
Revisiting One-vs-All Classifiers for Predictive Uncertainty and Out-of-Distribution Detection in Neural Networks

Shreyas Padhy^{1,2} Zachary Nado¹ Jie Ren¹ Jeremiah Liu¹ Jasper Snoek¹ Balaji Lakshminarayanan¹

Abstract

Accurate estimation of predictive uncertainty in modern neural networks is critical to achieve well calibrated predictions and detect out-of-distribution (OOD) inputs. The most promising approaches have been predominantly focused on improving model uncertainty (e.g. deep ensembles and Bayesian neural networks) and post-processing techniques for OOD detection (e.g. ODIN and Mahalanobis distance). However, there has been relatively little investigation into how the parametrization of the probabilities in discriminative classifiers affects the uncertainty estimates, and the dominant method, softmax cross-entropy, results in misleadingly high confidences on OOD data and under covariate shift. We investigate alternative ways of formulating probabilities using (1) a one-vs-all formulation to capture the notion of “none of the above”, and (2) a distance-based logit representation to encode uncertainty as a function of distance to the training manifold. We show that one-vs-all formulations can improve calibration on image classification tasks, while matching the predictive performance of softmax without incurring any additional training or test-time complexity.

1 Introduction

In recent years, the calibration of deep learning models used for predictive discrimination tasks has become very important, especially in the domains of healthcare (Esteva et al., 2017; Dusenberry et al., 2020), self-driving vehicles (Bojarski et al., 2016), and more generally, AI safety (Amodei et al., 2016). Estimating the *epistemic* uncertainty is especially valuable, as it captures the model’s lack of knowledge about data. Recent work in the field has dealt with three different paradigms of measuring predictive un-

certainty, namely (1) in-distribution calibration of models measured on an i.i.d. test set, (2) robustness under dataset shift, and (3) anomaly/out-of-distribution (OOD) detection, which measures the ability of models to assign low confidence predictions on inputs far away from the training data. Unfortunately, it has been recognized that vanilla predictive models fail to provide accurate quantification of predictive uncertainty (Guo et al., 2017; Hein et al., 2019). In recent years, many methods have been proposed to tackle this, including deep ensembles (Lakshminarayanan et al., 2017), Bayesian methods (MacKay, 1992; Neal, 2012; Blundell et al., 2015), and post-processing methods such as temperature scaling (Guo et al., 2017), ODIN (Liang et al., 2017), and Mahalanobis distance (Lee et al., 2018). However, the best performing methods under dataset shift either require additional memory and time complexity, or are hard to scale up for large datasets and architectures (Snoek et al., 2019).

In this work, we first study the contribution of the loss function used during training to the quality of predictive uncertainty of models. Specifically, we show why the parametrization of the probabilities underlying softmax cross-entropy loss are ill-suited for uncertainty estimation. We then propose two simple replacements to the parametrization of the probability distribution: (1) a one-vs-all normalization scheme that does not force all points in the input space to map to one of the classes, thereby naturally capturing the notion of “none of the above”, and (2) a distance-based logit representation to encode uncertainty as a function of distance to the training manifold.

2 Choice of parametrization of probabilities

In a supervised learning setting, we assume we have access to training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ drawn from a training distribution $p_{\text{in}}(\mathbf{x}, y)$, where \mathbf{x} denotes the feature vector and $y \in \{1, \dots, K\}$ denotes the class label. We train a model to make predictions on unlabeled testing data $\{\mathbf{x}_j\}_{j=1}^M$ by parametrizing the embeddings $f_{\theta}(\mathbf{x})$ of a neural network model into a predictive probability distribution $p(\hat{y}|\mathbf{x})$, and minimizing the empirical loss function calculated over a mini-batch.

In the traditional discriminative setting, the logit outputs of a neural network are calculated from the latent space embed-

¹Google Research, Brain Team ²Work done as an AI Resident. Correspondence to: Shreyas Padhy <shreyaspadhy@google.com>.

dings through an affine transformation $z_j = \mathbf{w}_j^\top f_\theta(\mathbf{x}) + b_j$. The probability distribution is then calculated through the softmax normalization,

$$p_S(\hat{y}^{(k)}|\mathbf{x}) = \frac{\exp(\mathbf{w}_k^\top f_\theta(\mathbf{x}) + b_k)}{\sum_{j=1}^K \exp(\mathbf{w}_j^\top f_\theta(\mathbf{x}) + b_j)}. \quad (1)$$

The parameters θ are learned by minimizing *cross-entropy* loss, which corresponds to maximizing the log-likelihood:

$$\mathcal{L}_S(\hat{y}^{(k)}|\mathbf{x}) = -\log(p_S(\hat{y}^{(k)}|\mathbf{x})). \quad (2)$$

This parametrization of the probabilities that are output by the model has been shown to result in misleadingly high-confidence predictions under dataset shift, and on OOD data (Guo et al., 2017; Hendrycks & Gimpel, 2016). Firstly, generating the class logits through an affine transformation of the embedding space has been shown to result in a confidence landscape that is invariant to the distance from the in-distribution embeddings in the model’s latent space (Hein et al., 2019). Secondly, the softmax normalization function explicitly maps the class logits into a K -dimensional simplex, for a K class problem, and this closed-world assumption of the input space results in highly confident, yet mis-classified predictions on entirely OOD points.

