

Identity Curvature Laplace Approximation for Improved Out-of-Distribution Detection

Maksim Zhdanov
AIRI*, NUST MISIS[†]
Moscow, Russia
zhdanov@airi.net

Stanislav Dereka
T-Bank, MIPT[‡]
Moscow, Russia
s.dereka@tbank.ru

Sergey Kolesnikov
T-Bank*
Moscow, Russia
scitator@gmail.com

Abstract

Uncertainty estimation is crucial in safety-critical applications, where robust out-of-distribution (OOD) detection is essential. Traditional Bayesian methods, though effective, are often hindered by high computational demands. As an alternative, Laplace approximation offers a more practical and efficient approach to uncertainty estimation. In this paper, we introduce the Identity Curvature Laplace Approximation (ICLA), a novel method that challenges the conventional posterior covariance formulation by using identity curvature and optimizing prior precision. This innovative design significantly enhances OOD detection performance on well-known datasets such as CIFAR-10, CIFAR-100, and ImageNet, while maintaining calibration scores. We attribute this improvement to the alignment issues between typical feature embeddings and curvature as measured by the Fisher information matrix. Our findings are further supported by demonstrating that incorporating Fisher penalty or sharpness-aware minimization techniques can greatly enhance the uncertainty estimation capabilities of standard Laplace approximation.

1. Introduction

Deep neural networks have become essential tools across various machine learning domains, including computer vision [6], natural language processing [41], and time-series forecasting [46]. Despite their wide application, these networks often exhibit overconfidence in their predictions [33], posing challenges in safety-critical environments. Uncertainty estimation, which evaluates a neural network’s confidence in its predictions, addresses this issue by providing a measure of trustworthiness. A key application is out-of-distribution (OOD) detection, which identifies samples

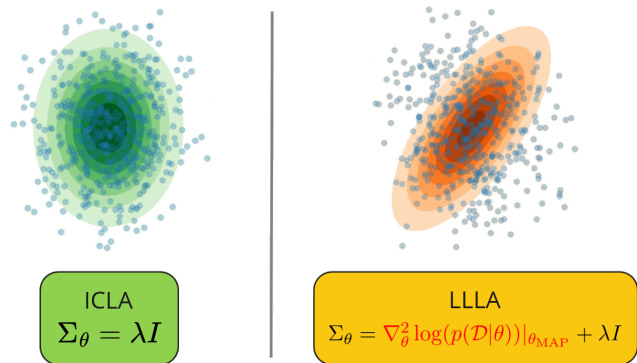


Figure 1. Identity Curvature Laplace Approximation (ICLA), without Hessian computation, produces smoother uncertainty landscapes than standard last-layer Laplace approximation (LLLA), and leads to improved OOD detection performance in real-world scenarios.

that significantly differ from the training data [12], typically resulting in unreliable predictions due to high overconfidence [45].

Bayesian methods [13, 32] have been popular for addressing uncertainty in neural networks [18]. The Laplace approximation [27], a computationally efficient Bayesian technique, approximates the posterior distribution of model weights with a Gaussian centered at the point of maximum likelihood. Efficient variations of this method [35, 36] make Bayesian inference feasible for large networks. For example, the last-layer Laplace approximation [21] reduces computational complexity by focusing on the final layer’s weights.

However, some researchers question the necessity of complex Bayesian methods. Deep ensembles (DE), a straightforward non-Bayesian approach, have shown comparable or superior performance in uncertainty estimation [23], suggesting that simpler methods might be sufficient for practical applications.

In this work, we identify an intrinsic limitation of model

*Work done while at T-Bank

[†]National University of Science and Technology MISIS

[‡]Moscow Institute of Physics and Technology




Figure 2. Uncertainty estimation on a toy binary classification dataset with test outliers for MAP network (*left*), Standard LLLA (*middle*), ICLA (*right*). For each setup, the average predictive uncertainty for outlier samples is displayed in the top-left corner. It can be seen that using identity curvature and prior precision as the source of covariance leads to a wider uncertainty surface while preserving the posterior landscape. Subsequently, it achieves better entropy estimates for outlier data points. See Section 4.1.

curvature in Laplace approximation for uncertainty estimation. We propose the Identity Curvature Laplace Approximation (ICLA¹), which replaces the Hessian with an identity matrix and optimizes prior precision, enhancing OOD detection while maintaining calibration scores. Unlike traditional methods, ICLA eliminates the need to compute the Hessian matrix during inference (see Figure 1). We empirically demonstrate that class separability in the embedding space explains ICLA’s effectiveness. Furthermore, we show that applying a Fisher trace penalty or sharpness-aware minimization during training smooths the model parameter landscape, closing the performance gap between standard Laplace approximation and ICLA.

