness. In contrast, low late-stage noise could enable convergence on a high-accuracy solution. The potential for combining noises from the same category should also be investigated further. The noises affected the entire architecture, but it may be possible to target noise injection methods that only affect specific layers or sections of the network, requiring more or less regularisation. Lastly, specific noise-based approaches for simultaneously exploring the generalisation and confidence calibration trade-off should be explored further. We hope our study and framework, embedded in our codebase, will assist further research in this area.

References

Jordan Ash and Ryan P Adams. On warm-starting neural network training. In *NeurIPS*, 2020.

Arthur Asuncion and David Newman. Uci machine learning repository, 2007.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *NeurIPS*, 2011.

Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1): 108–116, 1995.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *ICML*, 2015.

Alexander Camuto. *Understanding Gaussian noise injections in neural networks*. PhD thesis, University of Oxford, 2021.

Alexander Camuto, Matthew Willetts, Umut Simsekli, Stephen J Roberts, and Chris C Holmes. Explicit regularisation in gaussian noise injections. In *NeurIPS*, 2020.

Pratik Chaudhari and Stefano Soatto. On the energy landscape of deep networks. *arXiv*, 2015.

Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *ICML*, 2019.

Terrance DeVries and Graham W Taylor. Dataset augmentation in feature space. *arXiv*, 2017.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.

Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network optimization problems. *arXiv*, 2014.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *ICML*, 2017.

Hongyu Guo, Yongyi Mao, and Richong Zhang. Augmenting data with mixup for sentence classification: An empirical study. *arXiv*, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

Zhezhi He, Adnan Siraj Rakin, and Deliang Fan. Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In *CVPR*, 2019.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019.

Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2020.

Kevin Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. In *AISTATS*, 2016.

Hojin Jang, Devin McCormack, and Frank Tong. Noise-trained deep neural networks effectively predict human vision and its neural responses to challenging images. *PLoS biology*, 19(12):e3001418, 2021.

Yoon Kim. Convolutional neural networks for sentence classification. *arXiv*, 2014.

Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *NeurIPS*, 2015.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Anders Krogh and John Hertz. A simple weight decay can improve generalization. In *NeurIPS*, 1991.

Jan Kukačka, Vladimir Golkov, and Daniel Cremers. Regularization for deep learning: A taxonomy. *arXiv*, 2017.

Ken Lang. Newsweeder: Learning to filter netnews. In *ICML*, 1995.

Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, 2018.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.

Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *NeurIPS*, 2019.

Arvind Neelakantan, Luke Vilnis, Quoc V. Le, Lukasz Kaiser, Karol Kurach, Ilya Sutskever, and James Martens. Adding gradient noise improves learning for very deep networks. In *OpenReview*, 2017.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

Hyeonwoo Noh, Tackgeun You, Jonghwan Mun, and Bohyung Han. Regularizing deep neural networks by noise: Its interpretation and optimization. In *NeurIPS*, 2017.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.

Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. In *ICLR Workshop*, 2017.

Ben Poole, Jascha Sohl-Dickstein, and Surya Ganguli. Analyzing noise in autoencoders and deep networks. *arXiv*, 2014.

Lutz Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the trade*. Springer, 2002.

Jocelyn Sietsma and Robert JF Dow. Creating artificial neural networks that generalize. *Neural networks*, 4(1):67–79, 1991.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 2014.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.

Yusuke Tashiro, Yang Song, and Stefano Ermon. Diversity can be transferred: Output diversification for white-and black-box attacks. In *NeurIPS*, 2020.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *ICML*, 2013.

Colin Wei, Sham Kakade, and Tengyu Ma. The implicit and explicit regularization effects of dropout. In *ICML*, 2020.

Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011.

Jingfeng Wu, Wenqing Hu, Haoyi Xiong, Jun Huan, Vladimir Braverman, and Zhanxing Zhu. On the noisy gradient descent that generalizes as SGD. In *ICML*, 2020.

Huaxiu Yao, Yiping Wang, Linjun Zhang, James Y Zou, and Chelsea Finn. C-mixup: Improving generalization in regression. In *NeurIPS*, 2022.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.

Mo Zhou, Tianyi Liu, Yan Li, Dachao Lin, Enlu Zhou, and Tuo Zhao. Toward understanding the importance of noise in training neural networks. In *ICML*, 2019.

Appendix

In the Appendix, we first provide the experimental settings and the hyperparameter ranges for all the experiments in Section A. We report the OOD transferability of hyperparameters results in Section B. We then provide the full numerical results and visualisations for all the experiments in Section C.

A Settings

A.1 General Settings

We used stochastic gradient descent with a momentum of 0.9 to train all the networks. The learning rate and L2 regularisation were tuned and reused for each noise injection method. We used a cosine annealing learning rate schedule without restarts (Loshchilov & Hutter, 2017) for all experiments. We used gradient norm clipping of 20.0 to stabilise the training in most cases, with gradient clipping of 10.0 for tabular regression and 5.0 for WikiFace. The batch size was set to 256 for all experiments. The final results are the average of 3 runs with 3 different seeds. We used cross-entropy loss for all classification experiments. For regression, we used the Gaussian negative log-likelihood (NLL) loss, where we modelled the variance as an additional output passed through an exponential function to ensure positivity. We added a small ϵ of $1e^{-8}$ to the softmax probabilities to avoid NaNs. We clipped the variance between $1e^{-4}$ and $1e^{4}$ to avoid NaNs. The hyperparameter ranges, and the sampling scale for each dataset-architecture pair are in Table 4. The hyperparameters and implementations of all the noises and experiments can be found in the code, which will be open-sourced. We used the default PyTorch weight initialisation for all layers. For the tabular OOD experiments, we constructed custom augmentations where we applied Gaussian or Uniform noise scaled by the magnitude of the input features across 5 severities for addition: [0.02, 0.04, 0.06, 0.08, 0.1] or multiplication [0.04, 0.08, 0.12, 0.16, 0.2] where the severity scaled the range or the standard deviation of the noise applied to the input. Additionally, we zeroed out some input features with probability [0.04, 0.08, 0.12, 0.16, 0.2], denoting 5 severities. In total, there were 5 different input shifts across 5 severities each.

Hyperparameter (δ)	**Range**	**Scale**
Learning rate (LR)	$[10^{-4}, 10^{-1}]$	Log
L2 weight	$[10^{-7}, 10^{-1}]$	Log
Input Gaussian noise std.	$[10^{-4}, 10^{-1}]$	Log
Input AugMix alpha	$[0, 1]$	Linear
Input AugMix severity	$[1, 10]$	Linear
Input AugMix width	$[1, 5]$	Linear
Input AugMix chain-depth	$[-1, 3]$	Linear
Input ODS epsilon	$[10^{-4}, 10^{-1}]$	Log
Input ODS temperature	$[0.5, 5.0]$	Log
Input-Target MixUp alpha	$[0, 1]$	Linear
Input-Target CMixUp alpha	$[0, 1]$	Linear
Input-Target CMixUp sigma	$[10^{-4}, 10^{2}]$	Log
Target Label Smoothing	$[0, 0.25]$	Linear
Activation Gaussian noise std	$[10^{-4}, 10^{-1}]$	Log
Activation Dropout rate	$[0, 1]$	Linear
Gradient Gaussian noise η	$[0, 1]$	Linear
Gradient Gaussian noise γ	$[0, 1]$	Linear
Weight Gaussian noise std	$[10^{-4}, 10^{-1}]$	Log
Weight DropConnect rate	$[0, 1]$	Linear
Model noise shrink factor	$[0.0, 1.0]$	Linear
Model noise std	$[10^{-7}, 10^{-3}]$	Log
Model noise frequency	$[0, 20]$	Linear

Table 4: Hyperparameters (HPs) optimised for individual noises and their range.

Regarding noise implementation details, Dropout, DropConnect, additive weight or activation Gaussian noise, are applied to all linear and convolutional weights throughout the network, excluding the last layer and normalisation layers. Both model and Gaussian gradient noise are implemented on all weights within the network, encompassing affine parameters in normalisation layers. Rotation was omitted from AugMix, given that one of our tasks involved predicting the rotation angle.

A.2 Vision Experiments

For SVHN, we used a fully connected network with 4 hidden layers of 150 units followed by ReLU activations. When we used ResNet-18 we used it with [64, 128, 256, 512] channels in 4 stages with [2, 2, 2, 2] blocks with strides [1, 2, 2, 2]. When we used ResNet-34, we used it with [64, 128, 256, 512] channels in 4 stages with [3, 4, 6, 3] blocks with strides [1, 2, 2, 2]. In all cases, we trained the networks for 200 epochs. We only used 0-1 truncation followed by normalisation for each dataset without further data augmentations for training, validation and test sets. For rotation experiments, we enabled uniform rotations between (0, 90°) degrees, and we rescaled the targets accordingly to [-1, 1]. The selected OOD augmentations for visualisation experiments were Gaussian noise, motion blur, snow, elastic transformation and JPEG compression across all 5 severities.

A.3 NLP Classification Experiments

The NewsGroup20 was first pre-processed with respect to glove embeddings (Pennington et al., 2014) into embeddings of dimension 100 and sequence length 100. In all cases, we trained the networks for 200 epochs. We used the global-pooling convolutional network architecture from Kim (2014) with planes [128, 128, 128] and a transformer decoder (Vaswani et al., 2017) with embedding dimensions 100, 6 layers, 8 heads, 1024 hidden dimensions, 64 dimensions per head and no dropout. There was no OOD test set for the NLP task.

A.4 Tabular Regression Experiments

We used a fully connected network with [100, 100, 100, 100] hidden units and ReLU activations for the tabular experiments. In all cases, we trained the networks for 100 epochs. We normalised the input features and targets to zero mean and unit variance by using the training set statistics and applied the same normalisation to the validation and test sets. We used 20% of the data as the test set and 10% of the remaining data as the validation set. The regression targets were normalised to zero mean and unit variance.

B Out-of-Domain Transferability of Hyperparameters

We report the OOD results studying transferability of hyperparameters in Figure 8.

C Full Results

We provide full results of all experiments in the paper, where the main reported value is the mean across 3 repetitions, followed by the standard deviation. The ranks presented in the main body of the paper can be obtained by ranking the results in each table by the metric of interest. Following the tables, there are the visualisations of metric landscapes for CIFAR-10, Adult, WikiFace and Yacht datasets. We encourage the reader to look at our code for other datasets to regenerate them from there.

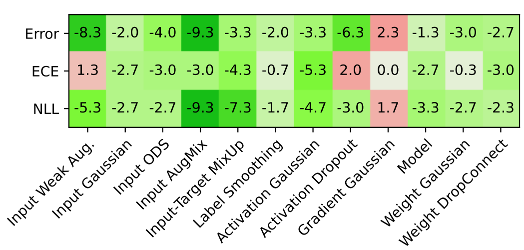

(a) OOD CV classification in dataset transfer.

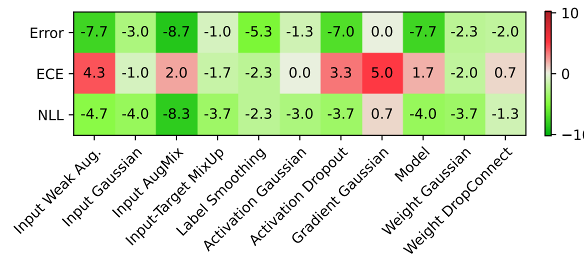

(b) OOD CV classification in architecture transfer.

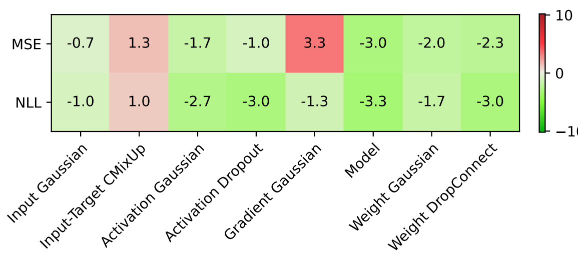

(c) OOD tabular regression in dataset transfer.

(d) OOD tabular regression in architecture transfer.

Figure 8: Transfer of hyperparameters.

Noise Type	**SVHN**		**CIFAR-10**		**CIFAR-100**		**TinyImageNet**	
	ID	**OOD**	**ID**	**OOD**	**ID**	**OOD**	**ID**	**OOD**
No Noise	$16.33_{\pm 0.07}$	$20.18_{\pm 0.09}$	$12.04_{\pm 0.21}$	$32.32_{\pm 0.24}$	$44.69_{\pm 0.84}$	$61.52_{\pm 0.52}$	$54.15_{\pm 0.36}$	$75.65_{\pm 0.19}$
Input Weak Aug.	$14.67_{\pm 0.14}$	$18.59_{\pm 0.14}$	$5.39_{\pm 0.17}$	$27.86_{\pm 0.37}$	$26.51_{\pm 0.10}$	$52.57_{\pm 0.13}$	$39.73_{\pm 0.27}$	$67.73_{\pm 0.13}$
Input Gaussian	$16.45_{\pm 0.23}$	$20.23_{\pm 0.22}$	$12.02_{\pm 0.05}$	$31.97_{\pm 0.31}$	$44.34_{\pm 1.70}$	$61.14_{\pm 1.60}$	$53.28_{\pm 0.24}$	$75.12_{\pm 0.06}$
Input ODS	$16.35_{\pm 0.20}$	$20.13_{\pm 0.15}$	$12.01_{\pm 0.16}$	$30.42_{\pm 0.45}$	$44.28_{\pm 0.64}$	$61.34_{\pm 0.56}$	$66.47_{\pm 15.18}$	$82.46_{\pm 8.72}$
Input AugMix	$12.28_{\pm 0.07}$	$15.67_{\pm 0.07}$	$7.48_{\pm 0.06}$	$18.75_{\pm 0.25}$	$30.09_{\pm 0.25}$	$46.05_{\pm 0.13}$	$42.49_{\pm 0.24}$	$61.58_{\pm 0.16}$
Input-Target MixUp	$13.95_{\pm 0.01}$	$17.71_{\pm 0.11}$	$10.97_{\pm 0.13}$	$28.38_{\pm 0.70}$	$46.15_{\pm 0.05}$	$63.00_{\pm 0.02}$	$54.02_{\pm 0.53}$	$75.33_{\pm 0.44}$
Label Smoothing	$16.35_{\pm 0.17}$	$20.10_{\pm 0.12}$	$11.88_{\pm 0.39}$	$31.63_{\pm 0.29}$	$42.48_{\pm 0.47}$	$58.41_{\pm 0.55}$	$52.47_{\pm 0.21}$	$73.96_{\pm 0.15}$
Activation Gaussian	$16.33_{\pm 0.17}$	$20.14_{\pm 0.13}$	$11.44_{\pm 0.12}$	$30.71_{\pm 0.20}$	$44.58_{\pm 0.28}$	$61.81_{\pm 0.15}$	$53.96_{\pm 0.16}$	$75.45_{\pm 0.14}$
Activation Dropout	$13.83_{\pm 0.13}$	$17.43_{\pm 0.10}$	$8.93_{\pm 0.25}$	$29.85_{\pm 0.86}$	$41.51_{\pm 0.83}$	$58.92_{\pm 0.80}$	$43.26_{\pm 0.35}$	$69.32_{\pm 0.32}$
Gradient Gaussian	$17.59_{\pm 0.15}$	$22.49_{\pm 0.07}$	$16.41_{\pm 0.15}$	$37.57_{\pm 0.12}$	$45.33_{\pm 0.52}$	$62.49_{\pm 0.24}$	$59.99_{\pm 0.38}$	$80.15_{\pm 0.21}$
Model	$16.17_{\pm 0.25}$	$20.08_{\pm 0.17}$	$10.65_{\pm 0.19}$	$32.83_{\pm 0.60}$	$35.88_{\pm 0.19}$	$56.73_{\pm 0.32}$	$49.66_{\pm 0.34}$	$72.44_{\pm 0.15}$
Weight Gaussian	$16.60_{\pm 0.16}$	$20.29_{\pm 0.11}$	$10.53_{\pm 0.24}$	$30.76_{\pm 0.55}$	$42.97_{\pm 0.54}$	$60.68_{\pm 0.24}$	$54.20_{\pm 0.11}$	$75.71_{\pm 0.13}$
Weight DropConnect	$15.82_{\pm 0.06}$	$19.53_{\pm 0.04}$	$12.20_{\pm 0.20}$	$31.29_{\pm 0.63}$	$42.08_{\pm 0.75}$	$59.16_{\pm 0.30}$	$54.33_{\pm 0.70}$	$75.70_{\pm 0.43}$
Top-2 Direct Combination	$12.57_{\pm 0.18}$	$16.04_{\pm 0.19}$	$4.78_{\pm 0.12}$	$16.47_{\pm 0.24}$	$24.69_{\pm 0.22}$	$43.04_{\pm 0.24}$	$37.18_{\pm 0.04}$	$58.16_{\pm 0.12}$
Top-3 Direct Combination	$13.49_{\pm 0.19}$	$16.95_{\pm 0.10}$	$5.07_{\pm 0.07}$	$17.72_{\pm 0.07}$	$24.88_{\pm 0.13}$	$42.99_{\pm 0.21}$	$37.43_{\pm 0.12}$	$58.73_{\pm 0.19}$
Top-2 Optimised Combination	$12.14_{\pm 0.10}$	$15.51_{\pm 0.11}$	$5.15_{\pm 0.22}$	$15.63_{\pm 0.27}$	$25.28_{\pm 0.15}$	$43.86_{\pm 0.16}$	$36.88_{\pm 0.08}$	$57.90_{\pm 0.24}$
Top-3 Optimised Combination	$12.66_{\pm 0.04}$	$16.08_{\pm 0.04}$	$6.60_{\pm 0.16}$	$17.91_{\pm 0.63}$	$26.19_{\pm 0.23}$	$43.84_{\pm 0.30}$	$36.59_{\pm 0.06}$	$60.36_{\pm 0.17}$

