

Feature Density Estimation for Out-of-Distribution Detection via Normalizing Flows

Evan D. Cook, Marc-Antoine Lavoie, Steven L. Waslander

Institute for Aerospace Studies

University of Toronto

Toronto, ON, Canada

{evan.cook, marc-antoine.lavoie, steven.waslander}@robotics.utoronto.ca

Abstract—Out-of-distribution (OOD) detection is a critical task for safe deployment of learning systems in the open world setting. In this work, we investigate the use of feature density estimation via normalizing flows for OOD detection and present a fully unsupervised approach which requires no exposure to OOD data, avoiding researcher bias in OOD sample selection. This is a post-hoc method which can be applied to any pretrained model, and involves training a lightweight auxiliary normalizing flow model to perform the out-of-distribution detection via density thresholding. Experiments on OOD detection in image classification show strong results for far-OOD data detection with only a single epoch of flow training, including 98.2% AUROC for ImageNet-1k vs. Textures, which exceeds the state of the art by 7.8%. We additionally explore the connection between the feature space distribution of the pretrained model and the performance of our method. Finally, we provide insights into training pitfalls that have plagued normalizing flows for use in OOD detection.

Keywords—normalizing flows; out of distribution detection;

I. INTRODUCTION

Machine learning has rapidly advanced in recent years, with state of the art models performing impressive tasks in a wide range of technical domains. However the standard workflow in machine learning is significantly less flexible than learning observed in animals in nature. While biological neural systems continually learn in uncontrolled environments, artificial neural networks are instead trained with a closed-world assumption [1] on a fixed corpus of training data, validated against a set of reserved data drawn from the same data distribution, and then deployed to perform roughly the same task. When deployed, these models can be exposed to inputs that are dissimilar to the in-distribution (ID) data they were trained and validated on, potentially leading to unpredictable behavior when encountering this out-of-distribution (OOD) data.

Addressing how artificial neural networks can be used in open world situations where they may be exposed to out-of-distribution data remains a challenge. Out-of-distribution detection is the task of identifying when a sample is not drawn from the training data distribution. This is especially important in safety critical applications such as autonomous vehicles; the statistical assurances on model performance provided by the validation dataset are no longer applicable.

In this work, we revisit using feature density estimation

(FDE) via normalizing flows for out-of-distribution detection in image classification. Prior works assert that normalizing flows are not effective for OOD detection when performing density estimation in pixel space [2], and density estimation in the feature space of pretrained models has been discussed but not thoroughly investigated [3]. We demonstrate that by performing density estimation in the feature space of a pretrained image classification backbone model, normalizing the feature representations, and under-training the normalizing flow we are able to achieve competitive results on both small and large datasets. The proposed method has the advantages of being fully unsupervised and requires no exposure to OOD training data, avoiding researcher bias from a specific definition of the OOD data. Finally, this is a post-hoc method that can be applied to any pretrained classification model, and it requires training a lightweight normalizing flow model for only a single epoch to perform the feature-space density estimation for out-of-distribution detection, making it a broadly applicable technique.

II. RELATED WORK

A. Out-of-Distribution Detection

OOD detection deals with identifying semantically distinct samples (from unseen classes) to avoid erroneously classifying them as one of the classes in the training distribution. Out-of-distribution detection performance is evaluated by attempting to discriminate between a validation dataset versus an out-of-distribution dataset. The most widely used metric is the area under the receiver operating characteristic (AUROC) [4], a threshold-free classification performance metric useful for comparing unbalanced datasets.

OOD detection is a rich field with many existing approaches. These can be divided into classification-based, distance-based, generative-based, and density-based methods [1]. Classification-based approaches define a classification output that identifies ID and OOD inputs at inference time, with common baseline methods including the max-softmax probability (MSP) [5], ODIN [6], the energy score [7], and post-hoc methods that modify the feature vector activations such as ASH [8] and ReAct [9]. MSP is a simple baseline method which thresholds on the maximum class probability. Energy score is a more modern development with stronger performance while remaining simple to implement,

calculating a metric inspired by thermodynamics (the free energy) from the classification logits. ReAct is used in conjunction with the energy score, but clips off the top 10% of feature vector activations prior to evaluation, resulting in state of the art performance on large scale datasets. Distance-based methods label OOD samples as those sufficiently far from ID training samples (in feature space), and include Euclidean and Mahalanobis distance [10]. Generative-based approaches employ generative models to reconstruct inputs, and assess samples with poor reconstruction or low likelihood under the generative model as OOD [11]. Examples of this approach include VAEs with modified priors [12], hierarchical VAEs [13], and diffusion models [14], [15]. These methods require training a generative network to model the data distribution, which can be computationally demanding.