The contributions of this work can be summarized as follows:

- We introduce ICLA, a simplified Laplace approximation method that optimizes prior precision without requiring Hessian computations at inference. ICLA outperforms standard Laplace methods in OOD detection on CIFAR-10, CIFAR-100, and ImageNet-200 datasets, while maintaining similar calibration metrics.
- We highlight the connection between curvature and class separability in the embedding space, explaining ICLA’s improved performance.
- We demonstrate that employing a Fisher penalty or sharpness-aware minimization techniques during training enhances the smoothness of the model parameters landscape, thereby improving the performance of traditional Laplace approximation methods.

2. Background

Out-of-Distribution Detection. We consider a multi-class classification problem where $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$ represents a dataset consisting of independently and identically distributed (i.i.d.) data points $x_i \in \mathbb{R}^D$ and labels

$y \in \{k_1, \dots, k_C\}$, each depicted as a one-hot vector in a C -dimensional space.

Given the defined multi-class classification problem, we have a particular set of classes y that we want to model. We refer to these classes and the corresponding data samples as *in-distribution* (ID) data. In an open world, however, other groups of classes do not belong to our training data \mathcal{D} . We can say that they come from an *out-of-distribution* (OOD) domain. The task of OOD detection is to measure how well a predictive model can identify whether a data sample at inference phase is from an ID or OOD domain [45].

Calibration. A neural network is considered calibrated when its output probabilities align with its real predictive performance. For instance, if for some set of samples $\{x_{i,i+n}\}$ a model predicts the label y with probability P , then we expect for a fraction of P of all data points from $\{x_{i,i+n}\}$ to have y as a true label. Common calibration measures are the expected calibration error (ECE) [31], negative log-likelihood (NLL) [9] and Brier score [30]. Refer to Appendix D for the definitions.

Bayesian Models. We define a neural network as $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^C$ parameterized by θ , with a likelihood $p(\mathcal{D}|\theta)$, as well as prior $p(\theta)$ and posterior distribution $p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta)$. According to Bayesian principles, the standard training of a neural network is a *maximum a posteriori* (MAP) estimation. This involves finding $\theta_{\text{MAP}} = \arg \max_\theta p(\theta|\mathcal{D})$, which can be further expanded as $\arg \max_\theta \log(p(\mathcal{D}|\theta)) + \log(p(\theta))$.

In practice, the log-likelihood $\log(p(\mathcal{D}|\theta))$ corresponds to cross-entropy loss for classification tasks, and the prior $\log(p(\theta))$ is typically an isotropic Gaussian distribution $\mathcal{N}(\theta|0, \sigma^2 I)$. In neural networks, a Gaussian prior can be achieved by applying weight decay regularization $w(\theta) = \frac{1}{2}\sigma^{-2}\|\theta\|^2$. If the prior is uniform, then training is equivalent to simple *maximum likelihood estimation* (MLE).

Since exact $p(\theta|\mathcal{D})$ posterior estimation is intractable for modern neural networks, many methods were proposed to

¹Source code: <https://github.com/maxnygma/icla>




Figure 3. Uncertainty estimation on a toy sinusoidal regression dataset for MAP network (*left*), Standard LLLA (*middle*), ICLA (*right*). The observed property of Laplace approximation and identity curvature generalizes to regression tasks and does not negatively impact uncertainty estimates. See Section 4.1.

approximate Bayesian inference, including the Monte Carlo dropout [8] and deep ensembles [23, 43]. Unfortunately, when compared to standard training methods, many of them have a significant computational cost.

Laplace Approximation. The simple, yet competitive approximation of the posterior $q(\theta|\mathcal{D})$ can be achieved using *Laplace approximation* [27]. This method computes a second-order expansion of $\log(p(\mathcal{D}|\theta))$ around θ_{MAP} , resulting in a Gaussian approximation to $p(\theta|\mathcal{D})$. $\log(p(\mathcal{D}|\theta))$ decomposes as

$$\begin{aligned} \log(p(\mathcal{D}|\theta)) &\approx \\ &\approx \log(p(\mathcal{D}|\theta_{\text{MAP}})) + \frac{1}{2}(\theta - \theta_{\text{MAP}})^T \Sigma_{\theta} (\theta - \theta_{\text{MAP}}), \end{aligned} \quad (1)$$

$$\Sigma_{\theta} \triangleq \nabla_{\theta}^2 \log(p(\mathcal{D}|\theta))|_{\theta_{\text{MAP}}} + \lambda I, \quad (2)$$

where Σ_{θ} is the Hessian of the log-likelihood with respect to θ evaluated at θ_{MAP} and with prior precision λ times identity matrix I . In practice, λ optimization is achievable using methods such as cross-validation or marginal likelihood [15, 16, 36]. Thus, we approximate $p(\theta|\mathcal{D})$ as $\mathcal{N}(\theta|\theta_{\text{MAP}}, \Sigma_{\theta}^{-1})$.