2.1 Distance-based logits

A natural way of mitigating the first issue is to explicitly encode a notion of distance in the formulation of the class logits from the learnt embeddings of a predictive model. This can be achieved by scaling the logits to be proportional to a function of distance of a point from the training manifold, and there has been previous work that enforces this parametrization in different ways (Zadeh et al., 2018; Xing et al., 2019). A particularly simple approach is one suggested by Macêdo et al. (2019) for OOD detection, where the logits z_j are defined as the negative Euclidean distance between the embeddings and the weights of the last dense layer, $z_j = -\|f_\theta(\mathbf{x}) - \mathbf{w}_j\|$. The probability distribution is calculated similarly with a softmax distribution, given as

$$p_D(\hat{y}^{(k)}|\mathbf{x}) = \frac{\exp(-\|f_\theta(\mathbf{x}) - \mathbf{w}_k\|)}{\sum_j \exp(-\|f_\theta(\mathbf{x}) - \mathbf{w}_j\|)} \quad (3)$$

and the loss function, called *Distinction Maximization (DM)* or *Isotropic Maximization* loss is calculated by maximizing the log-likelihood, given as

$$\mathcal{L}_D(\hat{y}^{(k)}|\mathbf{x}) = -\log(p_D(\hat{y}^{(k)}|\mathbf{x})) \quad (4)$$

According to Macêdo et al. (2019), the weights of the last layer \mathbf{w}_j end up approximately learning the class prototypes of the j th class during optimization in a mini-batch setting.

In practice, DM loss achieves better performance on OOD detection compared to the baseline of softmax cross-entropy. However, we observe that for in-distribution calibration and under dataset shift, a distance-based loss function that uses the softmax normalization performs worse, due to either a systematic under-confidence issue (which is very apparent on the CIFAR-10 dataset in particular, see Figure S2), or due to the underlying assumption of the softmax normalization that maps the domain to one of the K classes. Due to the relative normalization of the softmax function, even small differences in the distance of a point between class centers are highly exacerbated due to the exponentiation of the logits, resulting in points far away from the training manifold being predicted with a uniform confidence. In summary, the softmax normalization undoes the distance constraint imposed in this specific formulation, and we demonstrate this pathology in Figure 1 for a 2-dimensional toy example.




Figure 1. Confidence landscapes for a 2-dimensional toy example with 10 classes (left). We train a 2 layer fully-connected neural network with the different loss functions (further details in Appendix E.1). It can be seen that a distance-based loss function with the softmax normalization results in infinite regions of space with uniform confidence divided by the decision boundaries (center), whereas a distance-based method with one-vs-all normalization learns to predict higher confidences centered on the training manifold (right).

2.2 Revisiting Neural One-vs-All Classifiers

In order to avoid imposing the closed-world assumption of the softmax normalization, we explore an alternate formulation of the predictive discrimination task, where the K class problem is split into K independent binary tasks, each parametrized by a sigmoid normalization. More concretely, for affine-transformed logits $z_j = \mathbf{w}_j^\top f_\theta(\mathbf{x}) + b_j$, instead of a softmax normalization, we take K independent sigmoids, so that

$$p_{\text{OVA}}(\hat{y}^{(k)}|\mathbf{x}) = \frac{1}{1 + \exp(-(\mathbf{w}_k^\top f_\theta(\mathbf{x}) + b_k))} \quad (5)$$

whereas for a distance-based logit given by $z_j = -\|f_\theta(\mathbf{x}) - \mathbf{w}_j\|$, we have

$$p_{\text{OVADM}}(\hat{y}^{(k)}|\mathbf{x}) = \frac{2}{1 + \exp(-\|f_\theta(\mathbf{x}) - \mathbf{w}_k\|)}. \quad (6)$$

The factor of 2 here is a technical implementation detail that allows us to map the logits in a distance-based formu-

lation from the $(-\infty, 0]$ domain to the $[0, 1]$ domain. Since the logits are defined as the negative of the Euclidean distances, they are constrained to be in the negative domain, and the maxima of the sigmoid function in this range is 0.5 (at a value of 0), which corresponds to a distance of 0 to the learned class center. The factor of 2 therefore maps a distance of 0 to the learned class center to a confidence of 1.