Finally, for density-based approaches a density estimation model is built from the training data such that the ID data lies within high density regions, and OOD data encountered at inference time occupies low density regions. A threshold on the density can be added to transform a density estimator into an out-of-distribution detector, classifying low probability data as out-of-distribution. In this approach the density estimator is used as a proxy for model epistemic uncertainty [16]. Density estimation methods can be performed in the input data space or a transformed representation space, and include kernel methods, radial basis functions, and normalizing flows [16], [1]. It was observed in [3] that using a normalizing flow to perform density estimation in the feature space improves performance over density estimation in the pixel space, but their analysis is extremely limited and their results do not reach other state-of-the-art methods.

B. Normalizing Flows

Normalizing flows are a class of generative neural networks that are dimensionality preserving and fully invertible. They are trained to learn a diffeomorphism to map between two probability distributions, typically a data distribution and a known base probability distribution (such as the normal distribution). Normalizing flows have the dual function of being an exact density estimator (by measuring the probability of a datapoint mapped to the base distribution), and a generative model (by sampling from the base distribution, and then mapping into the data space). Mathematically, normalizing flows can be written as implementing a change of variables:

$$p(z) = q(f_{\theta}(z)) \left| \det \left(\frac{\partial f_{\theta}(z)}{\partial z^T} \right) \right|$$

where $p(z)$ is the data distribution, $q(z)$ is the known base distribution, and $f_{\theta}(z)$ is the mapping function between these two distributions, implemented as an invertible normalizing flow neural network parameterized by θ .

For a more thorough and formal review of the mathematics of normalizing flows, we refer readers to [17]. Implementing normalizing flows is often challenging, as the model must be entirely invertible and should have a Jacobian that can be efficiently calculated. However, they have shown impressive performance in many tasks, including generating realistic images of faces [18] and high quality density estimation on image data [19].

C. Normalizing Flows for Out-of-Distribution Detection

Normalizing flows have been applied to the task of OOD detection in several prior works with mixed success, but have historically performed very poorly for OOD detection in the image classification domain. When performing density estimation on pixel data in images, previous authors recommend against the use of normalizing flows, finding that they learn spurious pixel-level correlations and capture low-level statistics rather than high-level semantics [2], [3], [20].

In [21], [22], and [23] normalizing flows are applied to image segmentation anomaly detection by performing density estimation of multiscale feature map embeddings instead of pixel space. Results are promising, but limited to small scale datasets, and they use hand-tailored network architectures that do not generalize to other domains.

Flows have also been used for anomaly detection in video data. A Glow normalizing flow [18] is used by [24] to perform density estimation of the feature vectors produced by two autoencoders, one capturing spatial information and one capturing temporal information. This work highlights the importance of performing density estimation in the feature space and demonstrates competitive performance in this domain, but has a complex autoencoder architecture with a reconstruction loss term, limiting its potential applications. In [25] a normalizing flow is applied to the task of quantifying sample rareness, illustrating the value of feature space density estimation with normalizing flows for the downstream task of data mining and dataset balancing.

In [26] strong OOD detection performance is demonstrated using normalizing flows in image classification, but this is not a post-hoc method, as it requires jointly training the classifier backbone and normalizing flow together with additional hyperparameters. This approach is limited by the necessity to jointly learn the feature space, and results are only evaluated on small datasets. In [3], the concept of performing density estimation in the feature space of a pretrained classifier is briefly discussed, but their analysis is very limited and their OOD performance is not compelling. Our work carries the investigation of feature density estimation via normalizing flows much further, demonstrating that normalizing flows can achieve state of the art out-of-distribution detection in image classification using a simple, post-hoc method with no complex architecture changes or modifications to the backbone.