Hessian Approximations. For modern neural networks, computing the full Hessian is not feasible, since Equation 2 scales quadratically with the number of parameters θ . Fortunately, there are multiple approaches for efficient computation of the Hessian in deep neural networks. The most popular methods include the generalized Gauss-Newton (GGN) matrix [38] and the Fisher information matrix [1], as well as block-diagonal factorizations such as the Kronecker-factored approximate curvature (K-FAC) [29] (Refer to Appendix E for the definitions). There is also empirical Fisher (EF), which uses expectations over samples [2]. It is important to note that Fisher is the same as GGN when utilizing distributions from the exponential family [28], making them both well-justified approximations of the Hessian.

Last-Layer Laplace Approximation. Studies have shown that Bayesian methods applied solely to the final layer can rival the performance of those applied across all layers [21, 39]. In the context of last-layer Laplace approximation (LLLA) [21], the neural network f_{θ} is viewed as $f(x_i) = W\nu(x_i)$, where $W \in \mathbb{R}^{C \times L}$ is the weight matrix of the last layer, and $\nu \in \mathbb{R}^L$ represents the output from the penultimate layer. This results in a covariance $\Sigma_{\theta} \triangleq (\nabla_{\theta}^2 \log(p(\mathcal{D}|\theta))|_{W_{\text{MAP}}} + \lambda I) \in \mathbb{R}^{CL \times CL}$. In case of linearized predictive distribution [16] in the last-layer approximation, the output distribution is given by

$$\begin{aligned} p(f(x_i)|x_i, \mathcal{D}) &= \\ &= \mathcal{N}(f(x_i)|W_{\text{MAP}}\nu(x_i), J(x_i)^T \Sigma^{-1} J(x_i)), \end{aligned} \quad (3)$$

where J is the Jacobian at input x_i . Note that this approach does not require the Monte Carlo integration, which is often inaccurate in case of Laplace approximation [16].

3. Method

We put forward the hypothesis that, in widely adopted neural network models, *curvature may act as a bottleneck in terms of uncertainty estimation capabilities*. In particular, model curvature may conflict with class covariance typically present in classification tasks. Based on our observations, we further propose a simplified approach without these limitations.

Analysis approach. To confirm our hypothesis, we need a way to measure *model curvature* and *class covariance*. We analyze model curvature with the empirical Fisher information matrix, which is given by

$$F \triangleq \sum_{n=1}^N \nabla_{\theta} \log(p_{\theta}(y_n|x_n)) \nabla_{\theta} \log(p_{\theta}(y_n|x_n))^T. \quad (4)$$

We employ the lightweight diagonal representation [4] of the Fisher. This particular form of approximation is chosen for computational accessibility.

To evaluate class covariance we propose a mean class-wise cosine similarity (MCCS) measure:

$$\text{MCCS} = \frac{1}{\tilde{N}} \sum_{\substack{n>m \\ n,m \in 1, \dots, C}} \sum_{\substack{i>j \\ i,j \in 1, \dots, N_C}} \frac{\nu_{ni} \cdot \nu_{mj}}{\|\nu_{ni}\|_2 \|\nu_{mj}\|_2}, \quad (5)$$

where ν_{ni} represents the feature of i -th element of class n , N_C is the number of elements in class, and $\tilde{N} = CN_C(C-1)(N_C-1)/4$.

Our subsequent analysis in Section 5 shows that the Fisher information matrix usually has a long-tailed spectral distribution, where a few eigenvalues are incredibly large [19]. Moreover, we show that LLLA performance degrades on highly separable datasets. To connect the line of thought between class separability, curvature, and long-tailed Fisher eigenvalues distribution, we hypothesize that *penalizing Fisher long-tailness can improve LLLA performance*. Thus, we apply the Fisher penalty in the model training:

$$\text{FP} = \alpha \left\| \frac{1}{M} \sum_{i=1}^M \nabla_{\theta} \log(p_{\theta}(y_i|x_i)) \right\|_2, \quad (6)$$

where M is the batch size of $(x_i, y_i) \sim \mathcal{D}$ and α is a hyperparameter controlling the regularization strength. In essence, this penalty acts as model curvature regularization, which leads to a *flat minima* by explicitly lowering Fisher matrix long-tailness [7, 14, 17, 20, 34]. As an alternative to the Fisher penalty, we also apply Adaptive Sharpness-Aware Minimization (ASAM) [22], a widely adopted optimization technique for achieving a flatter curvature landscape. Our analysis shows that by applying these techniques (i.e. reducing Fisher long-tailness), LLLA performance can be improved on highly separable datasets, which confirms our intuition on model curvature and class separability.