Table 5: CV classification: 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**		**TinyImageNet**	
	ID	**OOD**	**ID**	**OOD**	**ID**	**OOD**	**ID**	**OOD**
No Noise	$13.10_{\pm0.05}$	$15.59_{\pm0.06}$	$4.02_{\pm0.15}$	$16.40_{\pm0.36}$	$5.76_{\pm0.25}$	$7.43_{\pm0.09}$	$16.43_{\pm0.27}$	$10.11_{\pm0.11}$
Input Weak Aug.	$6.35_{\pm0.19}$	$7.85_{\pm0.24}$	$1.92_{\pm0.23}$	$15.91_{\pm0.67}$	$4.83_{\pm0.19}$	$13.28_{\pm0.14}$	$6.49_{\pm0.28}$	$11.46_{\pm0.25}$
Input Gaussian	$13.24_{\pm0.17}$	$15.67_{\pm0.15}$	$4.22_{\pm0.38}$	$16.65_{\pm0.59}$	$5.42_{\pm0.46}$	$7.70_{\pm0.48}$	$15.96_{\pm1.37}$	$10.15_{\pm0.45}$
Input ODS	$13.03_{\pm0.13}$	$15.43_{\pm0.12}$	$4.26_{\pm0.13}$	$15.38_{\pm0.27}$	$5.86_{\pm0.30}$	$7.48_{\pm0.17}$	$26.97_{\pm19.11}$	$28.20_{\pm26.53}$
Input AugMix	$4.81_{\pm0.07}$	$6.03_{\pm0.09}$	$1.18_{\pm0.06}$	$6.27_{\pm0.39}$	$4.41_{\pm0.39}$	$11.38_{\pm0.24}$	$4.87_{\pm0.31}$	$15.78_{\pm0.31}$
Input-Target MixUp	$2.81_{\pm0.08}$	$3.49_{\pm0.07}$	$5.41_{\pm3.04}$	$7.31_{\pm0.66}$	$14.20_{\pm0.13}$	$8.57_{\pm0.42}$	$16.33_{\pm1.40}$	$10.26_{\pm0.32}$
Label Smoothing	$8.66_{\pm0.10}$	$10.61_{\pm0.08}$	$5.11_{\pm0.12}$	$9.46_{\pm0.02}$	$21.55_{\pm0.15}$	$17.10_{\pm0.15}$	$29.45_{\pm0.17}$	$16.62_{\pm0.03}$
Activation Gaussian	$13.09_{\pm0.14}$	$15.56_{\pm0.09}$	$5.50_{\pm0.13}$	$18.80_{\pm0.25}$	$5.12_{\pm0.32}$	$7.58_{\pm0.25}$	$14.77_{\pm0.76}$	$9.83_{\pm0.29}$
Activation Dropout	$5.33_{\pm0.19}$	$6.61_{\pm0.17}$	$4.48_{\pm0.15}$	$18.85_{\pm0.62}$	$5.48_{\pm0.91}$	$8.64_{\pm0.85}$	$10.31_{\pm0.66}$	$20.13_{\pm0.86}$
Gradient Gaussian	$14.84_{\pm0.09}$	$18.39_{\pm0.04}$	$6.03_{\pm0.27}$	$18.79_{\pm0.14}$	$5.54_{\pm0.42}$	$9.70_{\pm0.27}$	$23.56_{\pm0.24}$	$34.33_{\pm0.22}$
Model	$10.93_{\pm0.19}$	$12.82_{\pm0.13}$	$4.42_{\pm0.01}$	$18.43_{\pm0.54}$	$9.06_{\pm0.32}$	$11.00_{\pm0.39}$	$10.80_{\pm0.24}$	$9.01_{\pm0.03}$
Weight Gaussian	$13.38_{\pm0.12}$	$15.72_{\pm0.08}$	$6.48_{\pm0.24}$	$21.54_{\pm0.66}$	$5.99_{\pm0.26}$	$7.79_{\pm0.36}$	$14.95_{\pm1.03}$	$9.98_{\pm0.33}$
Weight DropConnect	$12.51_{\pm0.10}$	$14.87_{\pm0.05}$	$4.94_{\pm0.20}$	$16.99_{\pm0.66}$	$5.79_{\pm0.36}$	$8.21_{\pm0.47}$	$15.50_{\pm0.83}$	$10.12_{\pm0.35}$
Top-2 Direct Combination	$1.79_{\pm0.15}$	$2.76_{\pm0.17}$	$1.25_{\pm0.15}$	$6.91_{\pm0.26}$	$6.19_{\pm0.19}$	$14.21_{\pm0.44}$	$4.04_{\pm0.27}$	$15.69_{\pm0.34}$
Top-3 Direct Combination	$1.31_{\pm0.17}$	$1.86_{\pm0.09}$	$2.08_{\pm0.12}$	$8.57_{\pm0.33}$	$6.66_{\pm0.04}$	$14.67_{\pm0.21}$	$13.05_{\pm1.09}$	$22.67_{\pm1.96}$
Top-2 Optimised Combination	$2.55_{\pm0.07}$	$3.48_{\pm0.10}$	$0.96_{\pm0.13}$	$5.80_{\pm0.16}$	$6.68_{\pm0.22}$	$14.54_{\pm0.46}$	$3.21_{\pm0.39}$	$15.33_{\pm0.05}$
Top-3 Optimised Combination	$2.89_{\pm0.11}$	$3.98_{\pm0.08}$	$1.15_{\pm0.34}$	$4.97_{\pm1.05}$	$5.60_{\pm0.31}$	$12.56_{\pm0.61}$	$9.90_{\pm0.42}$	$20.11_{\pm0.26}$

Table 6: CV classification: 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**		**TinyImageNet**	
	ID	**OOD**	**ID**	**OOD**	**ID**	**OOD**	**ID**	**OOD**
No Noise	$1.43_{\pm0.00}$	$1.64_{\pm0.01}$	$0.42_{\pm0.01}$	$1.21_{\pm0.02}$	$1.85_{\pm0.03}$	$2.69_{\pm0.03}$	$2.80_{\pm0.02}$	$3.90_{\pm0.01}$
Input Weak Aug.	$0.62_{\pm0.01}$	$0.75_{\pm0.01}$	$0.20_{\pm0.00}$	$1.10_{\pm0.03}$	$1.07_{\pm0.01}$	$2.44_{\pm0.00}$	$1.84_{\pm0.02}$	$3.46_{\pm0.02}$
Input Gaussian	$1.43_{\pm0.01}$	$1.65_{\pm0.01}$	$0.42_{\pm0.01}$	$1.21_{\pm0.03}$	$1.84_{\pm0.08}$	$2.67_{\pm0.08}$	$2.75_{\pm0.06}$	$3.86_{\pm0.02}$
Input ODS	$1.40_{\pm0.02}$	$1.60_{\pm0.02}$	$0.42_{\pm0.01}$	$1.14_{\pm0.02}$	$1.84_{\pm0.03}$	$2.69_{\pm0.03}$	$4.79_{\pm2.79}$	$6.23_{\pm3.27}$
Input AugMix	$0.49_{\pm0.01}$	$0.61_{\pm0.00}$	$0.25_{\pm0.00}$	$0.63_{\pm0.01}$	$1.17_{\pm0.00}$	$2.04_{\pm0.01}$	$1.87_{\pm0.01}$	$3.14_{\pm0.01}$
Input-Target MixUp	$0.52_{\pm0.00}$	$0.64_{\pm0.00}$	$0.40_{\pm0.03}$	$0.91_{\pm0.02}$	$2.03_{\pm0.00}$	$2.77_{\pm0.01}$	$2.79_{\pm0.07}$	$3.88_{\pm0.04}$
Label Smoothing	$0.75_{\pm0.01}$	$0.90_{\pm0.01}$	$0.46_{\pm0.01}$	$1.09_{\pm0.01}$	$2.18_{\pm0.02}$	$2.87_{\pm0.03}$	$3.25_{\pm0.02}$	$4.15_{\pm0.01}$
Activation Gaussian	$1.41_{\pm0.00}$	$1.63_{\pm0.00}$	$0.43_{\pm0.00}$	$1.27_{\pm0.01}$	$1.84_{\pm0.01}$	$2.71_{\pm0.01}$	$2.74_{\pm0.02}$	$3.87_{\pm0.01}$
Activation Dropout	$0.51_{\pm0.01}$	$0.63_{\pm0.01}$	$0.33_{\pm0.01}$	$1.26_{\pm0.04}$	$1.71_{\pm0.03}$	$2.58_{\pm0.04}$	$1.97_{\pm0.02}$	$3.69_{\pm0.05}$
Gradient Gaussian	$1.76_{\pm0.01}$	$2.12_{\pm0.01}$	$0.56_{\pm0.01}$	$1.39_{\pm0.01}$	$1.86_{\pm0.02}$	$2.76_{\pm0.01}$	$3.20_{\pm0.02}$	$5.36_{\pm0.03}$
Model	$1.02_{\pm0.01}$	$1.17_{\pm0.01}$	$0.37_{\pm0.00}$	$1.27_{\pm0.03}$	$1.60_{\pm0.02}$	$2.63_{\pm0.02}$	$2.38_{\pm0.01}$	$3.64_{\pm0.01}$
Weight Gaussian	$1.44_{\pm0.02}$	$1.65_{\pm0.02}$	$0.44_{\pm0.01}$	$1.41_{\pm0.04}$	$1.79_{\pm0.03}$	$2.66_{\pm0.02}$	$2.76_{\pm0.03}$	$3.89_{\pm0.02}$
Weight DropConnect	$1.32_{\pm0.01}$	$1.52_{\pm0.01}$	$0.43_{\pm0.01}$	$1.21_{\pm0.03}$	$1.73_{\pm0.04}$	$2.58_{\pm0.02}$	$2.78_{\pm0.06}$	$3.90_{\pm0.04}$
Top-2 Direct Combination	$0.44_{\pm0.00}$	$0.54_{\pm0.01}$	$0.16_{\pm0.00}$	$0.57_{\pm0.01}$	$0.97_{\pm0.00}$	$1.96_{\pm0.02}$	$1.60_{\pm0.00}$	$2.98_{\pm0.01}$
Top-3 Direct Combination	$0.45_{\pm0.00}$	$0.55_{\pm0.00}$	$0.17_{\pm0.00}$	$0.63_{\pm0.01}$	$0.97_{\pm0.00}$	$1.97_{\pm0.01}$	$1.73_{\pm0.02}$	$3.25_{\pm0.10}$
Top-2 Optimised Combination	$0.43_{\pm0.00}$	$0.54_{\pm0.00}$	$0.17_{\pm0.01}$	$0.53_{\pm0.01}$	$0.98_{\pm0.01}$	$2.00_{\pm0.02}$	$1.58_{\pm0.01}$	$2.96_{\pm0.02}$
Top-3 Optimised Combination	$0.44_{\pm0.00}$	$0.55_{\pm0.00}$	$0.21_{\pm0.00}$	$0.57_{\pm0.00}$	$0.99_{\pm0.01}$	$1.93_{\pm0.03}$	$1.63_{\pm0.01}$	$3.25_{\pm0.01}$

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

Noise Type	CIFAR-10		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	$16.09_{\pm 0.18}$	$30.81_{\pm 0.76}$	$40.97_{\pm 0.40}$	$61.56_{\pm 0.11}$	$54.70_{\pm 0.81}$	$75.96_{\pm 0.52}$
Input Weak Aug.	$7.99_{\pm 0.07}$	$27.60_{\pm 0.42}$	$24.03_{\pm 0.09}$	$52.82_{\pm 0.13}$	$39.60_{\pm 0.26}$	$67.41_{\pm 0.02}$
Input Gaussian	$16.72_{\pm 0.13}$	$30.28_{\pm 0.85}$	$41.07_{\pm 0.25}$	$61.04_{\pm 0.14}$	$54.01_{\pm 0.60}$	$75.31_{\pm 0.33}$
Input ODS	$15.79_{\pm 0.13}$	$29.54_{\pm 0.71}$	$41.13_{\pm 0.43}$	$60.43_{\pm 0.39}$	$52.77_{\pm 0.54}$	$74.61_{\pm 0.11}$
Input AugMix	$10.26_{\pm 0.04}$	$20.95_{\pm 0.06}$	$30.18_{\pm 0.46}$	$45.94_{\pm 0.29}$	$40.04_{\pm 0.29}$	$60.67_{\pm 0.15}$
Input-Target MixUp	$16.90_{\pm 0.17}$	$32.62_{\pm 0.36}$	$39.04_{\pm 0.30}$	$58.48_{\pm 0.03}$	$51.49_{\pm 0.27}$	$72.87_{\pm 0.12}$
Label Smoothing	$17.18_{\pm 0.13}$	$31.10_{\pm 0.68}$	$42.00_{\pm 0.09}$	$61.49_{\pm 0.30}$	$52.33_{\pm 0.14}$	$73.95_{\pm 0.08}$
Activation Gaussian	$16.34_{\pm 0.15}$	$30.52_{\pm 0.79}$	$38.98_{\pm 0.19}$	$59.70_{\pm 0.22}$	$52.49_{\pm 0.23}$	$74.62_{\pm 0.21}$
Activation Dropout	$12.45_{\pm 0.12}$	$27.90_{\pm 0.10}$	$31.58_{\pm 0.54}$	$56.38_{\pm 0.36}$	$51.85_{\pm 0.10}$	$74.08_{\pm 0.01}$
Gradient Gaussian	$18.70_{\pm 0.15}$	$34.04_{\pm 0.49}$	$47.64_{\pm 0.31}$	$67.79_{\pm 0.22}$	$55.91_{\pm 0.25}$	$76.96_{\pm 0.06}$
Model	$13.11_{\pm 0.25}$	$32.40_{\pm 0.27}$	$79.31_{\pm 27.85}$	$87.90_{\pm 15.70}$	$48.86_{\pm 0.13}$	$72.07_{\pm 0.12}$
Weight Gaussian	$16.54_{\pm 0.11}$	$30.86_{\pm 0.29}$	$37.14_{\pm 0.19}$	$58.32_{\pm 0.48}$	$52.88_{\pm 0.45}$	$74.82_{\pm 0.18}$
Weight DropConnect	$16.79_{\pm 0.28}$	$28.80_{\pm 0.80}$	$41.18_{\pm 0.56}$	$61.39_{\pm 0.02}$	$53.67_{\pm 0.19}$	$75.28_{\pm 0.14}$

Table 8: CV classification: Error ($\downarrow, \%$) comparison on in-distribution (ID) test sets and with hyperparameters transferred across datasets.

Noise Type	CIFAR-10		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	$9.76_{\pm 0.07}$	$20.73_{\pm 0.67}$	$12.77_{\pm 0.86}$	$10.73_{\pm 0.38}$	$20.11_{\pm 2.25}$	$11.66_{\pm 0.63}$
Input Weak Aug.	$5.40_{\pm 0.09}$	$20.98_{\pm 0.62}$	$3.65_{\pm 0.38}$	$11.12_{\pm 0.59}$	$5.97_{\pm 0.40}$	$11.75_{\pm 0.48}$
Input Gaussian	$10.39_{\pm 0.15}$	$20.21_{\pm 0.76}$	$12.81_{\pm 0.25}$	$10.66_{\pm 0.16}$	$16.74_{\pm 0.67}$	$10.72_{\pm 0.16}$
Input ODS	$9.89_{\pm 0.12}$	$20.40_{\pm 0.75}$	$11.18_{\pm 2.18}$	$10.28_{\pm 1.07}$	$15.98_{\pm 0.37}$	$10.52_{\pm 0.22}$
Input AugMix	$5.53_{\pm 0.08}$	$11.65_{\pm 0.07}$	$8.94_{\pm 0.07}$	$8.86_{\pm 0.05}$	$3.97_{\pm 0.20}$	$15.43_{\pm 0.53}$
Input-Target MixUp	$6.27_{\pm 0.51}$	$8.12_{\pm 0.20}$	$14.76_{\pm 0.48}$	$10.99_{\pm 0.06}$	$19.33_{\pm 0.27}$	$10.08_{\pm 0.32}$
Label Smoothing	$1.65_{\pm 0.20}$	$6.05_{\pm 0.46}$	$21.53_{\pm 0.30}$	$15.64_{\pm 0.20}$	$27.18_{\pm 0.27}$	$15.49_{\pm 0.11}$
Activation Gaussian	$10.03_{\pm 0.06}$	$20.36_{\pm 0.69}$	$5.07_{\pm 0.19}$	$9.29_{\pm 0.35}$	$14.04_{\pm 0.72}$	$9.88_{\pm 0.11}$
Activation Dropout	$9.41_{\pm 0.17}$	$21.79_{\pm 0.11}$	$7.94_{\pm 0.39}$	$19.81_{\pm 0.47}$	$16.67_{\pm 0.07}$	$11.00_{\pm 0.12}$
Gradient Gaussian	$12.77_{\pm 0.20}$	$24.59_{\pm 0.50}$	$8.38_{\pm 0.33}$	$18.02_{\pm 0.19}$	$13.41_{\pm 0.73}$	$10.07_{\pm 0.14}$
Model	$7.53_{\pm 0.26}$	$21.33_{\pm 0.36}$	$0.80_{\pm 1.10}$	$4.07_{\pm 5.73}$	$14.45_{\pm 0.23}$	$10.68_{\pm 0.07}$
Weight Gaussian	$10.05_{\pm 0.11}$	$20.70_{\pm 0.23}$	$11.79_{\pm 1.29}$	$22.28_{\pm 1.77}$	$15.55_{\pm 0.55}$	$10.38_{\pm 0.30}$
Weight DropConnect	$10.93_{\pm 0.26}$	$19.75_{\pm 0.67}$	$11.24_{\pm 0.70}$	$9.93_{\pm 0.18}$	$18.10_{\pm 0.20}$	$11.11_{\pm 0.10}$

Table 9: CV classification: ECE ($\downarrow, \%$) comparison on in-distribution (ID) test sets and with hyperparameters transferred across datasets.