Figure 1: Feature density estimation for out-of-distribution detection.




Figure 2: (a) A t-SNE visualization of the feature space of a ResNet18 backbone model comparing ID data (CIFAR-10), OOD data (SVHN), and flow generated feature vectors. Clearly visible are 10 clusters corresponding to the 10 classes of CIFAR-10, and the coincidence between the ID data and the flow generated data points. (b) A histogram of the log-likelihood of ID features vs. OOD features under the normalizing flow model. Better separability of these distributions leads to higher AUROC for OOD detection.

III. METHOD

A. Feature Density Estimation

In this work we leverage a pretrained neural network backbone to provide a compressed, reduced representation of our input data that is rich in semantic information for the downstream task of image classification. We use the penultimate layer’s activations as feature vectors for density estimation. These feature vectors contain all of the necessary information for the backbone model to perform the output classification task, and are typically transformed to the final output logits using a linear projection head.

We perform density estimation on the feature representations using established normalizing flow architectures [18], [27], [28], learning an invertible mapping between the feature space and a normal probability distribution. Our normalizing flows are trained on the penultimate layer activations of a frozen pretrained image classifier, with the optimization criterion of minimizing the log-likelihood of

the transformed features. As an unsupervised method, the class labels of the original image data are unused. Once trained, the normalizing flow is a computationally efficient density estimator for the feature space of the pretrained backbone model. Out-of-distribution detection is achieved by applying a simple probability threshold to the density estimates for new samples, classifying low density samples as OOD. See Figures 1 and 2 for a block diagram and visualization of our method.

IV. EXPERIMENTAL SETUP

We evaluate the utility of normalizing flows for out-of-distribution detection on a range of image classification tasks, using a variety of backbone networks and in-distribution datasets.

Datasets: We use CIFAR-10 [29] and ImageNet-1k [30] as our in-distribution datasets. CIFAR-10 models are evaluated against random Gaussian noise, SVHN [31], Places365 [32], CelebA [33], and CIFAR-100 [29] as out-of-distribution datasets. ImageNet-1k models are evaluated against Textures [34] and reduced versions of iNaturalist [35], SUN [36], and Places [32] datasets. The latter three are filtered to ensure these datasets contain no common classes with ImageNet-1k, as done by [9]. Datasets can be considered as either *near-OOD* or *far-OOD* [11] depending on how semantically distinct they are from the ID dataset; in this work we consider Gaussian noise and SVHN to be far-OOD from CIFAR-10, and Textures to be far-OOD from ImageNet-1k.

Evaluation: Out-of-distribution detection performance is evaluated using AUROC, calculated between the in-distribution validation dataset and out-of-distribution dataset. AUROC is a threshold-free metric, and an AUROC of 50% indicates no separability between the distributions, while an AUROC of 100% indicates perfect separability between the distributions. We evaluate our method against MSP [5], ODIN [6], energy score [7], and ReAct [9].

Normalizing Flow Models: We use a 10 block Glow [18] flow for all experiments. For flow models trained on CIFAR-10, each block is composed of two linear layers with dimension [512, 2048, 512]. For flow models trained on ImageNet-1k, each block is composed of two linear layers

which do not alter the dimensionality of the feature space (2048 for ResNet50 and 768 for Swin-T). Flow models are trained using the Adam optimizer [37] for only a single epoch with a learning rate of $1e-4$ for CIFAR-10 and $1e-5$ for ImageNet-1k.

Backbone Model: For CIFAR-10, we train a ResNet18 classifier backbone using supervised learning to a validation accuracy of 92.1%. For ImageNet-1k, we use two PyTorch pretrained models as classifier backbones: ResNet50 and Swin-T, with top-1 validation accuracies of 76.1% and 81.5% respectively [38]. Backbone weights are frozen for all experiments.

V. RESULTS

We first present our main results for CIFAR-10, summarized Table I. Our method exceeds other approaches on SVHN and Gaussian noise (far-OOD), and is competitive with alternatives on more challenging datasets (Places365, CelebA, CIFAR-100). Further, we present results for the larger scale ImageNet-1k dataset in Table II. With a ResNet backbone, our method is able to achieve 98.2% AUROC on Textures [34], obtaining 7.8% better performance than the next best method ReAct [9]. With a transformer backbone, we again outperform ReAct by 7.5% on Textures. Our method consistently outperforms competing methods at detecting the more visually distinct far-OOD samples (CIFAR-10 vs. SVHN, and ImageNet vs. Textures).