Identity Curvature Laplace Approximation. Putting all the empirical observations together we come up with a simplification of LLLA, devoid of the above problems. Consider the covariance matrix Σ_{θ} from Equation 2, which is the sum of the Hessian and the product of prior precision with the identity matrix.

We propose constructing a posterior approximation based exclusively on identity curvature and optimized prior precision:

$$\Sigma_{\theta} = \lambda I. \quad (7)$$

In this study, we name the proposed method *Identity Curvature Laplace Approximation (ICLA)*. This approach doesn't utilize Hessian approximations and is not subject to the long-tailness problem by design. In subsequent experiments, we empirically show that ICLA yields

Algorithm 1 ICLA training

Input: Dataset $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$, neural network f , functions `fit_hessian_approximation` and `marginal_likelihood` from [2].

Output: Prior precision λ for Equation 7.

- 1: Initialize $\lambda \leftarrow 1$, Hessian $\mathcal{H} \leftarrow \mathbf{0}$
 - 2: $\mathcal{H} \leftarrow \text{fit_hessian_approximation}(f, \mathcal{D})$
 - 3: $\lambda \leftarrow \text{marginal_likelihood}(f, \mathcal{H})$
 - 4: $\mathcal{H} \leftarrow \mathbf{0}$
 - 5: **return** λ, \mathcal{H}
-

a broader uncertainty landscape and increased entropy for out-of-distribution samples, which facilitates OOD detection. Based on our observations in Section 5, *ICLA provides the highest performance boosts on highly separable datasets*.

Our Laplace approximation implementation is based on [2], and the algorithm for ICLA training is available in Algorithm 1. The function `fit_hessian_approximation` computes diagonal Fisher Hessian approximation for each batch in the validation dataset and returns the mean Hessian value of all batches. The function `marginal_likelihood` performs prior precision optimization (Equation 2) with marginal likelihood method [2], the algorithm is presented in Appendix A. Moreover, step 2 of the Algorithm 1 can be bypassed, i.e. the Hessian matrix is not fitted and initialized with zeros. We call this approach $\text{ICLA}_{\text{zero}}$. For the full details of implementation, refer to the released ICLA code implementation. We also report ICLA computational performance in Appendix B.

4. Experiments

4.1. Synthetic Examples

Half-Moons. To visually examine the changes ICLA brings to uncertainty estimates, we use a half-moons binary classification dataset with an additional 10 test outliers. In this experiment, we use a simple 2-layer MLP with ReLU activations and 20 hidden units. We compare the uncertainty landscapes produced by 3 methods: MAP, LLLA, and ICLA. The results are presented in Figure 2. We demonstrate that ICLA preserves the general structure of the predictive distribution and gives more strict and precise uncertainty estimates for outliers. In addition, one can see that ICLA's uncertainty surface leads to higher outlier entropy scores, which is crucial in OOD detection task.

Sinusoidal Regression. It is essential to make sure that our observations are not limited to toy classification settings and are capable of generalizing to synthetic regression.




Figure 4. Calibration scores. Expected calibration error (ECE), negative log-likelihood (NLL), and Brier score for CIFAR-10 (left), CIFAR-100 (middle) and ImageNet-200 (right) for different Laplace approximations. Identity curvature does not detrimentally affect model calibration. See Section 4.3 for details.

Table 1. OOD detection results for different Laplace approximations. All approximations are tested on multiple OOD sources. The mean and standard deviations’ AUROC values are reported. The best results among different methods are highlighted.

Method	CIFAR-10		CIFAR-100		ImageNet-200	
	Near OOD	Far OOD	Near OOD	Far OOD	Near OOD	Far OOD
LLA (GGN)	88.94 ± 0.30	91.54 ± 0.36	81.44 ± 0.09	80.00 ± 0.45	81.84 ± 0.10	89.39 ± 0.15
LLA (EF)	89.54 ± 0.36	92.09 ± 0.30	81.66 ± 0.10	80.17 ± 0.83	81.86 ± 0.11	89.45 ± 0.14
LLA (K-FAC)	88.02 ± 0.28	90.73 ± 0.52	80.30 ± 0.13	77.77 ± 0.53	81.73 ± 0.10	88.94 ± 0.14
ICLA	90.01 ± 0.21	92.50 ± 0.38	81.45 ± 0.10	80.79 ± 0.46	81.88 ± 0.10	89.45 ± 0.14

We visually demonstrate this in Figure 3 using a sinusoidal regression dataset. For this experiment, we use the same network as in the previous setup and compare the same set of methods: MAP, LLLA and ICLA. We show that ICLA produces uncertainty estimates similar to LLLA, which involves Hessian computation.