Noise Type	CIFAR-10		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	0.71±0.01	1.55±0.05	1.88±0.02	2.87±0.01	2.96±0.12	3.99±0.06
Input Weak Aug.	0.40±0.00	1.84±0.09	1.02±0.01	2.45±0.02	1.80±0.01	3.43±0.02
Input Gaussian	0.72±0.01	1.50±0.06	1.89±0.01	2.84±0.01	2.81±0.06	3.91±0.03
Input ODS	0.71±0.00	1.55±0.07	1.87±0.07	2.80±0.05	2.72±0.02	3.85±0.00
Input AugMix	0.39±0.00	0.84±0.00	1.30±0.01	2.02±0.01	1.73±0.01	3.10±0.02
Input-Target MixUp	0.56±0.01	1.04±0.02	1.71±0.01	2.59±0.00	2.62±0.02	3.65±0.01
Label Smoothing	0.56±0.00	1.01±0.02	2.16±0.01	3.04±0.01	3.16±0.00	4.10±0.01
Activation Gaussian	0.70±0.01	1.51±0.06	1.66±0.01	2.73±0.02	2.64±0.03	3.81±0.02
Activation Dropout	0.74±0.02	1.96±0.02	1.24±0.02	2.69±0.03	2.70±0.01	3.83±0.00
Gradient Gaussian	0.93±0.01	1.99±0.04	1.93±0.01	3.20±0.02	2.83±0.02	3.98±0.01
Model	0.52±0.00	1.43±0.02	3.59±1.44	4.07±0.75	2.46±0.01	3.69±0.01
Weight Gaussian	0.72±0.01	1.54±0.03	1.54±0.03	2.81±0.07	2.72±0.04	3.86±0.02
Weight DropConnect	0.79±0.01	1.48±0.06	1.86±0.01	2.85±0.00	2.83±0.01	3.92±0.00

Table 10: CV classification: NLL (↓) comparison on in-distribution (ID) test sets and with hyperparameters transferred across datasets.

Noise Type	SVHN		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	5.12±0.13	9.20±0.10	43.87±0.48	61.25±0.46	53.63±0.17	75.20±0.32
Input Weak Aug.	4.13±0.10	8.51±0.23	27.33±0.35	51.09±0.18	38.21±0.24	64.75±0.29
Input Gaussian	5.01±0.08	9.09±0.06	43.05±0.35	58.71±0.32	54.15±1.59	74.61±1.09
Input AugMix	3.51±0.05	8.27±0.04	30.23±0.06	45.51±0.17	42.05±0.29	59.93±0.10
Input-Target MixUp	5.58±0.12	12.75±0.11	43.93±0.68	59.98±0.23	52.95±1.26	73.82±0.54
Label Smoothing	5.04±0.03	8.88±0.01	42.99±0.36	57.27±0.43	53.77±0.32	74.25±0.11
Activation Gaussian	5.14±0.08	9.17±0.06	43.74±0.31	59.29±0.65	56.32±2.03	76.49±1.23
Activation Dropout	4.37±0.02	8.16±0.04	42.89±0.90	58.26±0.34	42.47±0.35	67.78±0.25
Gradient Gaussian	6.25±0.10	11.43±0.14	44.39±0.73	59.56±0.87	57.92±0.75	77.81±0.34
Model	3.98±0.02	8.11±0.04	37.14±0.33	56.97±0.16	47.11±0.53	69.76±0.09
Weight Gaussian	5.04±0.04	9.08±0.06	44.28±0.31	59.84±0.08	54.11±2.41	74.52±1.42
Weight DropConnect	5.11±0.11	9.03±0.11	45.75±0.89	59.71±1.15	56.44±3.41	76.12±2.55

Table 11: CV classification: Error (↓, %) comparison on in-distribution (ID) test sets and with hyperparameters transferred across architectures.

Noise Type	SVHN		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	2.73$_{\pm 0.10}$	5.13$_{\pm 0.06}$	5.76$_{\pm 0.34}$	7.50$_{\pm 0.11}$	14.68$_{\pm 0.47}$	9.93$_{\pm 0.08}$
Input Weak Aug.	2.66$_{\pm 0.05}$	5.75$_{\pm 0.09}$	6.33$_{\pm 0.32}$	15.26$_{\pm 0.80}$	8.20$_{\pm 1.23}$	9.88$_{\pm 0.38}$
Input Gaussian	2.68$_{\pm 0.05}$	5.09$_{\pm 0.03}$	5.07$_{\pm 0.61}$	9.59$_{\pm 1.63}$	13.79$_{\pm 0.49}$	8.75$_{\pm 0.46}$
Input AugMix	1.46$_{\pm 0.04}$	3.04$_{\pm 0.07}$	6.56$_{\pm 0.71}$	14.17$_{\pm 0.59}$	6.23$_{\pm 0.63}$	16.60$_{\pm 0.23}$
Input-Target MixUp	5.61$_{\pm 1.19}$	5.27$_{\pm 0.90}$	2.68$_{\pm 0.26}$	6.77$_{\pm 0.43}$	12.24$_{\pm 0.82}$	8.32$_{\pm 0.10}$
Label Smoothing	0.95$_{\pm 0.07}$	1.16$_{\pm 0.03}$	10.18$_{\pm 2.84}$	8.54$_{\pm 1.72}$	17.47$_{\pm 0.22}$	9.76$_{\pm 0.12}$
Activation Gaussian	2.75$_{\pm 0.05}$	5.11$_{\pm 0.07}$	5.45$_{\pm 1.23}$	10.46$_{\pm 1.78}$	10.66$_{\pm 3.63}$	8.42$_{\pm 0.89}$
Activation Dropout	3.03$_{\pm 0.02}$	5.72$_{\pm 0.02}$	5.49$_{\pm 0.32}$	10.00$_{\pm 0.41}$	14.68$_{\pm 0.17}$	25.16$_{\pm 0.75}$
Gradient Gaussian	3.61$_{\pm 0.11}$	7.07$_{\pm 0.15}$	5.59$_{\pm 1.20}$	11.16$_{\pm 1.73}$	21.10$_{\pm 0.45}$	31.33$_{\pm 0.35}$
Model	2.38$_{\pm 0.05}$	4.78$_{\pm 0.04}$	7.59$_{\pm 0.52}$	13.21$_{\pm 0.40}$	5.63$_{\pm 0.31}$	10.34$_{\pm 0.44}$
Weight Gaussian	2.67$_{\pm 0.03}$	5.06$_{\pm 0.02}$	5.30$_{\pm 0.20}$	9.87$_{\pm 0.53}$	12.42$_{\pm 0.87}$	8.24$_{\pm 0.34}$
Weight DropConnect	2.91$_{\pm 0.08}$	5.38$_{\pm 0.07}$	6.98$_{\pm 1.36}$	11.51$_{\pm 1.69}$	11.42$_{\pm 2.14}$	7.99$_{\pm 0.82}$

Table 12: CV classification: ECE ($\downarrow$, %) comparison on in-distribution (ID) test sets and with hyperparameters transferred across architectures.

Noise Type	SVHN		CIFAR-100		TinyImageNet	
	ID	OOD	ID	OOD	ID	OOD
No Noise	0.23$_{\pm 0.01}$	0.41$_{\pm 0.00}$	1.83$_{\pm 0.03}$	2.69$_{\pm 0.02}$	2.72$_{\pm 0.03}$	3.86$_{\pm 0.02}$
Input Weak Aug.	0.23$_{\pm 0.00}$	0.48$_{\pm 0.00}$	1.09$_{\pm 0.01}$	2.36$_{\pm 0.02}$	1.84$_{\pm 0.05}$	3.28$_{\pm 0.04}$
Input Gaussian	0.23$_{\pm 0.00}$	0.41$_{\pm 0.00}$	1.75$_{\pm 0.03}$	2.56$_{\pm 0.02}$	2.67$_{\pm 0.08}$	3.76$_{\pm 0.06}$
Input AugMix	0.16$_{\pm 0.00}$	0.31$_{\pm 0.00}$	1.19$_{\pm 0.01}$	2.05$_{\pm 0.00}$	1.90$_{\pm 0.01}$	3.05$_{\pm 0.01}$
Input-Target MixUp	0.24$_{\pm 0.01}$	0.46$_{\pm 0.01}$	1.76$_{\pm 0.03}$	2.58$_{\pm 0.01}$	2.56$_{\pm 0.03}$	3.68$_{\pm 0.03}$
Label Smoothing	0.19$_{\pm 0.00}$	0.31$_{\pm 0.00}$	1.99$_{\pm 0.06}$	2.65$_{\pm 0.06}$	2.82$_{\pm 0.02}$	3.86$_{\pm 0.01}$
Activation Gaussian	0.23$_{\pm 0.00}$	0.41$_{\pm 0.00}$	1.78$_{\pm 0.02}$	2.59$_{\pm 0.04}$	2.72$_{\pm 0.19}$	3.84$_{\pm 0.09}$
Activation Dropout	0.26$_{\pm 0.00}$	0.48$_{\pm 0.00}$	1.76$_{\pm 0.04}$	2.55$_{\pm 0.01}$	2.03$_{\pm 0.00}$	3.70$_{\pm 0.04}$
Gradient Gaussian	0.30$_{\pm 0.01}$	0.57$_{\pm 0.02}$	1.79$_{\pm 0.03}$	2.60$_{\pm 0.04}$	2.94$_{\pm 0.05}$	4.79$_{\pm 0.04}$
Model	0.18$_{\pm 0.00}$	0.35$_{\pm 0.00}$	1.59$_{\pm 0.01}$	2.67$_{\pm 0.01}$	2.10$_{\pm 0.02}$	3.46$_{\pm 0.00}$
Weight Gaussian	0.23$_{\pm 0.00}$	0.40$_{\pm 0.00}$	1.81$_{\pm 0.02}$	2.61$_{\pm 0.01}$	2.63$_{\pm 0.13}$	3.74$_{\pm 0.09}$
Weight DropConnect	0.24$_{\pm 0.01}$	0.43$_{\pm 0.01}$	1.87$_{\pm 0.04}$	2.61$_{\pm 0.04}$	2.73$_{\pm 0.16}$	3.82$_{\pm 0.15}$

Table 13: CV classification: NLL ($\downarrow$) comparison on in-distribution (ID) test set and with hyperparameters transferred across architectures.

Noise Type	Wine		Toxicity		Abalone		Students		Adult	
	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD
No Noise	35.73$_{\pm 2.20}$	64.16$_{\pm 0.62}$	47.06$_{\pm 8.32}$	42.78$_{\pm 5.83}$	43.30$_{\pm 0.29}$	47.32$_{\pm 0.13}$	65.40$_{\pm 4.88}$	68.52$_{\pm 2.38}$	15.54$_{\pm 0.23}$	17.03$_{\pm 0.19}$
Input Gaussian	35.94$_{\pm 2.43}$	64.15$_{\pm 0.68}$	47.06$_{\pm 8.32}$	42.51$_{\pm 5.84}$	43.30$_{\pm 0.43}$	47.37$_{\pm 0.32}$	65.40$_{\pm 4.88}$	68.42$_{\pm 1.97}$	15.58$_{\pm 0.42}$	16.99$_{\pm 0.29}$
Input ODS	35.83$_{\pm 1.64}$	63.00$_{\pm 0.59}$	48.04$_{\pm 5.55}$	42.78$_{\pm 5.74}$	43.38$_{\pm 0.73}$	46.65$_{\pm 0.43}$	64.56$_{\pm 3.73}$	68.64$_{\pm 2.48}$	15.47$_{\pm 0.39}$	16.97$_{\pm 0.22}$
Input-Target MixUp	38.02$_{\pm 0.82}$	63.21$_{\pm 0.77}$	50.98$_{\pm 5.00}$	44.20$_{\pm 3.34}$	43.46$_{\pm 0.45}$	47.25$_{\pm 0.31}$	65.40$_{\pm 2.60}$	69.05$_{\pm 2.38}$	15.28$_{\pm 0.33}$	16.66$_{\pm 0.25}$
Label Smoothing	35.52$_{\pm 2.37}$	64.07$_{\pm 0.86}$	50.00$_{\pm 4.16}$	44.20$_{\pm 4.48}$	43.22$_{\pm 0.20}$	47.28$_{\pm 0.33}$	64.98$_{\pm 4.30}$	68.56$_{\pm 2.23}$	15.55$_{\pm 0.35}$	16.97$_{\pm 0.22}$
Activation Gaussian	35.00$_{\pm 2.04}$	63.99$_{\pm 0.83}$	46.08$_{\pm 6.93}$	42.67$_{\pm 5.62}$	43.10$_{\pm 0.50}$	47.28$_{\pm 0.20}$	65.40$_{\pm 4.88}$	68.44$_{\pm 2.32}$	15.55$_{\pm 0.24}$	16.91$_{\pm 0.18}$
Activation Dropout	35.52$_{\pm 1.45}$	63.05$_{\pm 1.96}$	54.90$_{\pm 14.48}$	52.78$_{\pm 12.06}$	43.02$_{\pm 1.17}$	46.89$_{\pm 0.72}$	71.31$_{\pm 4.30}$	72.74$_{\pm 3.00}$	14.75$_{\pm 0.31}$	15.88$_{\pm 0.23}$
Gradient Gaussian	32.60$_{\pm 0.53}$	62.78$_{\pm 4.01}$	50.00$_{\pm 4.16}$	48.39$_{\pm 1.47}$	43.74$_{\pm 0.11}$	47.57$_{\pm 0.41}$	68.78$_{\pm 5.69}$	70.62$_{\pm 3.29}$	15.38$_{\pm 0.31}$	16.77$_{\pm 0.25}$
Model	37.92$_{\pm 1.28}$	63.49$_{\pm 0.91}$	45.10$_{\pm 6.04}$	42.71$_{\pm 5.46}$	43.42$_{\pm 0.20}$	47.43$_{\pm 0.35}$	65.82$_{\pm 4.51}$	68.39$_{\pm 2.16}$	14.72$_{\pm 0.46}$	15.90$_{\pm 0.16}$
Weight Gaussian	35.52$_{\pm 1.95}$	64.20$_{\pm 0.77}$	50.00$_{\pm 6.35}$	46.67$_{\pm 4.61}$	43.10$_{\pm 0.39}$	47.34$_{\pm 0.13}$	64.98$_{\pm 4.88}$	69.87$_{\pm 3.22}$	15.13$_{\pm 0.31}$	16.41$_{\pm 0.20}$
Weight DropConnect	38.85$_{\pm 0.53}$	61.50$_{\pm 1.15}$	47.06$_{\pm 8.32}$	43.22$_{\pm 5.81}$	43.42$_{\pm 0.78}$	46.75$_{\pm 0.28}$	64.56$_{\pm 4.51}$	69.57$_{\pm 3.15}$	14.93$_{\pm 0.32}$	16.21$_{\pm 0.18}$
Top-2 Direct Combination	37.29$_{\pm 0.97}$	63.19$_{\pm 0.84}$	39.22$_{\pm 3.67}$	37.18$_{\pm 1.33}$	43.62$_{\pm 0.71}$	47.21$_{\pm 0.88}$	69.62$_{\pm 2.73}$	72.34$_{\pm 1.56}$	14.67$_{\pm 0.26}$	15.83$_{\pm 0.19}$
Top-3 Direct Combination	38.65$_{\pm 0.53}$	63.17$_{\pm 1.06}$	31.37$_{\pm 7.72}$	31.37$_{\pm 7.72}$	43.30$_{\pm 0.68}$	47.10$_{\pm 0.71}$	70.46$_{\pm 2.15}$	72.46$_{\pm 1.41}$	14.71$_{\pm 0.44}$	15.83$_{\pm 0.20}$
Top-2 Optimised Combination	37.40$_{\pm 0.97}$	62.37$_{\pm 1.46}$	33.33$_{\pm 5.00}$	33.06$_{\pm 5.86}$	43.22$_{\pm 0.62}$	47.01$_{\pm 0.26}$	83.54$_{\pm 1.79}$	83.90$_{\pm 1.61}$	14.56$_{\pm 0.28}$	15.77$_{\pm 0.13}$
Top-3 Optimised Combination	37.08$_{\pm 1.26}$	63.17$_{\pm 1.02}$	48.04$_{\pm 7.34}$	42.67$_{\pm 3.24}$	43.22$_{\pm 0.39}$	47.13$_{\pm 0.31}$	74.68$_{\pm 10.18}$	76.49$_{\pm 8.65}$	14.67$_{\pm 0.30}$	15.84$_{\pm 0.14}$

Table 14: Tabular data classification: error ($\downarrow$, %) comparison on in-distribution (ID) test sets and with tuned hyperparameters.