In both cases, the loss function for a single example is then given by a sum over the K independent binary probabilities, maximizing the binary log-likelihood for the positive class and minimizing it for the negative classes

$$\mathcal{L}_{\text{OVA}} = -\log p(\hat{y}^{(k)}|\mathbf{x}) - \sum_{j=1, j \neq k}^K \log(1 - p(\hat{y}^{(j)}|\mathbf{x})). \quad (7)$$

One-vs-all formulations of loss functions have been studied before, especially in the context of multi-class classification (Duan et al., 2003; Rifkin & Klautau, 2004) and for Support Vector Machines (SVMs) (Liu & Zheng, 2005; Anthony et al., 2007; Mathur & Foody, 2008). More recently, one-vs-all formulations have been used to train classifiers for open-set recognition (Shu et al., 2017), due to their ability to encode a "none of the above" class naturally. The predictive uncertainty of one-vs-all classifiers were very recently explored by Franchi et al. (2020), where individual deep ensemble members were trained to specialize in independent one-vs-all tasks for a multiclass problem.

Formulating the loss function as in Equation (7) offers a few benefits, however. Firstly, by replacing the softmax normalization by K independent sigmoids, models are able to predict a higher confidence centered on the training manifold for each specific class versus the rest of the training manifold (see Figure 1). Secondly, by minimizing a linear combination of the binary log-likelihoods, we can train a single deterministic model to learn to optimize over the K binary tasks jointly, avoiding additional overhead during training or evaluation. Thirdly, a distance-based one-vs-all formulation learns more meaningful class-centers in a mini-batch fashion, compared to DM loss with a softmax normalization. We demonstrate this behavior for the 2-dimensional toy example in Figure S1, where it can be seen that softmax DM loss does not learn good representations of the centers of embeddings. Instead, due to the relative nature of the softmax normalization, it is more likely to learn a pivot point in embedding space, where if all points of a class are closest to their pivot point instead of the others, they will still get correctly classified. One-vs-all DM loss does not suffer from this issue, as the formulation in Equation 6 maps a distance of 0 to a class center in the embedding space, to a confidence score of exactly 1, whereas this is not strictly true for a softmax-based formulation.


3 Experimental Results

We report the performance of our proposed methods on a variety of different tasks that measure in-distribution calibration, robustness under covariate shift, and performance on OOD data. The quality of uncertainty for a predictive distribution is calculated using **Expected Calibration Error** (ECE) (Guo et al., 2017), which is a binned measure of the difference between the accuracy of a model and its predicted confidence, where the confidence is defined as the maximum predicted class probability for that example. For points with binned confidences B_i , if we define $\text{acc}(B_i)$ as the accuracy and $\text{conf}(B_i)$ as the confidence, the ECE is then defined as $\text{ECE} = \sum_i \frac{B_i}{N} |\text{acc}(B_i) - \text{conf}(B_i)|$. We use $N = 15$ bins to calculate ECE for our experiments.


3.1 Image Classification Tasks

To evaluate the behavior of discriminative models trained under our proposed loss functions, we benchmark the methods using the methodology first proposed in (Snoek et al., 2019) on the CIFAR-10-C and CIFAR-100-C datasets. These image datasets contain corrupted versions of the original test datasets, with 19 types of corruption applied with 5 levels of intensity each (Hendrycks & Dietterich, 2019b). In our experiments, models are trained with identical architectures and optimization setups, and for loss functions with a distance-based formulation, the last layer of the model is replaced with a distance-based layer, as defined in Equation (3) and Equation (6). For CIFAR-10, we train a ResNet 20 model using the different loss functions (further implementation details are given in Appendix E.2), where the distance-based layer weights are initialized by zeros as specified in (Macêdo et al., 2019), though we notice no difference if they are initialized using the standard random initialization in TensorFlow (Abadi et al., 2016). From Figure 2, we can see that one-vs-all methods consistently perform better under covariate shift while maintaining predictive accuracy compared to the baseline of softmax crossentropy, with a distance-based one-vs-all formulation performing the best amongst the methods. A distance-based softmax formulation suffers from consistently under-confident predictions, resulting in ECE improving as prediction accuracy drops at higher levels of shift, and this behavior is further discussed in Appendix B.

We also report results on the CIFAR-100 dataset using a Wide ResNet 28-10 model (Zagoruyko & Komodakis, 2016), with further implementation details in Appendix E.3. The results in Figure 3 show that one-vs-all methods consistently improve uncertainty while maintaining predictive accuracy compared to their softmax-based counterparts, with a distance-based formulation performing the best, especially at high levels of shift.



(a) Expected Calibration Error



(b) Accuracy


Figure 2. Expected Calibration Error and Accuracy under dataset shift for CIFAR-10-C. The box plots show the median, quartiles, minimum, and maximum performance per method. One-vs-all classifiers maintain similar predictive accuracy under covariate shift as the softmax cross-entropy baseline, while consistently achieving better calibration under covariate shift, especially as shift increases, and adding a distance-based logit representation further improves ECE performance. Data at [this url](#).