4.2. OOD Detection

For OOD detection experiments, we employ the well-established OpenOOD v1.5 [45] benchmark, which features the most state-of-the-art OOD detection methods. It uses CIFAR-10, CIFAR-100, and ImageNet-200 (a subset of ImageNet-1K) [3] as in-distribution datasets and evaluates methods against carefully chosen OOD datasets. It has “Near OOD” and “Far OOD” parts, where the former contains OOD datasets semantically more similar to ID and the latter contains OOD datasets of completely different categories from ID data. For all of our datasets, we use ResNet-18 [10] trained with the SGD optimizer with a momentum of 0.9. The initial learning rate is set to 0.1 and then reduced to 10^{-6} with the Cosine Annealing scheduler [26]. Training lasts 100 epochs on CIFAR-10 and CIFAR-100 and 90 epochs on ImageNet-200. Additionally, we perform a multi-seed evaluation with 3 random seeds and report the mean and standard deviation values for each experiment to obtain a fair comparison. The predictive distribution for OOD samples is obtained in the same manner as for estimation of ID uncertainty.

Laplace Approximations. We compare standard LLLA with GGN, EF, and K-FAC against ICLA in terms of OOD detection AUROC [45]. We tune prior precision for all of our methods via the marginal likelihood method [15]. The results are presented in Table 1. It can be seen from the results that ICLA achieves better performance on OOD detection problems than other Laplace approximation variations in 5 of 6 experiments.

Non-Bayesian Approaches. In addition to Laplace approximations, we compare ICLA with popular post-hoc non-Bayesian OOD detection methods from OpenOOD, including Maximum Softmax Probability (MSP) [11], ODIN [25], ReAct [40], VIM [42] SHE [44], and ASH [5]. The comparison is shown in Table 2. ICLA outperforms the mentioned methods on “Near OOD” CIFAR-10 and CIFAR-100 and achieves competitive performance on other setups.

4.3. Accuracy and Calibration

Employing identity curvature in Laplace approximation does not affect the in-distribution classification accuracy, as the method maintains an optimal solution around θ_{MAP} [21]. Most importantly, it does not significantly change in-distribution probabilities, impacting calibration scores (see Figure 4). To compare calibrations, we use the negative log-likelihood (NLL), expected calibration error (ECE) and Brier score. ICLA maintains calibration metrics relatively similar to more computationally complex versions of Laplace approximations.

Table 2. OOD detection results against non-Bayesian methods. All methods are trained on the ID dataset and tested on multiple OOD sources. The mean and standard deviations’ AUROC values are reported. The best results among different methods are highlighted and the 2nd best are underlined.

Method	CIFAR-10		CIFAR-100		ImageNet-200	
	Near OOD	Far OOD	Near OOD	Far OOD	Near OOD	Far OOD
MSP	88.03 ± 0.25	90.73 ± 0.43	80.27 ± 0.11	77.76 ± 0.44	83.34 ± 0.06	90.13 ± 0.09
ODIN	82.87 ± 1.85	87.96 ± 0.61	79.90 ± 0.11	79.28 ± 0.21	80.27 ± 0.08	91.71 ± 0.19
ReAct	87.11 ± 0.61	90.42 ± 1.41	80.77 ± 0.05	80.39 ± 0.49	81.87 ± 0.98	<u>92.31 ± 0.56</u>
VIM	88.68 ± 0.28	93.48 ± 0.24	74.98 ± 0.13	<u>81.70 ± 0.62</u>	78.68 ± 0.24	91.26 ± 0.19
SHE	81.54 ± 0.51	85.32 ± 1.43	78.95 ± 0.18	76.92 ± 1.16	80.18 ± 0.25	89.81 ± 0.61
ASH	75.27 ± 1.88	78.49 ± 2.31	78.20 ± 2.21	80.58 ± 2.56	79.38 ± 1.93	92.74 ± 1.91
ICLA	<u>90.01 ± 0.21</u>	92.50 ± 0.38	81.45 ± 0.10	80.79 ± 0.46	<u>81.88 ± 0.10</u>	89.45 ± 0.14
ICLA_{zero}	90.56 ± 0.23	<u>93.18 ± 0.45</u>	<u>81.38 ± 0.28</u>	82.49 ± 0.60	80.70 ± 0.10	89.82 ± 0.11