Noise Type	Wine		Toxicity		Abalone		Students		Adult	
	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD
No Noise	$8.75_{\pm1.34}$	$52.81_{\pm2.34}$	$45.78_{\pm8.12}$	$42.21_{\pm5.49}$	$3.94_{\pm0.59}$	$10.68_{\pm0.38}$	$12.46_{\pm2.92}$	$17.32_{\pm0.43}$	$3.47_{\pm0.29}$	$6.11_{\pm0.47}$
Input Gaussian	$8.56_{\pm0.60}$	$52.64_{\pm2.17}$	$45.84_{\pm8.40}$	$41.97_{\pm5.70}$	$3.60_{\pm0.41}$	$10.73_{\pm0.45}$	$13.36_{\pm0.44}$	$17.23_{\pm0.25}$	$3.50_{\pm0.44}$	$6.10_{\pm0.46}$
Input ODS	$5.75_{\pm0.71}$	$48.57_{\pm1.53}$	$42.82_{\pm4.17}$	$38.14_{\pm4.16}$	$3.04_{\pm0.75}$	$8.92_{\pm0.86}$	$14.04_{\pm2.38}$	$16.69_{\pm0.39}$	$3.64_{\pm0.46}$	$6.07_{\pm0.51}$
Input-Target MixUp	$4.43_{\pm0.77}$	$47.55_{\pm3.26}$	$43.18_{\pm5.18}$	$40.06_{\pm2.63}$	$3.41_{\pm1.00}$	$10.45_{\pm0.48}$	$11.07_{\pm3.74}$	$13.47_{\pm2.60}$	$2.79_{\pm0.36}$	$5.32_{\pm0.51}$
Label Smoothing	$4.82_{\pm0.61}$	$50.47_{\pm3.48}$	$42.08_{\pm4.15}$	$37.42_{\pm5.01}$	$2.77_{\pm0.74}$	$9.95_{\pm0.50}$	$12.69_{\pm2.59}$	$17.31_{\pm0.38}$	$3.32_{\pm0.32}$	$5.88_{\pm0.45}$
Activation Gaussian	$9.48_{\pm1.62}$	$52.68_{\pm2.58}$	$46.21_{\pm7.46}$	$41.90_{\pm5.45}$	$3.42_{\pm0.63}$	$10.63_{\pm0.47}$	$13.11_{\pm2.83}$	$17.27_{\pm0.37}$	$3.42_{\pm0.48}$	$5.94_{\pm0.45}$
Activation Dropout	$7.01_{\pm1.60}$	$49.44_{\pm5.72}$	$24.29_{\pm13.35}$	$21.99_{\pm13.93}$	$3.13_{\pm1.64}$	$8.68_{\pm0.82}$	$8.12_{\pm1.38}$	$8.41_{\pm1.46}$	$0.87_{\pm0.16}$	$2.74_{\pm0.09}$
Gradient Gaussian	$16.96_{\pm1.93}$	$55.49_{\pm5.17}$	$48.32_{\pm3.18}$	$47.65_{\pm1.35}$	$5.30_{\pm1.57}$	$13.13_{\pm2.01}$	$17.58_{\pm4.99}$	$20.88_{\pm4.69}$	$2.81_{\pm0.53}$	$5.29_{\pm0.50}$
Model	$4.73_{\pm1.73}$	$49.13_{\pm2.67}$	$39.66_{\pm2.95}$	$36.94_{\pm3.87}$	$2.80_{\pm1.09}$	$9.63_{\pm0.55}$	$12.35_{\pm2.63}$	$17.13_{\pm0.23}$	$1.49_{\pm0.51}$	$3.31_{\pm0.21}$
Weight Gaussian	$9.56_{\pm1.17}$	$53.75_{\pm2.61}$	$48.77_{\pm4.79}$	$44.02_{\pm3.87}$	$3.48_{\pm0.55}$	$10.66_{\pm0.19}$	$12.27_{\pm1.18}$	$17.80_{\pm0.92}$	$2.36_{\pm0.25}$	$4.65_{\pm0.25}$
Weight DropConnect	$5.65_{\pm3.34}$	$42.80_{\pm2.57}$	$46.47_{\pm7.40}$	$42.23_{\pm5.35}$	$3.72_{\pm1.47}$	$9.62_{\pm0.35}$	$12.73_{\pm1.63}$	$17.94_{\pm0.79}$	$1.70_{\pm0.40}$	$3.91_{\pm0.40}$
Top-2 Direct Combination	$4.06_{\pm1.00}$	$48.26_{\pm2.68}$	$18.01_{\pm10.53}$	$18.32_{\pm9.28}$	$3.85_{\pm1.66}$	$8.27_{\pm0.77}$	$12.41_{\pm1.48}$	$10.53_{\pm0.57}$	$0.88_{\pm0.21}$	$2.88_{\pm0.22}$
Top-3 Direct Combination	$3.36_{\pm0.63}$	$46.46_{\pm3.07}$	$11.48_{\pm10.31}$	$11.22_{\pm9.13}$	$3.35_{\pm1.17}$	$7.74_{\pm0.77}$	$10.22_{\pm0.72}$	$9.48_{\pm0.76}$	$1.03_{\pm0.20}$	$2.85_{\pm0.20}$
Top-2 Optimised Combination	$5.90_{\pm2.20}$	$45.81_{\pm1.46}$	$16.46_{\pm6.58}$	$17.43_{\pm5.01}$	$3.23_{\pm1.79}$	$9.67_{\pm0.45}$	$4.16_{\pm3.38}$	$4.38_{\pm2.64}$	$1.28_{\pm0.35}$	$3.10_{\pm0.14}$
Top-3 Optimised Combination	$5.58_{\pm1.06}$	$48.17_{\pm4.78}$	$24.36_{\pm13.90}$	$21.23_{\pm10.47}$	$3.53_{\pm0.87}$	$9.77_{\pm0.25}$	$8.82_{\pm2.45}$	$11.35_{\pm3.43}$	$1.22_{\pm0.12}$	$3.01_{\pm0.04}$

Table 15: Tabular data classification: ECE ($\downarrow, \%$) comparison on in-distribution (ID) test sets and with tuned hyperparameters.

Noise Type	Wine		Toxicity		Abalone		Students		Adult	
	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD
No Noise	$0.94_{\pm0.05}$	$6.25_{\pm0.81}$	$4.85_{\pm0.23}$	$4.95_{\pm0.34}$	$0.84_{\pm0.02}$	$1.04_{\pm0.02}$	$1.87_{\pm0.08}$	$2.41_{\pm0.11}$	$0.35_{\pm0.01}$	$0.44_{\pm0.02}$
Input Gaussian	$0.94_{\pm0.05}$	$6.17_{\pm0.80}$	$4.85_{\pm0.27}$	$4.95_{\pm0.35}$	$0.84_{\pm0.02}$	$1.04_{\pm0.02}$	$1.87_{\pm0.08}$	$2.39_{\pm0.11}$	$0.34_{\pm0.01}$	$0.44_{\pm0.02}$
Input ODS	$0.91_{\pm0.05}$	$5.28_{\pm0.49}$	$2.20_{\pm0.33}$	$2.61_{\pm0.41}$	$0.84_{\pm0.02}$	$0.97_{\pm0.02}$	$1.87_{\pm0.08}$	$2.41_{\pm0.11}$	$0.35_{\pm0.01}$	$0.44_{\pm0.02}$
Input-Target MixUp	$0.91_{\pm0.04}$	$4.69_{\pm0.59}$	$2.39_{\pm0.27}$	$2.59_{\pm0.15}$	$0.84_{\pm0.02}$	$1.03_{\pm0.01}$	$1.88_{\pm0.04}$	$2.25_{\pm0.10}$	$0.34_{\pm0.01}$	$0.43_{\pm0.01}$
Label Smoothing	$0.93_{\pm0.04}$	$5.41_{\pm0.97}$	$1.76_{\pm0.22}$	$1.86_{\pm0.36}$	$0.84_{\pm0.02}$	$1.01_{\pm0.01}$	$1.87_{\pm0.08}$	$2.41_{\pm0.12}$	$0.34_{\pm0.01}$	$0.43_{\pm0.01}$
Activation Gaussian	$0.94_{\pm0.04}$	$6.19_{\pm0.85}$	$4.76_{\pm0.23}$	$4.87_{\pm0.41}$	$0.84_{\pm0.02}$	$1.04_{\pm0.02}$	$1.87_{\pm0.08}$	$2.42_{\pm0.11}$	$0.34_{\pm0.01}$	$0.44_{\pm0.02}$
Activation Dropout	$0.92_{\pm0.05}$	$5.46_{\pm1.06}$	$1.05_{\pm0.25}$	$1.03_{\pm0.24}$	$0.84_{\pm0.02}$	$0.98_{\pm0.01}$	$2.18_{\pm0.05}$	$2.31_{\pm0.09}$	$0.32_{\pm0.01}$	$0.37_{\pm0.01}$
Gradient Gaussian	$1.26_{\pm0.09}$	$8.07_{\pm1.64}$	$5.41_{\pm0.16}$	$6.20_{\pm0.11}$	$0.86_{\pm0.03}$	$1.11_{\pm0.05}$	$1.95_{\pm0.06}$	$2.51_{\pm0.07}$	$0.34_{\pm0.01}$	$0.43_{\pm0.02}$
Model	$0.92_{\pm0.04}$	$4.75_{\pm0.48}$	$2.06_{\pm0.37}$	$2.34_{\pm0.45}$	$0.84_{\pm0.02}$	$1.01_{\pm0.02}$	$1.87_{\pm0.08}$	$2.42_{\pm0.11}$	$0.31_{\pm0.01}$	$0.37_{\pm0.01}$
Weight Gaussian	$0.95_{\pm0.05}$	$6.65_{\pm0.94}$	$3.02_{\pm0.20}$	$3.33_{\pm0.25}$	$0.84_{\pm0.02}$	$1.04_{\pm0.01}$	$1.91_{\pm0.08}$	$2.43_{\pm0.09}$	$0.33_{\pm0.01}$	$0.41_{\pm0.01}$
Weight DropConnect	$0.94_{\pm0.04}$	$4.19_{\pm0.33}$	$4.82_{\pm0.17}$	$4.91_{\pm0.30}$	$0.84_{\pm0.02}$	$0.99_{\pm0.01}$	$1.91_{\pm0.09}$	$2.43_{\pm0.09}$	$0.32_{\pm0.01}$	$0.39_{\pm0.01}$
Top-2 Direct Combination	$0.92_{\pm0.04}$	$4.60_{\pm0.41}$	$1.05_{\pm0.25}$	$1.12_{\pm0.30}$	$0.84_{\pm0.02}$	$0.96_{\pm0.02}$	$2.16_{\pm0.06}$	$2.31_{\pm0.06}$	$0.32_{\pm0.01}$	$0.38_{\pm0.01}$
Top-3 Direct Combination	$0.94_{\pm0.04}$	$4.29_{\pm0.40}$	$0.95_{\pm0.19}$	$0.95_{\pm0.19}$	$0.84_{\pm0.02}$	$0.95_{\pm0.02}$	$2.16_{\pm0.05}$	$2.30_{\pm0.09}$	$0.32_{\pm0.01}$	$0.37_{\pm0.01}$
Top-2 Optimised Combination	$0.91_{\pm0.04}$	$4.34_{\pm0.65}$	$0.97_{\pm0.18}$	$0.99_{\pm0.19}$	$0.84_{\pm0.02}$	$1.01_{\pm0.01}$	$2.50_{\pm0.10}$	$2.52_{\pm0.09}$	$0.31_{\pm0.01}$	$0.37_{\pm0.01}$
Top-3 Optimised Combination	$0.91_{\pm0.04}$	$4.80_{\pm0.91}$	$1.12_{\pm0.35}$	$1.22_{\pm0.42}$	$0.84_{\pm0.02}$	$1.01_{\pm0.01}$	$2.07_{\pm0.30}$	$2.33_{\pm0.14}$	$0.31_{\pm0.00}$	$0.36_{\pm0.01}$

Table 16: Tabular data classification: NLL ($\downarrow$) comparison on in-distribution (ID) test sets and with tuned hyperparameters.

Noise Type	Architecture	
	GP-CNN	Transformer
No Noise	$35.67_{\pm1.13}$	$36.56_{\pm0.83}$
Input Gaussian	$35.44_{\pm0.87}$	$36.59_{\pm0.84}$
Input ODS	$33.56_{\pm0.33}$	$34.44_{\pm0.74}$
Input-Target MixUp	$35.44_{\pm0.78}$	$36.70_{\pm0.92}$
Label Smoothing	$35.59_{\pm0.69}$	$36.56_{\pm1.10}$
Activation Gaussian	$35.56_{\pm1.16}$	$36.44_{\pm0.87}$
Activation Dropout	$39.19_{\pm0.92}$	$36.48_{\pm0.29}$
Gradient Gaussian	$40.56_{\pm0.18}$	$36.52_{\pm0.89}$
Model	$40.22_{\pm0.64}$	$36.52_{\pm0.76}$
Weight Gaussian	$35.48_{\pm0.53}$	$36.89_{\pm0.79}$
Weight DropConnect	$35.19_{\pm0.69}$	$37.04_{\pm0.68}$
Top-2 Direct Combination	$34.30_{\pm0.41}$	$34.33_{\pm0.68}$
Top-3 Direct Combination	$34.59_{\pm0.23}$	$34.41_{\pm0.67}$
Top-2 Optimised Combination	$33.63_{\pm0.76}$	$34.52_{\pm0.55}$
Top-3 Optimised Combination	$34.56_{\pm0.42}$	$36.22_{\pm0.79}$

Table 17: NewsGroup NLP classification: Error ($\downarrow, \%$) comparison on in-distribution (ID) test set and with tuned hyperparameters.

Noise Type	Architecture	
	GP-CNN	Transformer
No Noise	5.12±0.53	3.47±0.98
Input Gaussian	4.78±1.30	3.59±0.83
Input ODS	2.57±0.81	7.99±0.59
Input-Target MixUp	5.07±0.59	2.54±1.27
Label Smoothing	3.78±0.50	4.13±0.94
Activation Gaussian	5.75±0.50	3.42±1.05
Activation Dropout	7.19±1.29	2.26±0.20
Gradient Gaussian	24.26±1.01	3.24±1.07
Model	2.91±0.90	3.55±0.84
Weight Gaussian	4.57±0.48	4.01±0.79
Weight DropConnect	5.41±0.76	4.43±0.72
Top-2 Direct Combination	2.78±0.31	8.19±0.72
Top-3 Direct Combination	2.76±0.76	7.90±0.35
Top-2 Optimised Combination	5.77±0.52	6.26±0.62
Top-3 Optimised Combination	3.12±0.48	3.31±0.70

Table 18: NewsGroup NLP classification: ECE ($\downarrow$, %) comparison on in-distribution (ID) test set and with tuned hyperparameters.

Noise Type	Architecture	
	GP-CNN	Transformer
No Noise	1.14±0.01	1.13±0.02
Input Gaussian	1.12±0.01	1.13±0.02
Input ODS	1.03±0.00	1.10±0.01
Input-Target MixUp	1.13±0.01	1.13±0.02
Label Smoothing	1.12±0.01	1.13±0.02
Activation Gaussian	1.14±0.01	1.13±0.02
Activation Dropout	1.18±0.01	1.11±0.02
Gradient Gaussian	1.87±0.07	1.13±0.02
Model	1.21±0.01	1.13±0.02
Weight Gaussian	1.14±0.01	1.13±0.02
Weight DropConnect	1.13±0.01	1.13±0.02
Top-2 Direct Combination	1.02±0.01	1.10±0.01
Top-3 Direct Combination	1.03±0.01	1.10±0.01
Top-2 Optimised Combination	1.03±0.00	1.09±0.01
Top-3 Optimised Combination	1.02±0.01	1.12±0.01

Table 19: NewsGroup NLP classification: NLL ($\downarrow$) comparison on in-distribution (ID) test set and with tuned hyperparameters.

Noise Type	Rotated CIFAR-100		WikiFace	
	ID	OOD	ID	OOD
No Noise	$0.03_{\pm 0.00}$	$0.13_{\pm 0.01}$	$0.03_{\pm 0.00}$	$0.04_{\pm 0.00}$
Input Weak Aug.	$0.03_{\pm 0.00}$	$0.13_{\pm 0.01}$	$0.03_{\pm 0.00}$	$0.04_{\pm 0.00}$
Input Gaussian	$0.03_{\pm 0.00}$	$0.14_{\pm 0.01}$	$0.04_{\pm 0.00}$	$0.05_{\pm 0.00}$
Input AugMix	$0.03_{\pm 0.00}$	$0.07_{\pm 0.00}$	$0.03_{\pm 0.00}$	$0.04_{\pm 0.00}$
Input-Target CMixUp	$0.03_{\pm 0.00}$	$0.09_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Activation Gaussian	$0.03_{\pm 0.00}$	$0.14_{\pm 0.01}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Activation Dropout	$0.03_{\pm 0.00}$	$0.14_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.05_{\pm 0.00}$
Gradient Gaussian	$0.04_{\pm 0.00}$	$0.11_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Model	$0.04_{\pm 0.00}$	$0.16_{\pm 0.01}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Weight Gaussian	$0.03_{\pm 0.00}$	$0.15_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Weight DropConnect	$0.03_{\pm 0.00}$	$0.12_{\pm 0.01}$	$0.10_{\pm 0.04}$	$0.11_{\pm 0.04}$
Top-2 Direct Combination	$0.03_{\pm 0.00}$	$0.08_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Top-3 Direct Combination	$0.03_{\pm 0.00}$	$0.08_{\pm 0.01}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$
Top-2 Optimised Combination	$0.03_{\pm 0.00}$	$0.06_{\pm 0.00}$	$0.03_{\pm 0.00}$	$0.04_{\pm 0.00}$
Top-3 Optimised Combination	$0.24_{\pm 0.15}$	$0.29_{\pm 0.10}$	$0.04_{\pm 0.00}$	$0.04_{\pm 0.00}$

Table 20: Rotated CV regression: MSE (↓) comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with tuned hyperparameters.