3.2 OOD Performance with Accuracy as a Function of Confidence


In order to evaluate the performance of one-vs-all methods in situations where it is preferable for discriminative models to avoid incorrect yet overconfident predictions, we consider plotting accuracy vs. confidence plots (Lakshminarayanan et al., 2017), where the trained models are evaluated on cases where the confidence estimates are above a user-specified threshold. In Figure 4, we plot confidence vs accuracy curves for a Wide ResNet model trained on CIFAR-100, and evaluated on a test set containing both CIFAR-100 and an OOD dataset (SVHN (Netzer et al., 2011), CIFAR-10). Test examples higher than a confidence threshold $0 \leq \tau \leq 1$ are filtered and their accuracy is calculated. From these curves, we would expect a well-calibrated model to have a monotonically increasing curve, with higher accuracies for higher confidence predictions, and we can see that one-vs-all methods consistently result in more robust accuracies across a larger range of confidence thresholds. In Appendix D, we also report more traditional OOD detection metrics such as AUROC and AUPRC, and show how one-vs-all DM loss performs poorly due to a high overlap between incorrectly classified in-distribution points and OOD points, especially at lower confidences.

4 Conclusions

In order to improve the predictive uncertainty of discriminative models, we analyze the different failure modes of softmax crossentropy, and show that the softmax function contributes significantly to overconfident predictions under covariate shift and on OOD data. We propose a one-vs-all formulation instead that improves the performance of neural



(a) Expected Calibration Error



(b) Accuracy

Figure 3. Expected Calibration Error and Accuracy under dataset shift for CIFAR-100-C. One-vs-all classifiers maintain similar predictive accuracy under covariate shift as the softmax cross-entropy baseline, while consistently achieving better calibration under covariate shift, especially as shift increases, and adding a distance-based logit representation further improves ECE performance. Data at [this url](#).




Figure 4. Accuracy vs Confidence curves: Wide ResNet 28-10 model trained on CIFAR-100 and tested on both in-distribution CIFAR-100 test and OOD CIFAR-10 (left) and SVHN (right) test data. Softmax cross-entropy tends to produce more overconfident incorrect predictions, whereas one-vs-all methods are significantly more robust to OOD data.

networks on image classification tasks. We also show how the linear transformations of the embeddings that result in the logits also contribute to miscalibration, and demonstrate how distance-based logits followed by a softmax normalization do not mitigate this issue. However, combining a distance-based formulation with a one-vs-all scheme results in even better performance on certain image datasets, and on a language intent classification task (Appendix C.2). However, one-vs-all distance-based losses suffer from a few setbacks; we observe that training from scratch is challenging for datasets with a large value of K such as Imagenet (Appendix C.1), and further exploration is required to scale these losses to such large datasets, and to non-image domains. However, we believe that decomposing a multiclass problem into binary tasks using one-vs-all formulations allows for further exploration into binary losses that have been extensively studied for calibration tasks (Gneiting & Raftery, 2007; Reid & Williamson, 2010; Mukhoti et al., 2020), which would make for an interesting avenue of further research.

References

- TensorFlow Datasets, a collection of ready-to-use datasets. <https://www.tensorflow.org/datasets>.
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- Anthony, G., Gregg, H., and Tshilidzi, M. Image classification using svms: one-against-one vs one-against-all. *arXiv preprint arXiv:0711.2914*, 2007.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- DeGroot, M. H. and Fienberg, S. E. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2):12–22, 1983.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Duan, K., Keerthi, S. S., Chu, W., Shevade, S. K., and Poo, A. N. Multi-category classification by soft-max combination of binary classifiers. In *International Workshop on Multiple Classifier Systems*, pp. 125–134. Springer, 2003.
- Dusenberry, M. W., Tran, D., Choi, E., Kemp, J., Nixon, J., Jerfel, G., Heller, K., and Dai, A. M. Analyzing the role of model uncertainty for electronic health records. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pp. 204–213, 2020.
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., and Thrun, S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, 2017.
- Franchi, G., Bursuc, A., Aldea, E., Dubuisson, S., and Bloch, I. One versus all for deep neural network incertitude (ovnni) quantification. *arXiv preprint arXiv:2006.00954*, 2020.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hein, M., Andriushchenko, M., and Bitterwolf, J. Why ReLU networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 41–50, 2019.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019a. URL <https://openreview.net/forum?id=HJz6tiCqYm>.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019b.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, pp. 6402–6413, 2017.
- Larson, S., Mahendran, A., Peper, J. J., Clarke, C., Lee, A., Hill, P., Kummerfeld, J. K., Leach, K., Laurenzano, M. A., Tang, L., et al. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*, 2019.
- Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pp. 7167–7177, 2018.
- Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017.