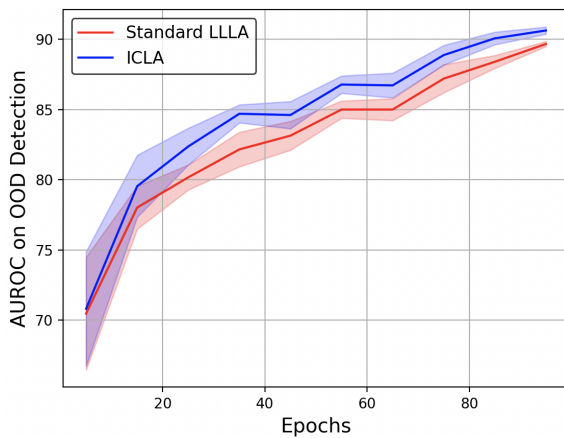


Figure 5. Epoch-wise comparison between standard LLLA and ICLA for OOD detection. Based on this experiment, we can conclude that the observed phenomenon is unrelated to the model training stage. See Section 5.1 for details.

5. Analysis

To understand the roots of the efficacy of ICLA, we apply our analysis approach from Section 3. In our work, we aim to answer the following questions:

1. Does the training stage have an impact on the performance of common Laplace approximation and ICLA?
2. What is the connection between the structure of the Fisher matrix (i.e. model curvature) and the data it is used to model?
3. Can we explain the outstanding results of ICLA by linking them to the separability of classes in the embedding space?




Figure 6. Diagonal empirical Fisher eigenvalues. Only a small percentage of eigenvalues represent large values, leading us to conclude that the spectral distribution has a long tail. Index stands for the *i*-th eigenvalue. See Section 5.2 for details.

4. How do the smoother solutions (with lowered long-tailness), made possible by applying the Fisher penalty, change the comparison between LLLA and ICLA?

5.1. Training Stage Independence

At first thought, it might seem like the performance gap between LLLA and ICLA might be connected to the training stage, since curvature continuously changes throughout the training process. However, our empirical observations suggest that the phenomena is not linked to the training stage of a model. In Figure 5, we build approximations for different training epochs and show that even at the start of training, there is a positive performance gap that persists between the standard last-layer Laplace approximation and ICLA.




Figure 7. MCCS and performance gap comparison for CIFAR-10, CIFAR-100, C100-5-SP and C100-5-NSP. It can be seen that CIFAR-10 and C100-5-SP have a lower MCCS and larger gap, while CIFAR-100 and C100-5-NSP have a higher MCCS and lower gap. This points to a correlation between feature separability and Laplace approximation performance. See Section 5.3 for details.

5.2. Fisher Eigenvalues

Firstly, we examine the eigenvalues λ_i of empirical Fisher used in LLLA. We observe that almost all eigenvalues are small, and only a tiny part of them exhibits extremely large values, demonstrating clear long-tailness. We show the spectral distributions in Figure 6. Our observation perfectly aligns with the past research on analyzing the curvature of neural networks [19, 24, 37]. This is a crucial property to consider, as matrices with long-tailed eigenvalues are used to model data covariance.

5.3. Feature Embeddings

As one might recall from Equation 3, the curvature information directly influences resulting uncertainty through covariance. This can be seen since Σ^{-1} holds the *congruence relation* to the resulting covariance of output distribution, i.e., the covariance of output distribution and Fisher represent the same linear transformation, but in different bases.

Naturally, this raises the following question: *How does long-tailed curvature align with class covariance?* To answer this, we select 4 datasets with different inter-class separability: **CIFAR-10**, **CIFAR-100**, **C100-5-SP**. We use 5 classes from different super classes of CIFAR-100, namely “flowers”, “food containers”, “fruit and vegetables”, “household electrical devices” and “insects” (more separable case). From **C100-5-NSP**, we use 5 classes from the “trees” super class (less separable case).

We use Equation 5 as a measure of similarity and sepa-




Figure 8. Data separability experiment. Gradually moving from a separable (C100-5-SP) to less separable dataset (C100-5-NSP), we observe how standard LLLA begins to estimate uncertainty better than its alternative with identity curvature. This can be seen as a sign of high data separability and curvature information being linked. See Section 5.4 for details.

rability between clusters of classes in the embedding space and compare it to the performance gap in OOD detection between LLLA and ICLA. The performance gap is calculated simply as the average of “Near OOD” and “Far OOD” performance differences between ICLA and LLLA.

Our observations in Figure 7 show a connection between class separability and the informativeness of curvature about data covariance. Feature embeddings illustrating separability are in Appendix F.