Noise Type	Rotated CIFAR-100		WikiFace	
	ID	OOD	ID	OOD
No Noise	$-4.81_{\pm 0.00}$	$6.90_{\pm 1.52}$	$27.03_{\pm 2.94}$	$31.43_{\pm 3.80}$
Input Weak Aug.	$-4.60_{\pm 0.06}$	$3.41_{\pm 0.53}$	$-0.82_{\pm 0.24}$	$0.31_{\pm 0.66}$
Input Gaussian	$-4.67_{\pm 0.11}$	$7.57_{\pm 2.14}$	$27.96_{\pm 2.66}$	$31.50_{\pm 2.15}$
Input AugMix	$-4.82_{\pm 0.01}$	$-1.70_{\pm 0.04}$	$0.78_{\pm 0.27}$	$-0.12_{\pm 0.11}$
Input-Target CMixUp	$-4.62_{\pm 0.04}$	$1.86_{\pm 0.35}$	$21.83_{\pm 2.35}$	$25.62_{\pm 0.93}$
Activation Gaussian	$-4.27_{\pm 0.07}$	$2.94_{\pm 0.18}$	$14.93_{\pm 1.09}$	$17.61_{\pm 1.01}$
Activation Dropout	$-3.81_{\pm 0.55}$	$1.04_{\pm 0.59}$	$-1.35_{\pm 0.02}$	$-0.53_{\pm 0.36}$
Gradient Gaussian	$-3.70_{\pm 0.00}$	$-0.44_{\pm 0.00}$	$25.87_{\pm 2.19}$	$29.69_{\pm 3.75}$
Model	$-4.36_{\pm 0.05}$	$3.44_{\pm 0.76}$	$-1.08_{\pm 0.03}$	$-1.08_{\pm 0.03}$
Weight Gaussian	$-4.23_{\pm 0.12}$	$2.53_{\pm 0.36}$	$4.42_{\pm 0.10}$	$5.65_{\pm 0.60}$
Weight DropConnect	$-2.28_{\pm 1.93}$	$39.88_{\pm 24.73}$	$4.83_{\pm 3.37}$	$6.00_{\pm 4.13}$
Top-2 Direct Combination	$-4.14_{\pm 0.23}$	$-1.93_{\pm 0.13}$	$-1.34_{\pm 0.01}$	$-1.15_{\pm 0.02}$
Top-3 Direct Combination	$-4.21_{\pm 0.07}$	$-1.76_{\pm 0.01}$	$-1.32_{\pm 0.01}$	$-1.16_{\pm 0.02}$
Top-2 Optimised Combination	$-4.27_{\pm 0.03}$	$-1.71_{\pm 0.05}$	$-1.39_{\pm 0.01}$	$-1.12_{\pm 0.02}$
Top-3 Optimised Combination	$1.84_{\pm 2.73}$	$0.57_{\pm 0.85}$	$-1.34_{\pm 0.03}$	$-1.07_{\pm 0.10}$

Table 21: Rotated CV regression: NLL (↓) comparison on in-distribution (ID) and out-of-distribution (OOD) test sets and with tuned hyperparameters.

Noise Type	Energy		Boston		Wine		Yacht		Concrete	
	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD
No Noise	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.13_{\pm 0.03}$	$0.50_{\pm 0.06}$	$0.25_{\pm 0.04}$	$5600.91_{\pm 1269.15}$	$0.07_{\pm 0.07}$	$1.08_{\pm 0.54}$	$0.10_{\pm 0.01}$	$0.70_{\pm 0.05}$
Input Gaussian	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.12_{\pm 0.02}$	$0.46_{\pm 0.03}$	$0.25_{\pm 0.04}$	$5596.92_{\pm 1263.18}$	$0.02_{\pm 0.02}$	$0.86_{\pm 0.61}$	$0.10_{\pm 0.01}$	$0.73_{\pm 0.08}$
Input-Target CMixUp	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.13_{\pm 0.02}$	$0.51_{\pm 0.03}$	$0.25_{\pm 0.04}$	$5544.15_{\pm 1208.32}$	$0.15_{\pm 0.17}$	$0.59_{\pm 0.22}$	$0.10_{\pm 0.00}$	$0.65_{\pm 0.02}$
Activation Gaussian	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.13_{\pm 0.03}$	$0.53_{\pm 0.07}$	$0.25_{\pm 0.04}$	$5611.73_{\pm 1267.34}$	$0.05_{\pm 0.05}$	$0.93_{\pm 0.65}$	$0.10_{\pm 0.01}$	$0.72_{\pm 0.07}$
Activation Dropout	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.16_{\pm 0.04}$	$0.44_{\pm 0.03}$	$0.38_{\pm 0.06}$	$3018.21_{\pm 1098.31}$	$0.02_{\pm 0.01}$	$0.94_{\pm 0.63}$	$0.10_{\pm 0.01}$	$0.74_{\pm 0.04}$
Gradient Gaussian	$0.04_{\pm 0.00}$	$0.21_{\pm 0.01}$	$0.17_{\pm 0.03}$	$0.68_{\pm 0.06}$	$0.25_{\pm 0.04}$	$5426.31_{\pm 1188.42}$	$0.01_{\pm 0.01}$	$0.99_{\pm 0.54}$	$0.10_{\pm 0.01}$	$0.72_{\pm 0.07}$
Model	$0.04_{\pm 0.00}$	$0.22_{\pm 0.02}$	$0.14_{\pm 0.03}$	$0.54_{\pm 0.04}$	$0.25_{\pm 0.04}$	$5599.59_{\pm 1267.68}$	$0.54_{\pm 0.70}$	$0.70_{\pm 0.59}$	$0.10_{\pm 0.01}$	$0.64_{\pm 0.07}$
Weight Gaussian	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.13_{\pm 0.03}$	$0.55_{\pm 0.03}$	$0.25_{\pm 0.04}$	$5623.27_{\pm 1266.02}$	$0.02_{\pm 0.01}$	$0.61_{\pm 0.30}$	$0.10_{\pm 0.01}$	$0.73_{\pm 0.03}$
Weight DropConnect	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.14_{\pm 0.05}$	$0.57_{\pm 0.07}$	$0.25_{\pm 0.04}$	$5619.57_{\pm 1275.47}$	$0.05_{\pm 0.05}$	$0.50_{\pm 0.10}$	$0.10_{\pm 0.01}$	$0.72_{\pm 0.04}$
Top-2 Direct	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.12_{\pm 0.02}$	$0.56_{\pm 0.04}$	$0.25_{\pm 0.04}$	$5588.08_{\pm 1258.38}$	$0.03_{\pm 0.03}$	$0.80_{\pm 0.61}$	$0.10_{\pm 0.01}$	$0.76_{\pm 0.04}$
Top-3 Direct	$0.04_{\pm 0.00}$	$0.21_{\pm 0.02}$	$0.14_{\pm 0.03}$	$0.52_{\pm 0.05}$	$0.25_{\pm 0.04}$	$5561.37_{\pm 1206.28}$	$0.04_{\pm 0.03}$	$0.75_{\pm 0.51}$	$0.10_{\pm 0.01}$	$0.73_{\pm 0.00}$
Top-2 Optimised	$0.04_{\pm 0.00}$	$0.22_{\pm 0.02}$	$0.14_{\pm 0.01}$	$0.45_{\pm 0.09}$	$0.25_{\pm 0.04}$	$5581.54_{\pm 1246.90}$	$0.02_{\pm 0.01}$	$1.15_{\pm 0.56}$	$0.10_{\pm 0.01}$	$0.71_{\pm 0.06}$
Top-3 Optimised	$0.04_{\pm 0.00}$	$0.22_{\pm 0.02}$	$0.15_{\pm 0.04}$	$0.44_{\pm 0.03}$	$0.25_{\pm 0.04}$	$5576.97_{\pm 1246.68}$	$0.02_{\pm 0.00}$	$0.70_{\pm 0.12}$	$0.10_{\pm 0.01}$	$0.71_{\pm 0.07}$

Table 22: Tabular regression: MSE (↓) comparison on in-distribution (ID) test sets and with tuned hyperparameters.

Noise Type	Energy		Boston		Wine		Yacht		Concrete	
	ID	OOD	ID	OOD	ID	OOD	ID	OOD	ID	OOD
No Noise	$-1.52_{\pm 0.05}$	$1817.31_{\pm 702.69}$	$-0.30_{\pm 0.06}$	$2.53_{\pm 0.94}$	$-0.22_{\pm 0.08}$	$7170328.29_{\pm 3990504.14}$	$-1.18_{\pm 0.22}$	$58.54_{\pm 81.78}$	$-0.56_{\pm 0.09}$	$328.73_{\pm 284.77}$
Input Gaussian	$-1.55_{\pm 0.05}$	$2457.74_{\pm 1324.99}$	$-0.54_{\pm 0.09}$	$1.45_{\pm 0.94}$	$-0.22_{\pm 0.08}$	$7145441.15_{\pm 3994378.22}$	$-1.45_{\pm 0.19}$	$4.77_{\pm 5.63}$	$-0.55_{\pm 0.10}$	$320.13_{\pm 239.89}$
Input-Target CMixUp	$-1.53_{\pm 0.02}$	$3059.13_{\pm 3345.28}$	$-0.43_{\pm 0.11}$	$1.61_{\pm 0.20}$	$-0.19_{\pm 0.09}$	$7761911.37_{\pm 5363752.00}$	$-1.05_{\pm 0.20}$	$1.54_{\pm 1.31}$	$-0.63_{\pm 0.01}$	$140.19_{\pm 87.70}$
Activation Gaussian	$-1.55_{\pm 0.02}$	$3745.28_{\pm 2255.98}$	$-0.07_{\pm 0.25}$	$4.40_{\pm 1.12}$	$-0.22_{\pm 0.08}$	$7104338.64_{\pm 3975769.54}$	$-1.24_{\pm 0.18}$	$5.57_{\pm 5.90}$	$-0.54_{\pm 0.07}$	$384.49_{\pm 393.94}$
Activation Dropout	$-1.53_{\pm 0.06}$	$2313.01_{\pm 980.36}$	$-0.64_{\pm 0.03}$	$-0.05_{\pm 0.10}$	$0.01_{\pm 0.05}$	$180520.25_{\pm 248366.62}$	$-1.41_{\pm 0.29}$	$19.59_{\pm 19.52}$	$-0.57_{\pm 0.08}$	$376.76_{\pm 214.94}$
Gradient Gaussian	$-1.56_{\pm 0.08}$	$4408.72_{\pm 1930.86}$	$0.15_{\pm 0.22}$	$19.01_{\pm 13.05}$	$-0.22_{\pm 0.08}$	$7760431.18_{\pm 2807573.55}$	$-1.86_{\pm 0.47}$	$17.10_{\pm 19.19}$	$-0.56_{\pm 0.08}$	$472.63_{\pm 487.85}$
Model	$-1.56_{\pm 0.03}$	$4409.60_{\pm 1591.54}$	$-0.02_{\pm 0.21}$	$3.43_{\pm 0.70}$	$-0.22_{\pm 0.08}$	$7152203.36_{\pm 3998182.36}$	$-0.77_{\pm 1.13}$	$237.51_{\pm 335.67}$	$-0.59_{\pm 0.06}$	$213.64_{\pm 229.61}$
Weight Gaussian	$-1.56_{\pm 0.03}$	$3777.23_{\pm 1513.61}$	$-0.19_{\pm 0.14}$	$3.79_{\pm 1.58}$	$-0.22_{\pm 0.08}$	$7170268.09_{\pm 4161089.62}$	$-2.04_{\pm 0.45}$	$7.53_{\pm 6.81}$	$-0.53_{\pm 0.09}$	$499.07_{\pm 497.88}$
Weight DropConnect	$-1.55_{\pm 0.04}$	$4191.70_{\pm 2467.82}$	$-0.53_{\pm 0.09}$	$1.59_{\pm 0.54}$	$-0.22_{\pm 0.08}$	$7165628.28_{\pm 4153223.44}$	$-1.31_{\pm 0.46}$	$3.03_{\pm 3.13}$	$-0.58_{\pm 0.09}$	$293.66_{\pm 228.45}$
Top-2 Direct	$-1.52_{\pm 0.02}$	$339.14_{\pm 54.93}$	$-0.14_{\pm 0.23}$	$4.30_{\pm 0.63}$	$-0.22_{\pm 0.08}$	$91454.37_{\pm 45902.04}$	$-1.69_{\pm 0.31}$	$4.59_{\pm 3.78}$	$-0.57_{\pm 0.07}$	$399.03_{\pm 215.22}$
Top-3 Direct	$-1.55_{\pm 0.02}$	$338.68_{\pm 31.63}$	$-0.48_{\pm 0.11}$	$1.09_{\pm 0.72}$	$-0.22_{\pm 0.08}$	$88988.14_{\pm 43835.81}$	$-1.46_{\pm 0.29}$	$19.70_{\pm 28.16}$	$-0.60_{\pm 0.05}$	$276.95_{\pm 189.09}$
Top-2 Optimised	$-1.55_{\pm 0.02}$	$371.09_{\pm 76.51}$	$-0.36_{\pm 0.34}$	$1.83_{\pm 0.84}$	$-0.22_{\pm 0.08}$	$90698.30_{\pm 45723.43}$	$-1.41_{\pm 0.31}$	$84.90_{\pm 88.19}$	$-0.56_{\pm 0.08}$	$274.05_{\pm 255.85}$
Top-3 Optimised	$-1.55_{\pm 0.02}$	$390.96_{\pm 26.40}$	$-0.48_{\pm 0.12}$	$0.09_{\pm 0.05}$	$-0.22_{\pm 0.08}$	$90023.16_{\pm 43206.88}$	$-1.46_{\pm 0.70}$	$8.01_{\pm 6.52}$	$-0.61_{\pm 0.05}$	$256.70_{\pm 217.90}$

Table 23: Tabular regression: NLL (↓) comparison on in-distribution (ID) test sets and with tuned hyperparameters.

Noise Type	Energy		Wine		Concrete	
	ID	OOD	ID	OOD	ID	OOD
No Noise	$0.03_{\pm 0.01}$	$0.18_{\pm 0.04}$	$0.21_{\pm 0.01}$	$7969.85_{\pm 1464.68}$	$0.09_{\pm 0.01}$	$0.68_{\pm 0.21}$
Input Gaussian	$0.03_{\pm 0.01}$	$0.17_{\pm 0.04}$	$0.21_{\pm 0.02}$	$8479.83_{\pm 2666.80}$	$0.08_{\pm 0.01}$	$0.64_{\pm 0.14}$
Input-Target CMixUp	$0.03_{\pm 0.00}$	$0.15_{\pm 0.02}$	$0.18_{\pm 0.01}$	$10948.25_{\pm 355.18}$	$0.10_{\pm 0.01}$	$0.79_{\pm 0.24}$
Activation Gaussian	$0.03_{\pm 0.01}$	$0.15_{\pm 0.03}$	$0.20_{\pm 0.02}$	$8718.85_{\pm 3534.05}$	$0.09_{\pm 0.01}$	$0.58_{\pm 0.06}$
Activation Dropout	$0.03_{\pm 0.00}$	$0.15_{\pm 0.03}$	$0.23_{\pm 0.02}$	$9143.26_{\pm 885.77}$	$0.09_{\pm 0.02}$	$0.66_{\pm 0.16}$
Gradient Gaussian	$0.03_{\pm 0.01}$	$0.19_{\pm 0.06}$	$311.38_{\pm 405.17}$	$380875.57_{\pm 398196.36}$	$0.09_{\pm 0.02}$	$0.81_{\pm 0.20}$
Model	$0.03_{\pm 0.00}$	$0.16_{\pm 0.02}$	$0.21_{\pm 0.01}$	$7392.87_{\pm 1504.76}$	$0.11_{\pm 0.01}$	$0.48_{\pm 0.10}$
Weight Gaussian	$0.03_{\pm 0.01}$	$0.15_{\pm 0.02}$	$0.20_{\pm 0.02}$	$8641.30_{\pm 2279.35}$	$0.09_{\pm 0.02}$	$0.53_{\pm 0.08}$
Weight DropConnect	$0.03_{\pm 0.01}$	$0.15_{\pm 0.02}$	$0.20_{\pm 0.02}$	$9051.78_{\pm 3747.02}$	$0.09_{\pm 0.02}$	$0.53_{\pm 0.05}$

Table 24: Tabular regression: MSE (↓) comparison on in-distribution (ID) test sets and with hyperparameters transferred across datasets.