- Liu, Y. and Zheng, Y. F. One-against-all multi-class svm classification using reliability measures. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pp. 849–854. IEEE, 2005.
- Macêdo, D., Ren, T. I., Zanchettin, C., Oliveira, A. L., Tapp, A., and Ludermit, T. Isotropic maximization loss and entropic score: Fast, accurate, scalable, unexposed, turnkey, and native neural networks out-of-distribution detection. *arXiv 1908.05569*, 2019.
- MacKay, D. J. *Bayesian methods for adaptive models*. PhD thesis, California Institute of Technology, 1992.
- Mathur, A. and Foody, G. M. Multiclass and binary svm classification: Implications for training and classification users. *IEEE Geoscience and remote sensing letters*, 5(2): 241–245, 2008.
- Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P., and Dokania, P. The intriguing effects of focal loss on the calibration of deep neural networks, 2020. URL <https://openreview.net/forum?id=SJxTZehFPH>.
- Neal, R. M. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Niculescu-Mizil, A. and Caruana, R. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 625–632, 2005.
- Reid, M. D. and Williamson, R. C. Composite binary losses. *Journal of Machine Learning Research*, 11(Sep):2387–2422, 2010.
- Rifkin, R. and Klautau, A. In defense of one-vs-all classification. *Journal of machine learning research*, 5(Jan): 101–141, 2004.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015a. doi: 10.1007/s11263-015-0816-y.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. ImageNet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015b.
- Shorten, C. and Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60, 2019.
- Shu, L., Xu, H., and Liu, B. DOC: Deep open classification of text documents. *arXiv preprint arXiv:1709.08716*, 2017.
- Snoek, J., Ovadia, Y., Fertig, E., Lakshminarayanan, B., Nowozin, S., Sculley, D., Dillon, J., Ren, J., and Nado, Z. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pp. 13969–13980, 2019.
- Tran, D., Hoffman, M. D., Moore, D., Suter, C., Vasudevan, S., Radul, A., Johnson, M., and Saurous, R. A. Simple, distributed, and accelerated probabilistic programming. In *Neural Information Processing Systems*, 2018.
- Tsai, H., Riesa, J., Johnson, M., Arivazhagan, N., Li, X., and Archer, A. Small and practical bert models for sequence labeling. *arXiv preprint arXiv:1909.00100*, 2019.
- Xing, C., Arik, S., Zhang, Z., and Pfister, T. Distance-based learning from errors for confidence calibration. *arXiv preprint arXiv:1912.01730*, 2019.
- You, Y., Gitman, I., and Ginsburg, B. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- Zadeh, P. H., Hosseini, R., and Sra, S. Deep-RBF networks revisited: Robust classification with rejection. *arXiv preprint arXiv:1812.03190*, 2018.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

A Comparison of learned class centers

In order to better portray the optimization issues introduced by the softmax normalization in learning meaningful class centers using a DM Loss formulation, we train models on a 2D toy dataset using distance-based logits, but with differing normalization schemes. In both cases, the learnt embedding space is 16-dimensional, and the weight matrix of the last layer is given by $\mathbf{W} \in \mathbb{R}^{16 \times 10}$, where the j th column $\mathbf{W}_{:,j}$ corresponds to the learned class center for the j th class. In order to plot the learned class centers in Figure S1, the 16-dimensional hidden embedding space obtained from a forward pass of the training data is projected onto a 2-dimensional PCA subspace, along with the columns of the weight matrix of the last layer.




Figure S1. Penultimate Layer Embeddings and Learned Class Centers (solid circles) for DM Loss (left) and One-vs-all DM Loss (right). It can be seen that the weights of the last layer learnt by DM Loss do not accurately represent the class centers of the embeddings, due to the softmax normalization. However, the one-vs-all formulation learns more meaningful class-centers that better approximate the centers of class-specific embeddings. Implementation details are mentioned in Appendix E.1.

B Reliability Diagrams

To further understand the interplay of behaviors between one-vs-all normalizations and distance-based logits, we plot reliability diagrams (DeGroot & Fienberg, 1983; Niculescu-Mizil & Caruana, 2005) which show accuracy as a function of confidence in the form of binned histograms of confidences over the in-distribution test set. We plot reliability diagrams for models trained on CIFAR-10 in Figure S2 and CIFAR-100 in Figure S3. From these plots, we can see that softmax cross-entropy tends to result in models that have very high confidences on the test set, often with lower accuracies than expected. One-vs-all methods tend to give better calibrated predictions, especially for points with lower accuracy corresponding to lower confidence predictions. Interestingly, we see somewhat anomalous behavior of DM loss on the CIFAR-10 dataset, which has been shown before in Macêdo et al. (2019), where all confidences for points in the in-distribution test set are lower than a certain threshold. One hypothesis for this behavior is the fact that it is easy to obtain high accuracy on the training data without necessarily learning to cluster points close to their class centers, as long as all points of a class are closest to the corresponding

weight in the last layer, and this sort of behavior is observed even in the 2D toy example in Appendix A. This anomalous histogram of confidences also explains the decreasing trend of ECE seen under dataset shift in Figure 2, as under higher intensity of shift, models would tend to get less accurate, while the predictions made by a model trained with the DM loss are consistently underconfident. We do not see this behavior for CIFAR-100 however.