We conducted another experiment to explore the trend between LLLA performance and data separability. Starting with C100-5-SP, we trained 5 models, progressively adding one “trees” super-class and replacing a C100-5-SP class. The first model used C100-5-SP, and the last used C100-5-NSP. This setup highlights the relationship between separability changes and performance gaps for ICLA and LLLA. As shown in Figure 8, moving from the most to least separable cases improves LLLA over ICLA, demonstrating the role of data separability in identity curvature efficacy.

5.4. Smooth Curvature Gap Reduction

Previously, we demonstrated the connection between data separability, curvature, and uncertainty estimation. Let us take a final step and connect the line between the performance of LLLA and smoothness.

Fisher Trace Penalty. As was shown before, Fisher with long-tailed spectral distribution is not informative for modeling class covariance, which leads to poor LLLA performance. Thus, for setups with highly separable embeddings, we would prefer the curvature to be smoother, i.e. lower Fisher long-tailness.

Using the CIFAR-100 and C100-5-SP datasets with




Figure 9. Effects of Fisher penalty. (Left) OOD detection performance gap for ICLA and LLLA. (Middle) OOD detection AUROC for the C100-5-SP dataset. (Right) OOD detection AUROC for the CIFAR-100 dataset. As regularization increases, the gap vanishes. While the Fisher penalty improves both LLLA and ICLA, the effect on LLLA is more visible. See Section 5.4 for details.

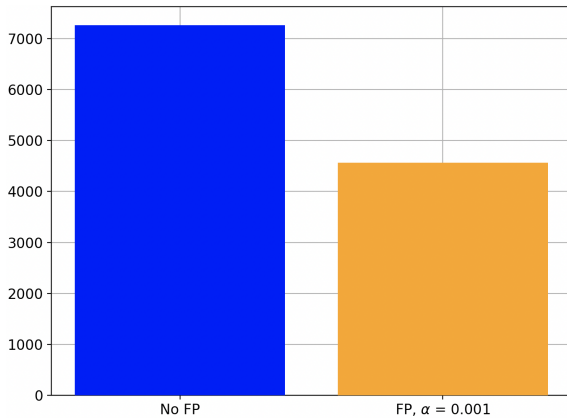


Figure 10. Mean eigenvalues of Fisher for a model trained on CIFAR-10 without the Fisher penalty and a model trained with the Fisher penalty, $\alpha = 0.001$. The model trained with the Fisher penalty achieves a flatter solution. See Section 5.4 for details.

lower separability (Section 5.3), we apply the Fisher penalty (Equation 6) while progressively increasing α , affecting LLLA and ICLA performance. Figure 9 illustrates the vanishing gap between LLLA and ICLA, highlighting the link between curvature and separability. Figure 10 shows the mean Fisher eigenvalues for unregularized and Fisher-penalty-trained models, demonstrating that the penalty reduces eigenvalue magnitude, leading to flatter curvature.

Sharpness-Aware Minimization. In the context of our study, it can be applied for the same reason as the Fisher Penalty, as we hypothesize that achieving flatter solutions should reduce the OOD detection performance gap. We employ training with ASAM for both LLLA and ICLA and compare its performance with Fisher Penalty. The results are presented in Table 3. Our observations show that applying ASAM also reduces the performance gap.

The results for the Fisher Penalty and ASAM confirm

Table 3. OOD detection results with Adaptive Sharpness-Aware Minimization (ASAM) for the CIFAR-10 dataset.

	Near OOD	Far OOD
LLA+FP	89.37	91.91
LLA+ASAM	91.92	93.88
Gap with FP	0.62	0.60
Gap with ASAM	0.40	0.10

that the performance gap between LLLA and ICLA is curvature-related. Highly separable data leads to high curvature (long-tailed Fisher), and achieving a flatter curvature improves LLLA performance, reducing the gap with ICLA.

6. Conclusion

In this work, we demonstrated the effectiveness of Identity Curvature Laplace Approximation (ICLA) in improving Laplace approximation for out-of-distribution (OOD) detection. Our empirical results indicate that ICLA significantly enhances OOD detection performance on the CIFAR-10, CIFAR-100, and ImageNet-200 datasets, while maintaining adequate calibration scores. We established a link between ICLA performance and class separability in the embedding space. Our findings suggest that improving the OOD detection capability of Laplace approximation can have significant implications for developing safe and deployable AI systems.

Limitations. While our study suggests that using empirical Fisher (EF) Hessian approximation to study the curvature is inconsequential for the obtained results, more work can be done to confidently confirm it.