Noise Type	Energy		Wine		Concrete	
	ID	OOD	ID	OOD	ID	OOD
No Noise	-1.70±0.11	71.13±33.74	4.78±1.31	184737.06±89012.05	-0.54±0.15	163.20±158.58
Input Gaussian	-1.74±0.07	77.67±39.81	5.24±1.58	108313.49±97244.80	-0.40±0.15	75.81±71.04
Input-Target CMixUp	-1.66±0.10	45.14±28.41	-0.06±0.15	17570843.36±24372949.06	-0.65±0.05	439.78±527.26
Activation Gaussian	-1.70±0.10	19.95±6.87	2.92±1.96	13323470.95±18820335.63	-0.47±0.09	43.37±9.14
Activation Dropout	-1.73±0.07	21.38±7.77	-0.29±0.06	50.93±65.83	-0.56±0.12	154.96±164.41
Gradient Gaussian	-1.70±0.11	147.90±111.83	1.97±1.94	39.66±26.95	-0.11±0.10	102.71±64.24
Model	-1.64±0.02	37.29±22.13	4.68±1.33	61705.40±64591.04	-0.64±0.03	25.73±32.25
Weight Gaussian	-1.71±0.06	28.59±12.19	4.51±1.62	28605057.88±18501115.57	-0.55±0.07	35.82±32.72
Weight DropConnect	-1.73±0.12	32.32±14.57	0.37±0.43	7213717.63±10190125.70	-0.67±0.13	7.66±2.88

Table 25: Tabular regression: NLL (↓) comparison on in-distribution (ID) test sets and with hyperparameters transferred across datasets.

Noise Type	Boston		Yacht		Concrete	
	ID	OOD	ID	OOD	ID	OOD
No Noise	0.13±0.04	0.58±0.10	0.56±0.74	0.91±0.50	0.11±0.01	0.61±0.06
Input Gaussian	0.13±0.03	0.51±0.09	0.55±0.75	1.08±0.45	0.10±0.01	0.64±0.07
Input-Target CMixUp	0.50±0.47	0.72±0.32	0.05±0.02	0.70±0.50	0.11±0.01	0.63±0.02
Activation Gaussian	0.14±0.03	0.55±0.09	0.56±0.75	0.98±0.46	0.11±0.01	0.67±0.08
Activation Dropout	0.16±0.04	0.42±0.05	0.03±0.00	0.55±0.09	0.10±0.01	0.61±0.07
Gradient Gaussian	0.16±0.04	0.52±0.05	0.02±0.01	0.71±0.21	0.11±0.01	0.61±0.06
Model	0.15±0.06	0.55±0.03	0.11±0.04	0.36±0.08	0.12±0.01	0.59±0.02
Weight Gaussian	0.12±0.04	0.48±0.06	0.03±0.00	0.54±0.11	0.11±0.01	0.64±0.08
Weight DropConnect	0.13±0.03	0.52±0.05	0.04±0.02	0.80±0.62	0.11±0.01	0.63±0.06

Table 26: Tabular regression: MSE (↓) comparison on in-distribution (ID) test sets and with hyperparameters transferred across architectures.

Noise Type	Boston		Yacht		Concrete	
	ID	OOD	ID	OOD	ID	OOD
No Noise	-0.25±0.21	3.45±0.36	-0.42±0.88	3.02±4.02	-0.57±0.06	145.83±127.35
Input Gaussian	-0.35±0.10	2.45±0.94	-0.94±1.25	3.59±1.96	-0.58±0.08	179.80±132.61
Input-Target CMixUp	-0.09±0.29	1.06±0.52	-1.20±0.07	1.03±1.20	-0.60±0.04	118.96±75.71
Activation Gaussian	-0.23±0.18	2.83±1.99	-0.46±0.91	1.60±1.55	-0.59±0.08	207.32±127.13
Activation Dropout	-0.57±0.02	-0.13±0.05	-1.67±0.35	3.44±2.50	-0.58±0.06	163.41±151.15
Gradient Gaussian	0.97±1.13	5.51±2.97	-1.52±0.43	5.11±7.37	-0.57±0.06	145.83±127.35
Model	-0.31±0.25	3.08±0.43	-0.57±0.10	-0.24±0.10	-0.58±0.03	54.63±27.81
Weight Gaussian	-0.42±0.18	1.79±0.32	-1.62±0.34	8.84±5.65	-0.56±0.07	188.86±173.50
Weight DropConnect	-0.46±0.09	1.36±0.18	-1.72±0.31	7.51±7.67	-0.55±0.06	168.03±139.58

Table 27: Tabular regression: NLL (↓) comparison on in-distribution (ID) test sets and with hyperparameters transferred across architectures.

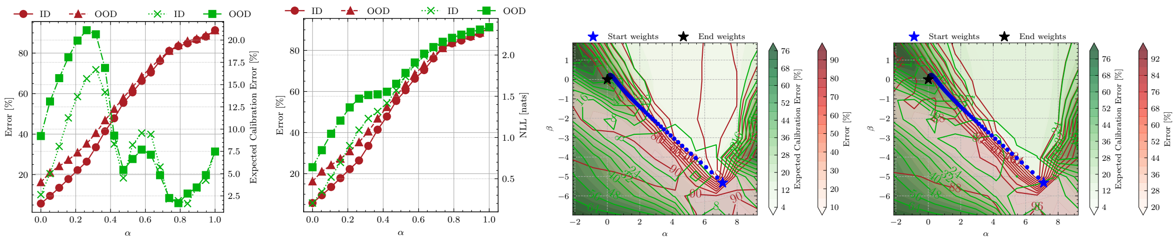

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 9: Input Random Crop, Horizontal Flip on CIFAR-10. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

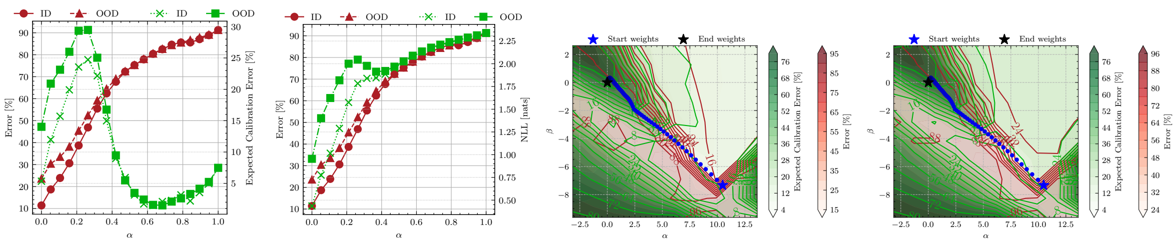

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 10: Input Additive Gaussian on CIFAR-10. *Observations*: Changed the smoothness of the 1D curves where NLL became less smooth and removed the bumps in ECE for α approaching the initial model. The 2D metric landscape trajectory did not change in comparison to no noise.

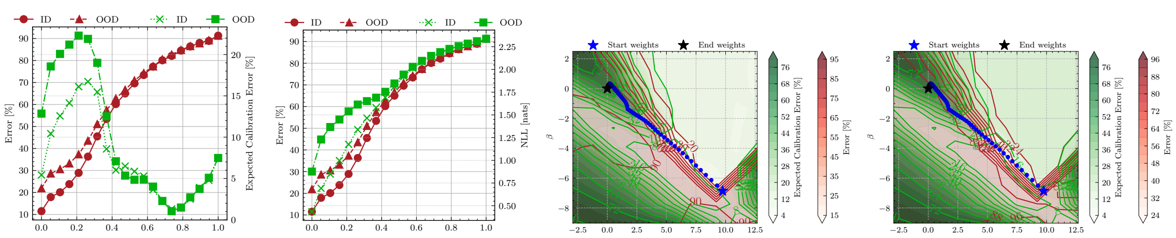

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 11: Input ODS on CIFAR-10. *Observations*: Marginally changed the smoothness of the 1D curves. The 2D metric landscape trajectory did not change in comparison to no noise.

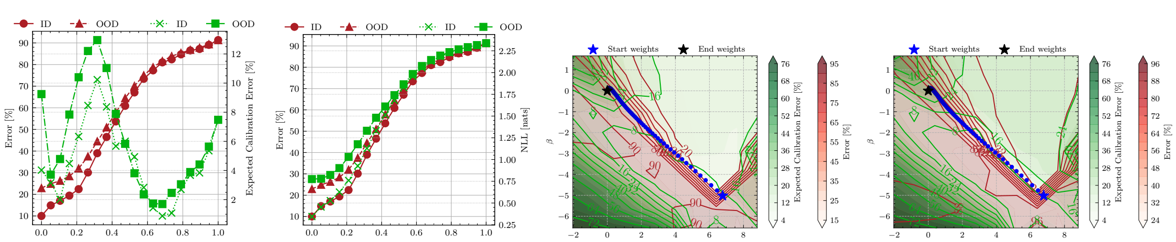

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 12: Input-Target MixUp on CIFAR-10. *Observations*: Both the NLL and ECE 1D curves changed in comparison to no noise, and the 2D plots seem to explore wider valleys compared to no noise.

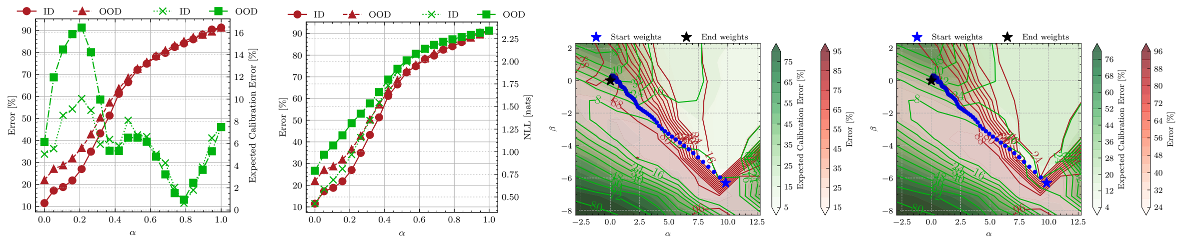

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 13: Target Smoothing on CIFAR-10. *Observations*: The NLL became more aligned with the error, not the ECE. The 2D plots show slightly more variation in the trajectory than no noise.

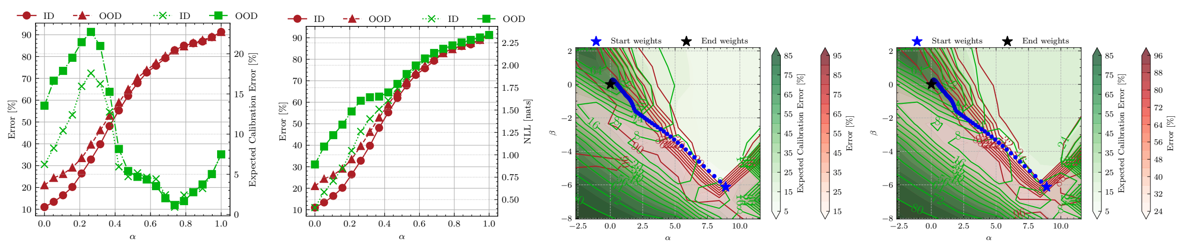

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 14: Activation Additive Gaussian on CIFAR-10. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

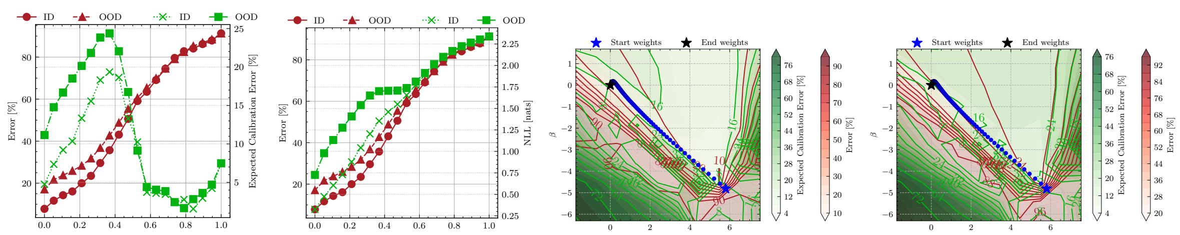

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 15: Activation Dropout on CIFAR-10. *Observations*: Dropout narrowed the gap between ID and OOD results; nevertheless, the shape of the 1D curves is similar to no noise. The trajectories in 2D plots did not seem to converge into a narrow local minimum.

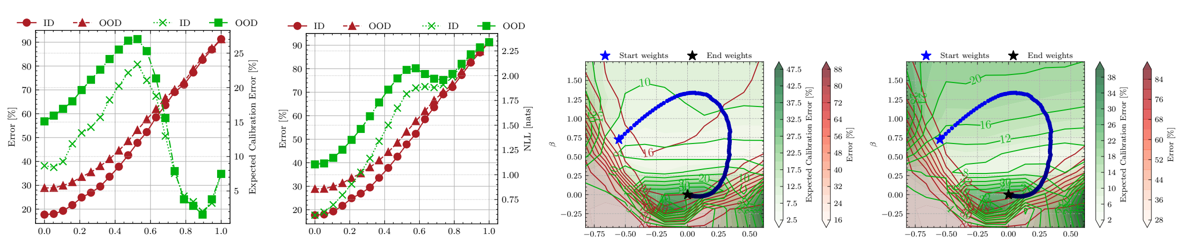

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 16: Gradient Gaussian on CIFAR-10. *Observations*: The 1D and 2D figures changed curvature and shape drastically, and NLL and ECE follow a non-linear pattern. The 2D plots show a circular curvature, perhaps suggesting difficulty in convergence.

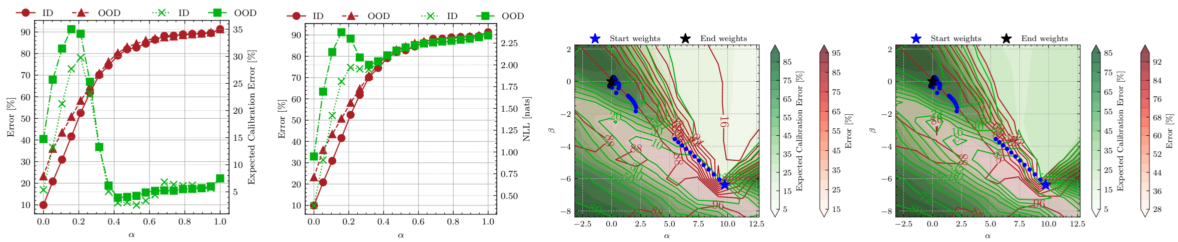

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 17: Model Shrink and Perturb on CIFAR-10. *Observations*: The 1D and 2D figures changed curvature and shape drastically, and all metrics show a non-linear optimisation path as hypothesised. The point cluster around centres created by shrinking and perturbing the weights.

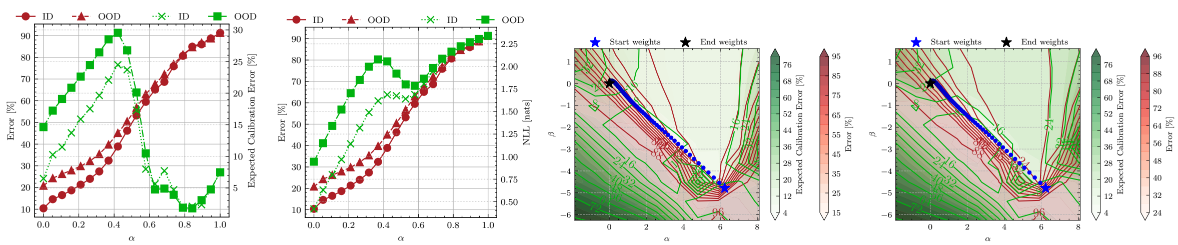

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 18: Weight Additive Gaussian on CIFAR-10. *Observations*: The 1D curves marginally changed their shape. However, the difference between ID and OOD metrics became more profound. The 2D plots suggest that the optimisation was not able to converge.

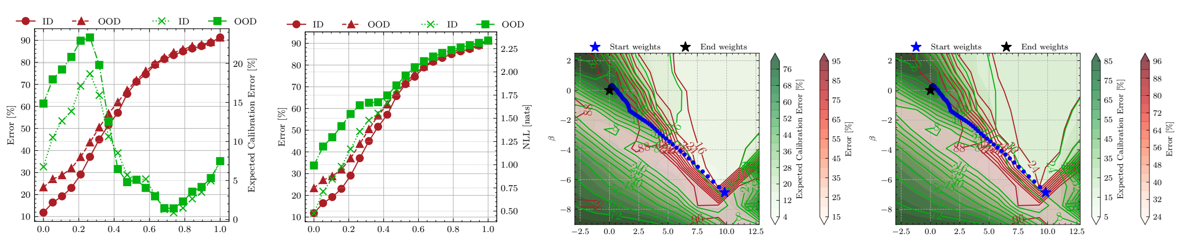

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 19: Weight DropConnect on CIFAR-10. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

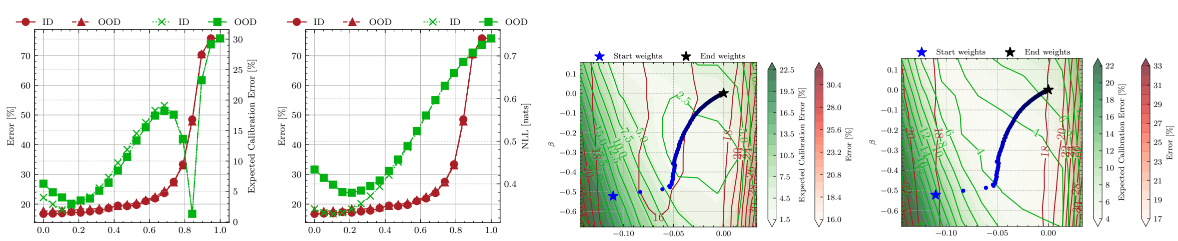

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 20: No noise on Adult.