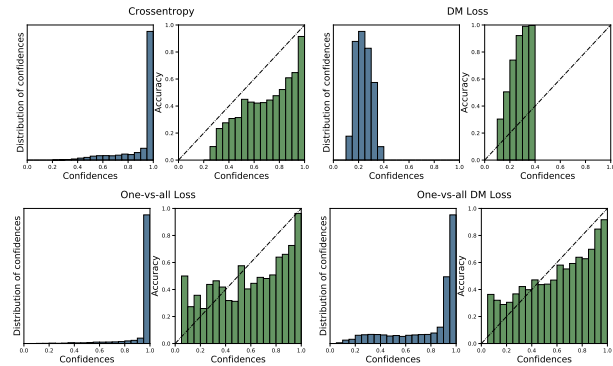


Figure S2. Reliability diagrams for ResNet 20 models trained on the CIFAR-10 dataset, for softmax cross-entropy (top-left), DM loss (top-right), One-vs-all loss (bottom-left) and One-vs-all DM loss (bottom-right). One-vs-all methods show better-calibrated reliability curves, especially showing a larger number of points with lower confidence on the in-distribution test set, with correspondingly low accuracies. Of particular note is the behavior of DM loss, where there is a systematic issue of underconfidence, as all points lie below a threshold. This behavior has been reported in Macêdo et al. (2019), and though this results in better OOD performance overall, the exceedingly low confidences result in very miscalibrated predictions in-distribution and under dataset shift.

C Additional Results

C.1 ImageNet

We report the performance of the proposed loss functions on the ImageNet dataset (Russakovsky et al., 2015b), using a ResNet-50 model (He et al., 2016) trained using the setup described in Appendix E.4. From Figure S4, we see that a one-vs-all formulation with affine-transformed logits has similar predictive accuracy to softmax cross-entropy, and performs more robustly under dataset shift, especially at higher intensities of shift. The distance-based one-vs-all formulation however has difficulty matching the predictive performance of softmax cross-entropy, while still having lower ECE under dataset shift. We hypothesize that the large number of classes in ImageNet make a distance-based formulation difficult to scale, especially since the gradients to learn class centers are calculated in a mini-batch setting, and even with a batch size of 4096, the number of points per class can be heavily imbalanced, resulting in poor gradients




Figure S3. Reliability diagrams for Wide ResNet 28-10 models trained on the CIFAR-100 dataset, for softmax cross-entropy (**top-left**), DM loss (**top-right**), One-vs-all loss (**bottom-left**) and One-vs-all DM loss (**bottom-right**). One-vs-all methods consistently have more datapoints with lower confidences, which correspond to lower accuracy predictions as well, as opposed to the softmax-based methods. Interestingly, DM loss does not suffer the same issue of underconfidence as in CIFAR-10 (Figure S2).

for optimization.




Figure S4. Mean Accuracy (left) and Mean Expected Calibration Error (right) for a ResNet-50 model trained on ImageNet and evaluated on the ImageNet-C corrupted data, under 5 levels of shift.

C.2 CLINC Intent Classification Dataset

We also evaluate the performance of one-vs-all methods on out-of-distribution data on modalities beyond common image tasks, namely a language understanding task using the CLINC out-of-scope intent detection benchmark dataset (Larson et al., 2019). The benchmark dataset contains 150 types of in-domain services and 1500 out-of-domain sentences, and it is important for real-world, goal-oriented dialog systems to be able to distinguish between in-domain and out-of-domain sentences. For this task, we train a BERT-Mini model (Devlin et al., 2018; Tsai et al., 2019) on the in-domain data from scratch (with further details in Appendix E.5), and report both the in-distribution ECE, and performance on out-of-distribution data. From Table C.2, we see that one-vs-all formulations are competitive with the predictive accuracy achieved by softmax cross-entropy, while improving in-distribution calibration. Simultaneously, we can see from Figure S5 that one-vs-all DM Loss performs more robustly at different confidence thresholds, however

regular one-vs-all loss struggles at higher confidences to distinguish in-distribution sentences from OOD sentences.

Loss function	Test Accuracy	Test ECE
Crossentropy	90.16	0.079
DM Loss	89.89	0.083
One-vs-all Loss	90.67	0.078
One-vs-all DM Loss	88.89	0.052

Table S1. In-distribution test accuracy and test ECE for a BERT-Mini model trained on the Intent Classification dataset. Results are averaged over 5 runs.




Figure S5. Accuracy vs Confidence curves: BERT-Mini model trained on the CLINC intent classification dataset. Softmax crossentropy tends to produce more overconfident incorrect predictions, whereas one-vs-all DM loss is significantly more robust to out-of-distribution data, especially at lower confidence thresholds.

D On OOD detection metrics: AUROC versus Confidence-Accuracy curves

In this section, we report the Area Under the Receiver Operating Characteristic curve (AUROC) and Area under the Precision-Recall curve (AUPRC) for OOD tasks in the image domain. Both the AUROC and AUPRC are calculated for the binary task of classifying in-distribution test points from OOD test points. From Tables S2 and S3, we can see that one-vs-all DM loss performs poorly compared to the other methods for a scalar metric such as AUROC and AUPRC. However, from the confidence versus accuracy plots in Section 3.2, we can see that one-vs-all DM loss is more sensitive to OOD inputs, especially at lower confidence values. In order to explain this apparent disparity, we plot the histograms of the predicted confidence for three sets of points:

1. correctly classified in-distribution test points,
2. incorrectly classified in-distribution test points, and
3. OOD test points.