References

- [1] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998. 3

- [2] Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux—effortless Bayesian deep learning. *Advances in Neural Information Processing Systems*, 34:20089–20103, 2021. 3, 4
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [4] John Denker and Yann LeCun. Transforming neural-net output levels to probability distributions. *Advances in neural information processing systems*, 3, 1990. 3
- [5] Andrija Djurisic, Nebojsa Bozanic, Arjun Ashok, and Rosanne Liu. Extremely simple activation shaping for out-of-distribution detection. *arXiv preprint arXiv:2209.09858*, 2022. 5
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [7] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017. 4
- [8] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. 3
- [9] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [11] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016. 5
- [12] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018. 1
- [13] Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13, 1993. 1
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 9(1):1–42, 1997. 4
- [15] Alexander Immer, Matthias Bauer, Vincent Fortuin, Gunnar Rätsch, and Khan Mohammad Emtiyaz. Scalable marginal likelihood estimation for model selection in deep learning. In *International Conference on Machine Learning*, pages 4563–4573. PMLR, 2021. 3, 5
- [16] Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural nets via local linearization. In *International conference on artificial intelligence and statistics*, pages 703–711. PMLR, 2021. 3
- [17] Stanislaw Jastrzebski, Devansh Arpit, Oliver Astrand, Giancarlo B Kerg, Huan Wang, Caiming Xiong, Richard Socher, Kyunghyun Cho, and Krzysztof J Geras. Catastrophic Fisher explosion: Early phase Fisher matrix impacts generalization. In *International Conference on Machine Learning*, pages 4772–4784. PMLR, 2021. 4
- [18] Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun. Hands-on Bayesian neural networks—a tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2):29–48, 2022. 1
- [19] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. Universal statistics of Fisher information in deep neural networks: Mean field approach. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1032–1041. PMLR, 2019. 4, 7
- [20] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. 4
- [21] Agustinus Kristiadi, Matthias Hein, and Philipp Hennig. Being Bayesian, even just a bit, fixes overconfidence in relu networks. In *International conference on machine learning*, pages 5436–5446. PMLR, 2020. 1, 3, 5
- [22] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, pages 5905–5914. PMLR, 2021. 4
- [23] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017. 1, 3
- [24] Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–50. Springer, 2002. 7
- [25] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017. 5
- [26] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5
- [27] David JC MacKay. A practical Bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992. 1, 3
- [28] James Martens. New insights and perspectives on the natural gradient method. *The Journal of Machine Learning Research*, 21(1):5776–5851, 2020. 3
- [29] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR, 2015. 3
- [30] Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. *Advances in Neural Information Processing Systems*, 34:15682–15694, 2021. 2

- [31] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using Bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29, 2015. 2
- [32] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012. 1
- [33] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. 1
- [34] Matthew D Norton and Johannes O Royset. Diametrical risk minimization: Theory and computations. *Machine Learning*, pages 1–19, 2021. 4
- [35] Hippolyt Ritter, Aleksandar Botev, and David Barber. Online structured Laplace approximations for overcoming catastrophic forgetting. *Advances in Neural Information Processing Systems*, 31, 2018. 1
- [36] Hippolyt Ritter, Aleksandar Botev, and David Barber. A scalable Laplace approximation for neural networks. In *6th International Conference on Learning Representations, ICLR 2018-Conference Track Proceedings*, volume 6. International Conference on Representation Learning, 2018. 1, 3
- [37] Levent Sagun, Utku Evci, V Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017. 7
- [38] Nicol N Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural computation*, 14(7):1723–1738, 2002. 3
- [39] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nandathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. Scalable Bayesian optimization using deep neural networks. In *International conference on machine learning*, pages 2171–2180. PMLR, 2015. 3
- [40] Yiyu Sun, Chuan Guo, and Yixuan Li. React: Out-of-distribution detection with rectified activations. *Advances in Neural Information Processing Systems*, 34:144–157, 2021. 5
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1
- [42] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4921–4930, 2022. 5
- [43] Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020. 3
- [44] Jinsong Zhang, Qiang Fu, Xu Chen, Lun Du, Zelin Li, Gang Wang, Shi Han, Dongmei Zhang, et al. Out-of-distribution detection based on in-distribution data patterns memorization with modern hopfield energy. In *The Eleventh International Conference on Learning Representations*, 2022. 5
- [45] Jingyang Zhang, Jingkang Yang, Pengyun Wang, Haoqi Wang, Yueqian Lin, Haoran Zhang, Yiyu Sun, Xuefeng Du, Kaiyang Zhou, Wayne Zhang, et al. Openood v1. 5: Enhanced benchmark for out-of-distribution detection. *arXiv preprint arXiv:2306.09301*, 2023. 1, 2, 5
- [46] Michael Zhang, Khaled K Saab, Michael Poli, Tri Dao, Karan Goel, and Christopher Ré. Effectively modeling time series with simple discrete state spaces. *arXiv preprint arXiv:2303.09489*, 2023. 1