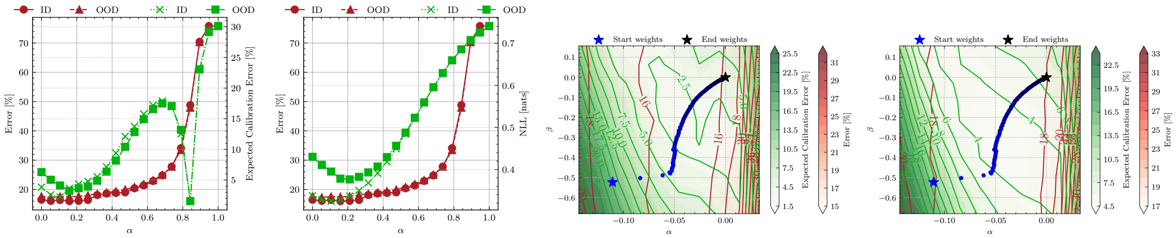

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 21: Input Additive Gaussian on Adult. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

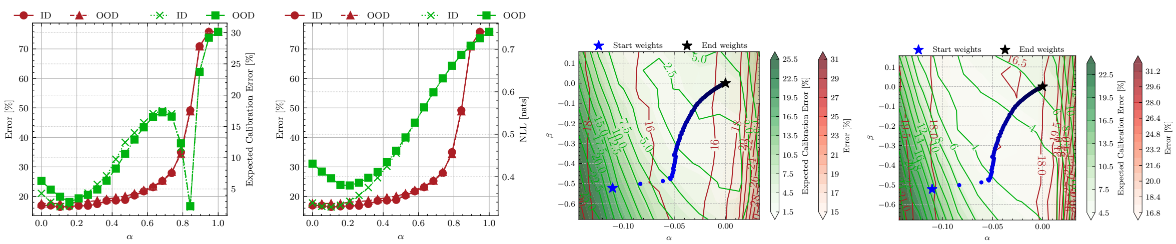

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 22: Input ODS on Adult. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

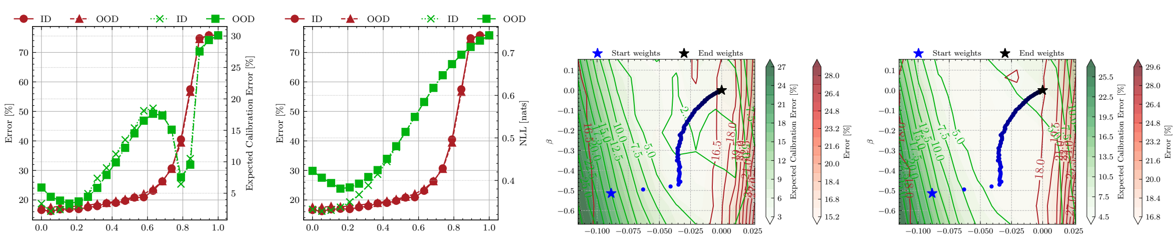

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 23: Input-Target MixUp on Adult. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

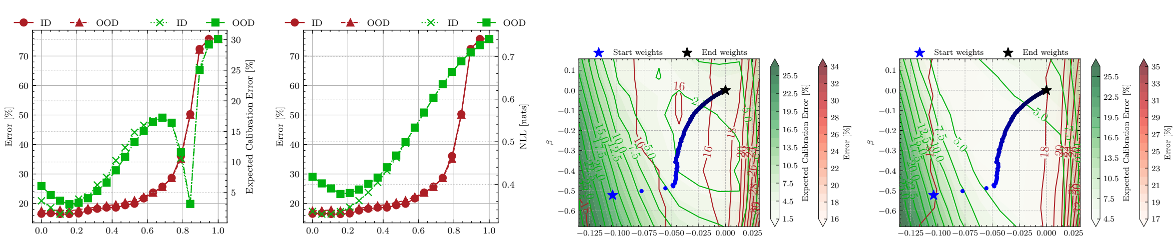

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 24: Target Smoothing on Adult. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

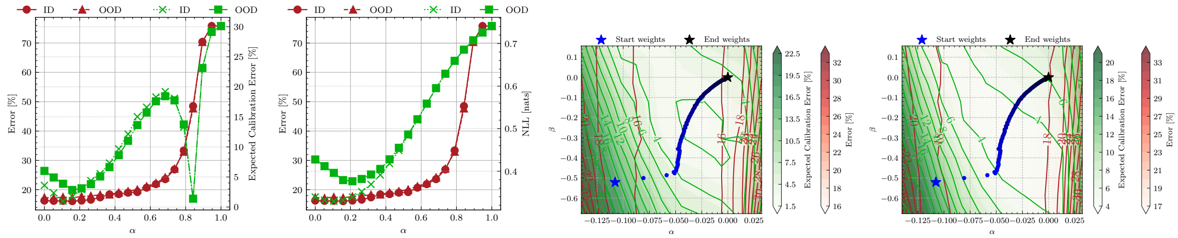

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 25: Activation Additive Gaussian on Adult. *Observations*: Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

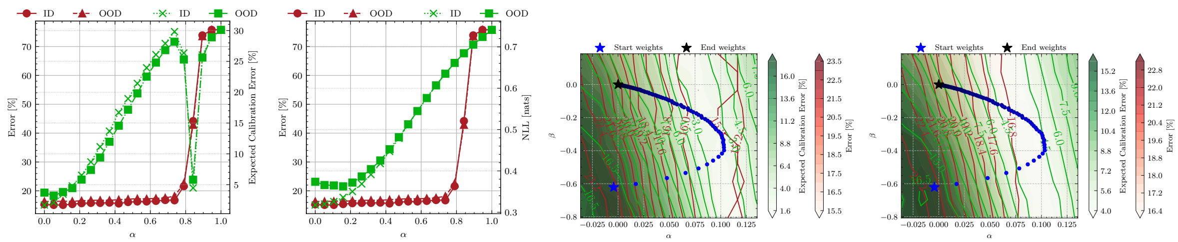

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 26: Activation Dropout on Adult. *Observations*: Changed the ECE curvature and made the NLL plots smoother in the 1D case. In the 2D plots, the ECE and error appear aligned during optimisation.

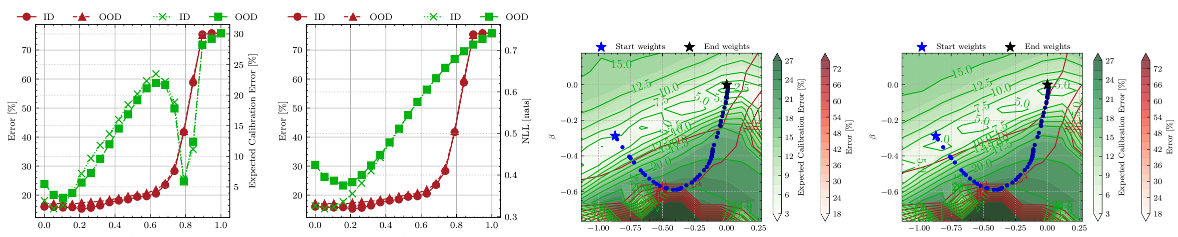

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 27: Gradient Gaussian on Adult. *Observations*: Did not change the smoothness of the 1D curves, but the 2D trajectory appears more exploratory compared to no noise.

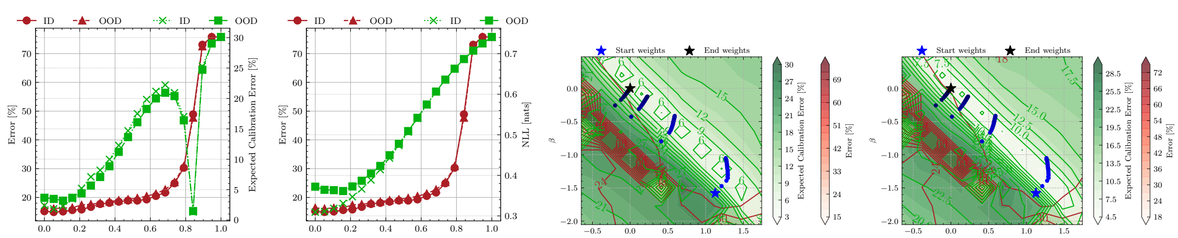

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 28: Model Shrink and Perturb on Adult. *Observations*: Did not change the smoothness of the 1D curves, but the 2D trajectory appears more exploratory compared to no noise.

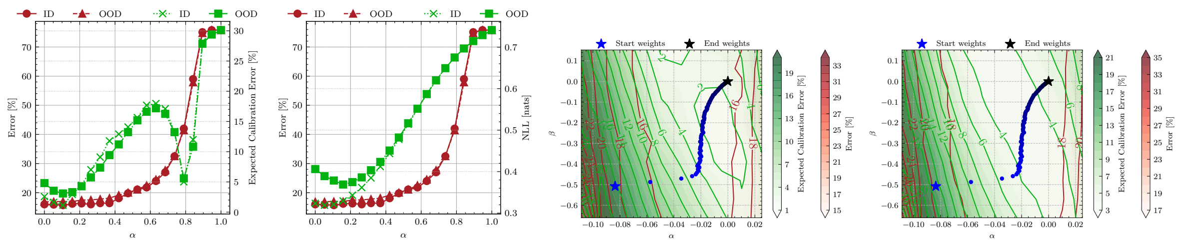

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 29: Weight Additive Gaussian on Adult. Did not change the smoothness of the 1D curves or the 2D metric landscape trajectory compared to no noise.

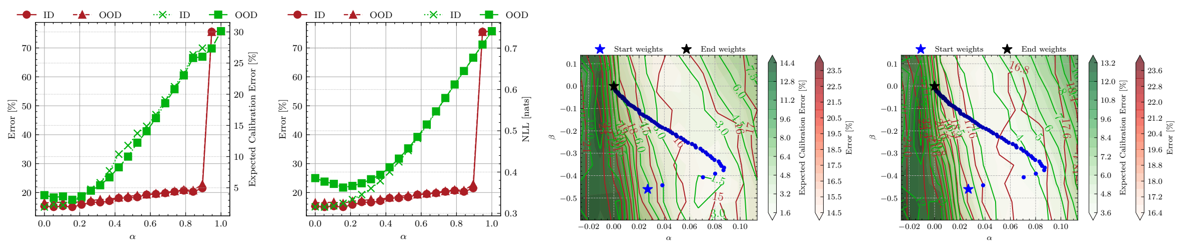

(a) Error and ECE. (b) Error and NLL. (c) Error and ECE on ID. (d) Error and ECE on OOD.

Figure 30: Weight DropConnect on Adult. Changed the ECE curvature and made the NLL plots smoother in the 1D case. In the 2D plots, the ECE and error appear aligned during optimisation.

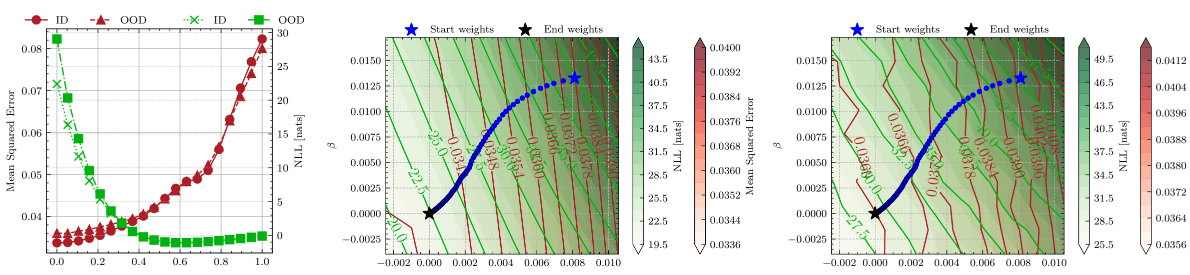

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 31: Input Additive Gaussian on WikiFace. *Observations*: Did not change the smoothness of the 1D curves, or the 2D trajectory.

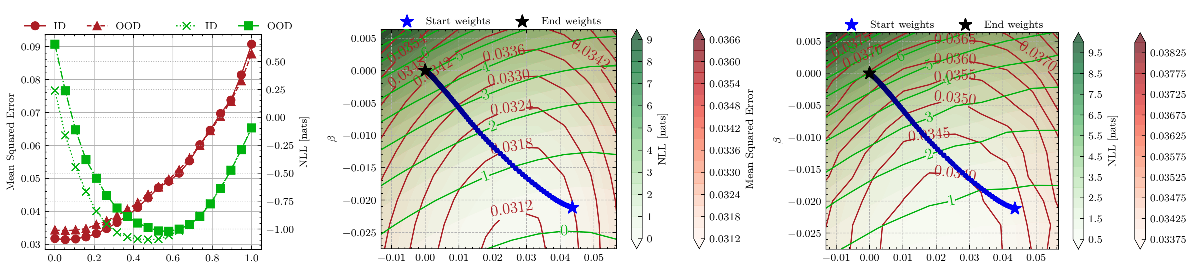

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 32: Input Random Crop, Horizontal Flip on WikiFace. *Observations*: Surprisingly, the NLL starts decreasing compared to MSE as the model is interpolated between the final and the initial model in the 1D plots. The 2D plots demonstrate that the model was able to explore a deeper optimal from the start where NLL was slower to converge than MSE.

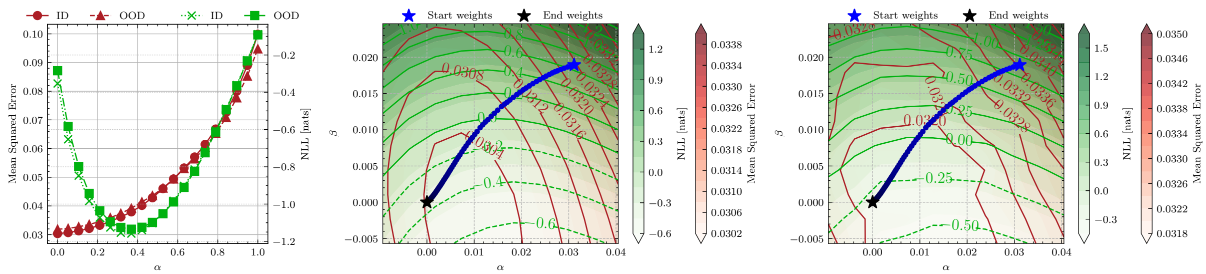

(a) MSE and NLL.

(b) MSE and NLL on ID.

(c) MSE and NLL on OOD.

Figure 33: Input AugMix on WikiFace. *Observations*: Surprisingly, the NLL starts decreasing compared to MSE as the model is interpolated between the final and the initial model in the 1D plots. The 2D plots demonstrate that the model was able to explore a deeper optimal from the start where NLL was slower to converge than MSE, and it did not converge in the optima from the perspective of NLL.

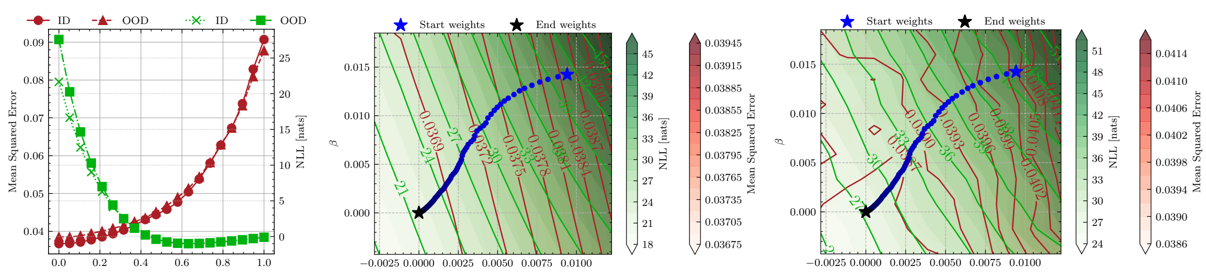

(a) MSE and NLL.

(b) MSE and NLL on ID.

(c) MSE and NLL on OOD.

Figure 34: Input-Target CMixUp on WikiFace. *Observations*: Did not change the smoothness of the 1D curves, or the 2D trajectory appears more exploratory compared to no noise.

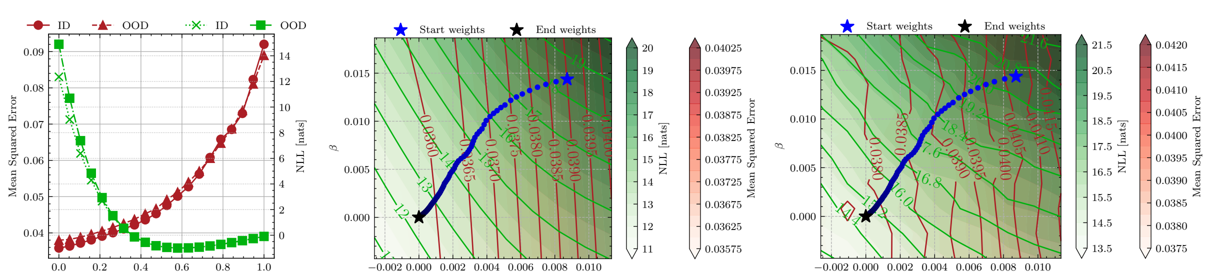

(a) MSE and NLL.

(b) MSE and NLL on ID.

(c) MSE and NLL on OOD.

Figure 35: Activation Additive Gaussian on WikiFace. *Observations*: Did not change the smoothness of the 1D curves, but the 2D trajectory appears more exploratory compared to no noise.