The disparity above can be explained by the fact that confidence versus accuracy curves measure the ability of a

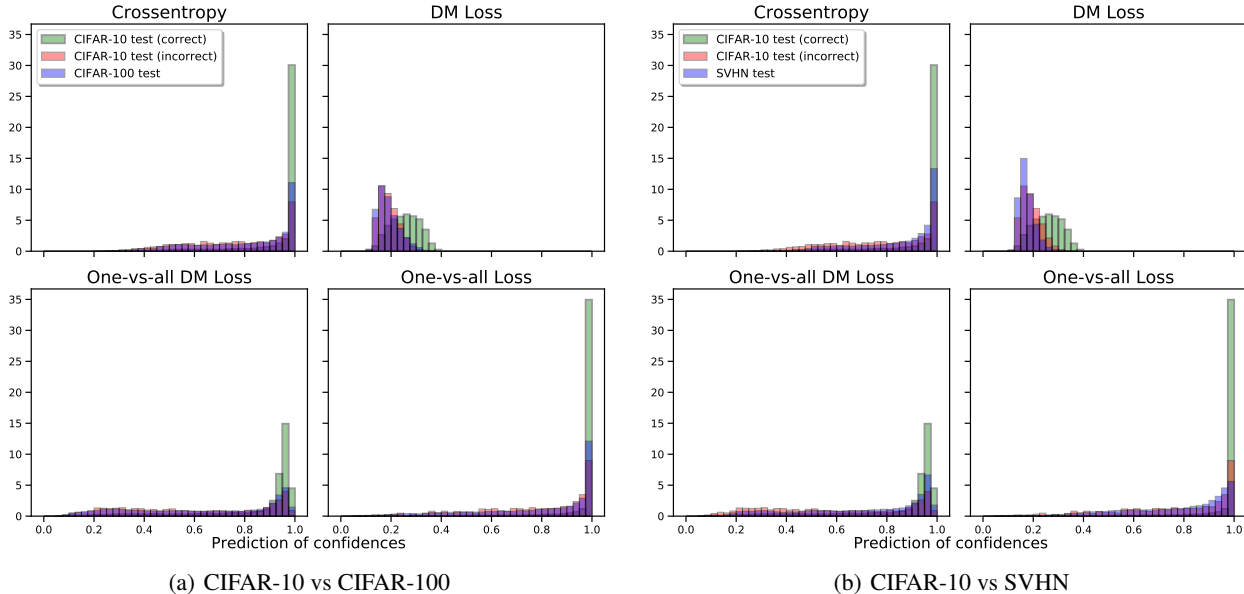


Figure S6. Histograms of predicted confidences for a ResNet-20 architecture trained with different loss functions on the CIFAR-10 dataset, and tested on the OOD CIFAR-100 and SVHN datasets. We can see that there is a much higher overlap in the confidence distributions of incorrectly classified in-distribution test points and OOD test points, which results in poorer performance when using a scalar metric such as AUROC and AUPRC that aggregate over all thresholds in $[0, 1]$.

method to separate set 1 from sets 2+3, whereas AUROC and AUPRC measure the ability of a method to separate sets 1+2 (in-distribution) from set 3 (OOD). Specifically, the confidence versus accuracy curve measures accuracy only on the unrejected inputs, hence it does not penalize a method for assigning the same low confidence to incorrectly classified in-distribution test points and OOD test inputs. On the other hand, AUROC measures the ability to separate in-distribution from OOD inputs, so it does not penalize a method for assigning the same confidence to correctly and incorrectly classified in-distribution inputs.

From Figure S7, we can see that one-vs-all DM loss outputs much lower confidences for incorrectly classified in-distribution test points, while also outputting low confidences for OOD test points. Due to the overlap in the confidence histograms for these two subsets of data, a scalar metric such as AUROC and AUPRC will report lower performance, due to difficulty in distinguishing between the two at low confidences.

\mathcal{D}_{in}	Loss function	AUROC (\uparrow) / AUPRC (\uparrow)	
		$\mathcal{D}_{ood} = \text{CIFAR-100}$	$\mathcal{D}_{ood} = \text{SVHN}$
CIFAR-10	Crossentropy	0.7334 / 0.6789	0.7026 / 0.8058
	DM Loss	0.7478 / 0.7104	0.8264 / 0.8940
	One-vs-all Loss	0.8449 / 0.8037	0.8934 / 0.9248
	One-vs-all DM Loss	0.6805 / 0.6486	0.6191 / 0.7725

Table S2. AUROC and AUPRC scores for OOD detection tasks for a Wide ResNet 28-10 model trained on the CIFAR-10 dataset and evaluated on the CIFAR-100 and SVHN OOD datasets.