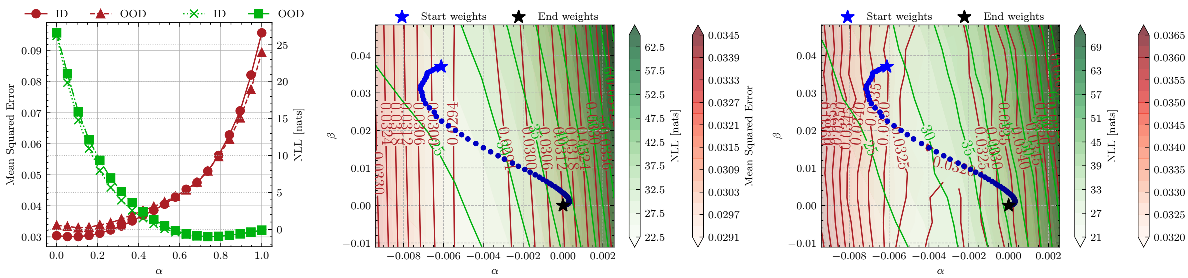

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 36: Gradient Gaussian on WikiFace. *Observations*: Did not change the smoothness of the 1D curves, but the 2D trajectory appears to align MSE and NLL. However, it seems that the optimisation missed a local minimum during training.

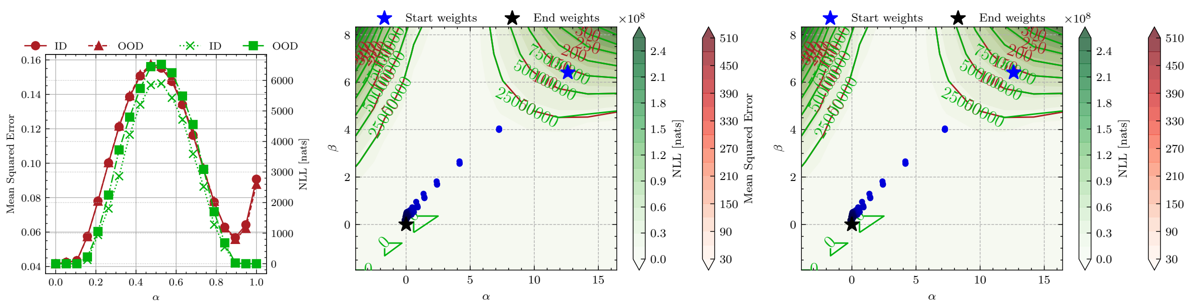

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 37: Model Shrink and Perturb on WikiFace. *Observations*: Due to shrinking and perturbation, the experiment appears to converge in a narrow basin and as seed in the 1D plots, the optimisation was completely non-linear and unrecoverable.

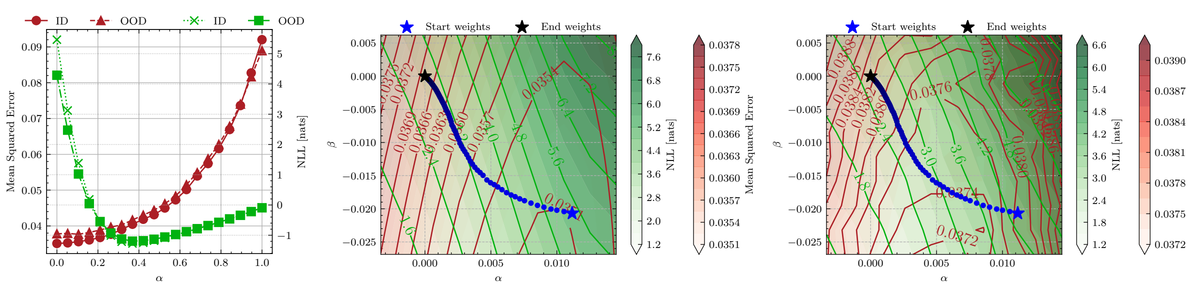

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 38: Weight Additive Gaussian on WikiFace. *Observations*: The 1D curves look similar to no noise, although with respect to a different scale for NLL. The 2D plots explore a similar trajectory to no noise; however, the 2D landscape appears more distorted.

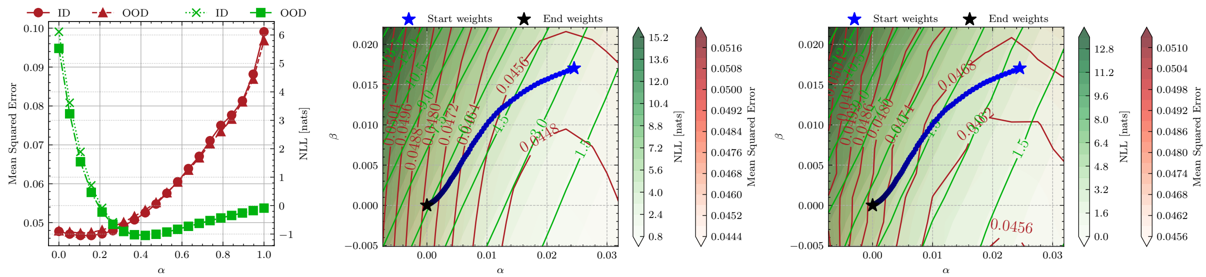

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 39: Weight DropConnect on WikiFace. *Observations*: The 1D curves look similar to no noise, although with respect to a different scale for NLL. The 2D plots explore a similar trajectory to no noise.

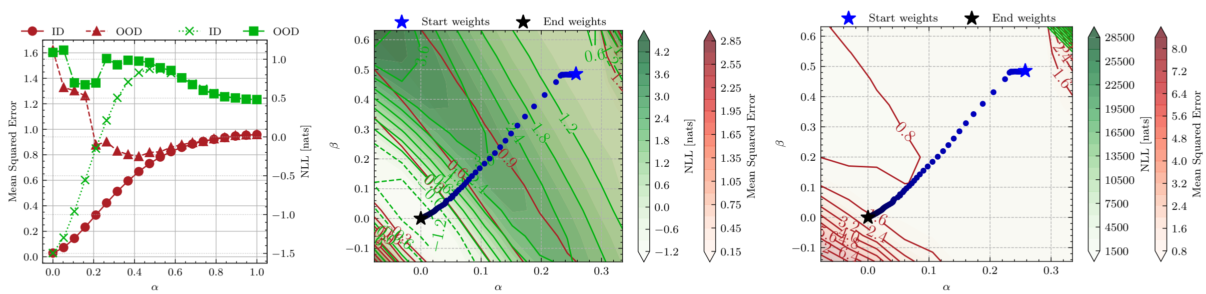

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 40: No noise on Yacht.

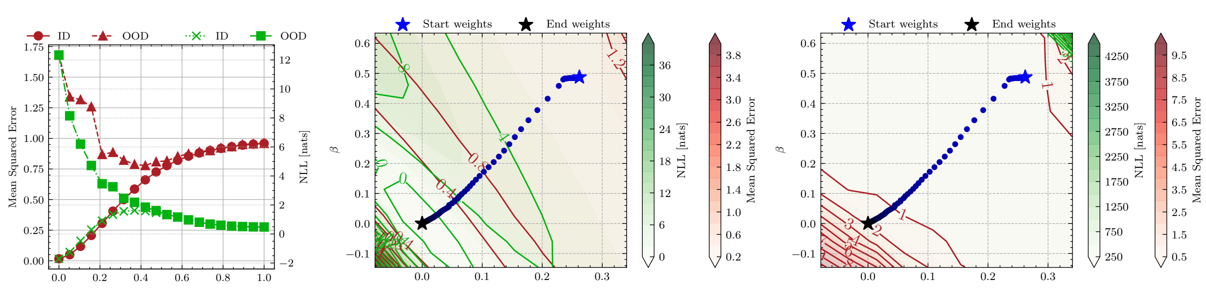

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 41: Input Additive Gaussian on Yacht. *Observations*: While the shape of the 1D curves looks similar to no noise, the MSE and NLL magnitudes are different. The 2D plots demonstrate a wider landscape of feasible solutions than no noise.

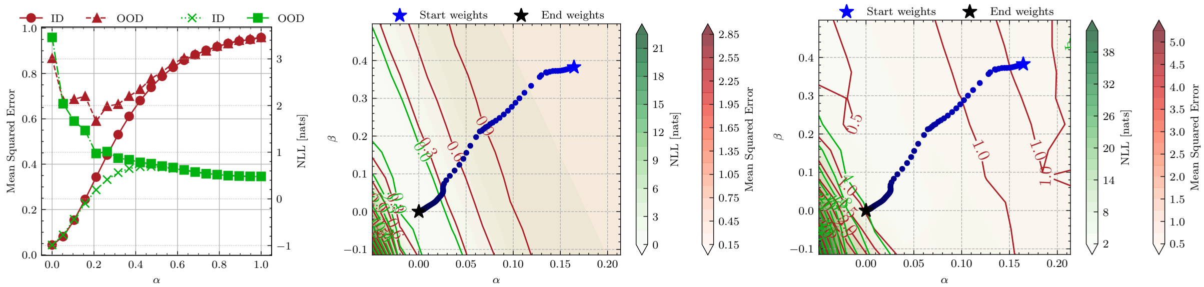

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 42: Input-Target CMixUp on Yacht. *Observations*: While the shape of the 1D curves looks similar to no noise, the MSE and NLL magnitudes are different. The 2D plots demonstrate a wider landscape of feasible solutions than no noise.

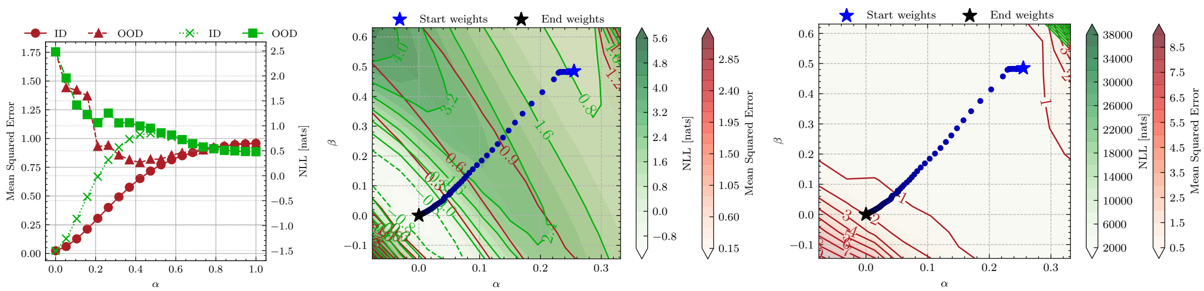

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 43: Activation Additive Gaussian on Yacht. *Observations*: While the shape of the 1D curves looks similar to no noise, the MSE and NLL magnitudes are different. The 2D plots are close to the no-noise ones, showing marginal differences.

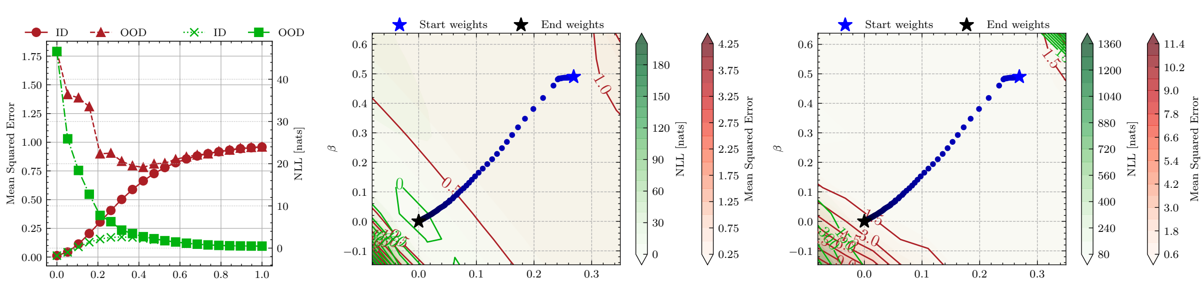

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 44: Activation Dropout on Yacht. *Observations*: Dropout converged in a narrow valley, as demonstrated in the 2D plots, but also in the 1D plot where a small interpolation disturbs especially the OOD performance of the model.

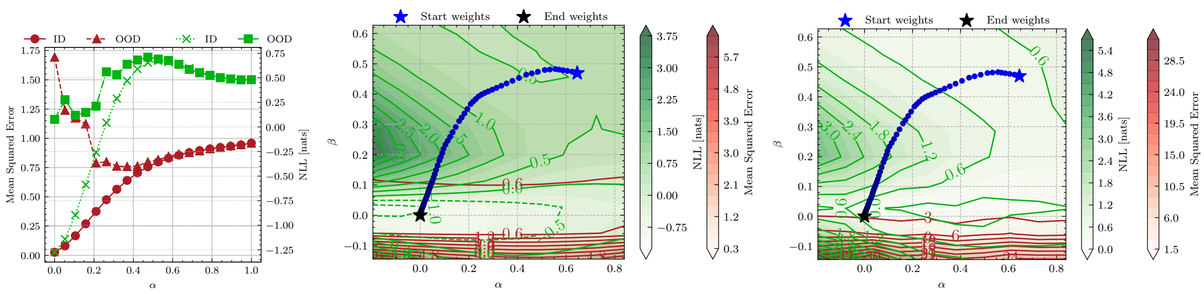

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 45: Gradient Gaussian on Yacht. *Observations*: The 1D curves remained unchanged except for the magnitude of NLL or MSE. Nevertheless, the 2D plots show us that the optimisation trajectory significantly differed from no noise where the landscape of potential optimal solutions was wider.

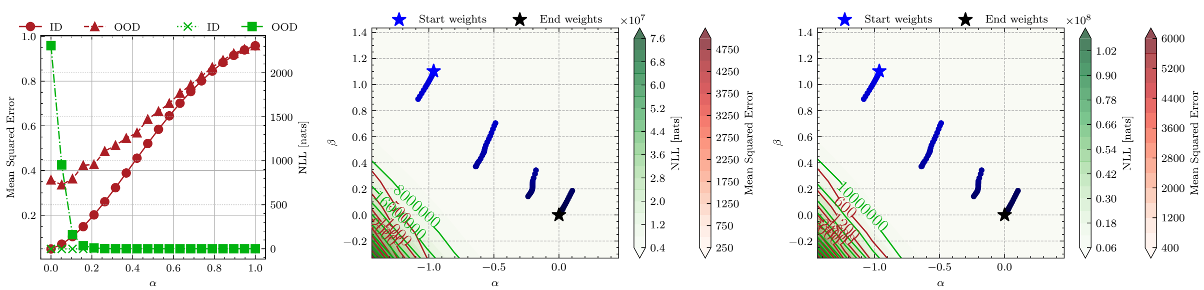

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 46: Model Shrink and Perturb on Yacht. *Observations*: The model jumped between narrow valleys as seed in the 2D plots and the 1D plots show smoother behaviour from the OOD perspective for MSE but no NLL.

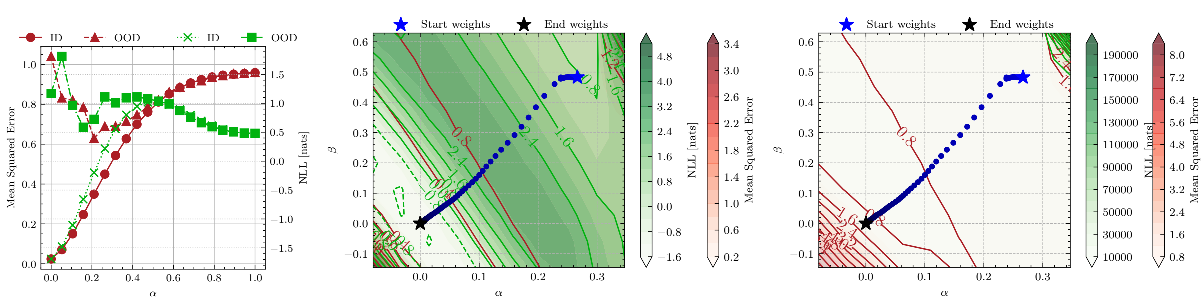

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 47: Weight Additive Gaussian on Yacht. *Observations*: Did not change the smoothness of the 1D curves or the 2D trajectory.

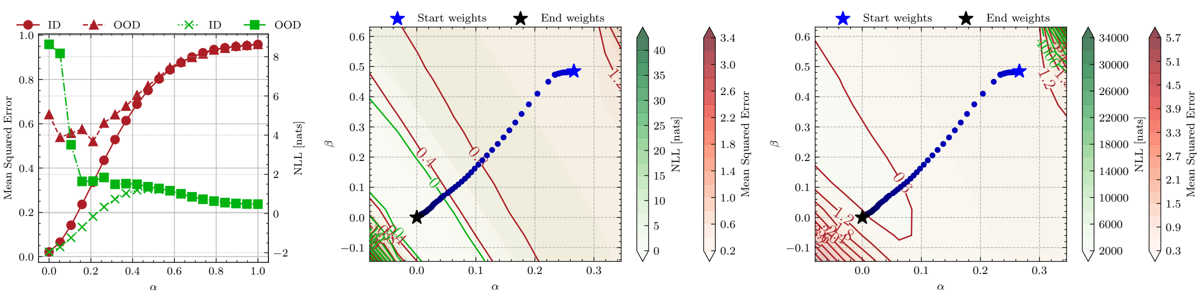

(a) MSE and NLL. (b) MSE and NLL on ID. (c) MSE and NLL on OOD.

Figure 48: Weight DropConnect on Yacht. *Observations*: While the shape of the 1D curves looks similar to no noise, the MSE and NLL magnitudes are different. The 2D plots demonstrate a wider landscape of feasible solutions than no noise.