\mathcal{D}_{in}	Loss function	AUROC (\uparrow) / AUPRC (\uparrow)	
		$\mathcal{D}_{ood} = \text{CIFAR-10}$	$\mathcal{D}_{ood} = \text{SVHN}$
CIFAR-100	Crossentropy	0.8003 / 0.7587	0.6420 / 0.7779
	DM Loss	0.7862 / 0.7488	0.6744 / 0.8047
	One-vs-all Loss	0.8019 / 0.7601	0.6700 / 0.8028
	One-vs-all DM Loss	0.7669 / 0.7274	0.5966 / 0.7948

Table S3. AUROC and AUPRC scores for OOD detection tasks for a Wide ResNet 28-10 model trained on the CIFAR-100 dataset and evaluated on the CIFAR-10 and SVHN OOD datasets.

E Implementation Details

E.1 2D Toy Dataset

The synthetic 2-dimensional toy dataset is generated by drawing 1000 points per class from a 2-dimensional normal distribution, where the points for the j th class are given by $\mathcal{X}_j \sim \mathcal{N}_2 \left(\left(20 \cos\left(\frac{j}{10 \times 2\pi}\right), 20 \sin\left(\frac{j}{10 \times 2\pi}\right) \right), 2\mathbf{I} \right)$. For training, we use a fully-connected neural network with 2 hidden layers containing 16 units each, trained on a batch size of 128 using SGD for 10000 steps. For all examples, the trained models achieve a training accuracy of 100%.

E.2 CIFAR-10

We use the Tensorflow Datasets (TFD) implementations of both the CIFAR-10 (Krizhevsky, 2009) and CIFAR-10-C (Hendrycks & Dietterich, 2019a) datasets. For training, the CIFAR-10 dataset is augmented by padding with zeros, followed by random cropping, random flipping, and rescaling




Figure S7. Histograms of predicted confidences for a Wide ResNet 28-10 architecture trained with different loss functions on the CIFAR-100 dataset, and tested on the OOD CIFAR-10 and SVHN datasets. We can see that there is a much higher overlap in the confidence distributions of incorrectly classified in-distribution test points and OOD test points, which results in poorer performance when using a scalar metric such as AUROC and AUPRC that aggregate over all thresholds in $[0, 1]$.

to $[-1, 1]$. We train the Resnet-20 v1 architecture on a batch size of 512 using the Adam optimizer with a learning rate schedule, and tune the following hyperparameters using random search with 100 trials within ranges specified as follows: learning rate (1×10^{-3} to 1), momentum (0.85 to 0.99), and Adam’s ϵ (1×10^{-8} to 1×10^{-5}).

E.3 CIFAR-100

We use the Tensorflow Datasets implementation of CIFAR-100 (Krizhevsky, 2009) and generate the CIFAR-100-C dataset by calculating the 17 types of corruptions (with the exception of `motion_blur` and `snow`) over 5 levels of intensity as defined in Hendrycks & Dietterich (2019a). The CIFAR-100 dataset is augmented in a similar manner as CIFAR-10 during training, and we use a Wide ResNet 28-10 architecture with L2 regularization (Zagoruyko & Komodakis, 2016) trained on a batch size of 512 for 200 epochs using SGD with Nesterov momentum (Tran et al., 2018). We tune the following hyperparameters using random search with 50 trials with ranges specified as follows: learning rate (1×10^{-4} , 1×10^{-1}) and L2 weighing factor (1×10^{-5} , 1×10^{-3}).

E.4 ImageNet

We use the Tensorflow Datasets implementation of the ImageNet (Russakovsky et al., 2015a) and ImageNet-C (Hendrycks & Dietterich, 2019a) datasets. During training, the ImageNet dataset is preprocessed using the data augmen-

tation pipeline commonly used by the Inception architecture (Shorten & Khoshgoftaar, 2019). We use a ResNet-50 "v1.5" architecture¹ with L2 regularization trained on a batch size of 4096 for 90 epochs using SGD with Nesterov momentum and a warmup learning rate schedule (You et al., 2017). We tune the following hyperparameters using random search with 50 trials with ranges specified as follows: learning rate (5×10^{-1} , 1×10^1) and L2 weighing factor (1×10^{-5} , 1×10^{-3}).

E.5 Intent Classification

We use the CLINC out-of-scope (OOS) intent classification data as defined in Larson et al. (2019), where the BERT tokenizer is used to pre-tokenize sentences with a maximum sequence length of 32. For training, we use the BERT-Mini architecture (Devlin et al., 2018; Tsai et al., 2019), which is a 4 layer transformer with 256 hidden units per layer and 4 attention heads, trained on a batch size of 256 for 1000 epochs using the Adam optimizer, where the learning rate is tuned for 30 trials sampled uniformly between (1×10^{-5} , 1×10^{-3}).

¹The v1.5 architecture differs slightly from the v1 architecture more commonly defined in He et al. (2016). For further details, please see <http://torch.ch/blog/2016/02/04/resnets.html>.