VI. DISCUSSION


To practically implement normalizing flows for OOD detection, we discuss several key considerations.

A. Flow Regularization

During training, the goal is to optimize a flow model that fits the training distribution, generalizes to the validation set, and still separates OOD data. It is critically important to manage overfitting: the separability of validation and OOD data is directly impacted by the normalizing flow’s generalization gap between the training and validation data distributions (see Figure 3 for a visualization of training, validation, and OOD likelihood distributions).

A variety of standard regularizations techniques can be used to avoid overfitting on the training data. Especially important is data augmentation. We train our flows on feature vectors obtained with dataset augmentation identical to those used to train the backbone model.

Under-training was found to be critically important. Our experiments demonstrated a surprising trend: training a normalizing flow model to minimize the validation loss may actually be detrimental to OOD detection performance. Counterintuitively, loss on a test OOD dataset *decreases* during training (OOD data becomes more likely as the flow model fits to ID data), and AUROC for this test OOD dataset peaks early, then drops with additional epochs as the flow



(a) Epoch 0, AUROC=96.1% (b) Epoch 999, AUROC=84.4%

Figure 3: Feature likelihood histograms for the same flow model at 0 epochs and 999 epochs. With further training, the likelihood increases for all distributions, but the training and validation distributions begin to separate due to overfitting, while the separability of the ID/OOD distributions degrades.




Figure 4: Normalizing flow AUROC (CIFAR-10 vs. SVHN) and loss during training. Loss of ID and OOD test datasets decreases during training. AUROC peaks early then declines.

model fits to the training data (see Figure 4). The optimal number of epochs to train a flow model for depends on the OOD dataset, flow architecture, and backbone model, and is far lower than when the validation loss begins to rise (classic overfitting). This is thus distinct from early stopping, and represents a novel and beneficial form of under-training. In our work, we report all results using normalizing flows trained to only a single epoch to ensure consistent evaluations across models and datasets.

B. Normalizing Feature Vectors

We find that normalizing the feature vectors strongly improves OOD detection performance on far-OOD datasets (CIFAR-10 vs. SVHN and ImageNet-1k vs. Textures). We can write the final linear head of the classification model as:

$$y = W^T z = \|z\| (W^T \hat{z})$$

where W is the weight matrix (ignoring the bias term), z is our feature vector, $\|z\|$ is the Euclidean norm of z , and \hat{z} is the normalized (unit-length) feature vector. We interpret the product $W^T \hat{z}$ as the semantic agreement between \hat{z} and each logit’s class, while $\|z\|$ relates to the network’s overall confidence in the output. Larger feature norms, $\|z\|$, correlate with larger logits and higher classification probabilities. Training a flow density estimator on unnormalized

Table I: Out-of-distribution detection performance results, with CIFAR-10 as in-distribution.

Backbone	Method	Out-of-Distribution Dataset (AUROC \uparrow)				
		Gaussian	SVHN	Places365	CelebA	CIFAR-100
ResNet18 + CIFAR-10	MSP	90.86	88.62	88.78	90.62	85.86
	Energy Score	87.33	87.12	92.77	94.62	88.35
	ODIN	88.95	87.47	95.13	95.16	87.43
	FDE (ours)	99.41	96.06	92.61	93.23	85.93

Table II: Out-of-distribution detection performance results, with ImageNet-1k as in-distribution.

Backbone	Method	Out-of-Distribution Dataset (AUROC \uparrow)			
		Textures	iNaturalist	SUN	Places
ResNet50 + ImageNet-1k	MSP	80.49	88.39	81.64	80.53
	Energy Score	86.80	90.62	86.57	83.96
	ReAct	90.44	96.43	94.42	91.93
	ODIN	87.72	91.41	86.82	84.36
	FDE (ours)	98.20	81.75	78.61	72.46
Swin-T + ImageNet-1k	MSP	81.72	89.86	81.26	80.33
	Energy Score	78.05	85.13	73.34	70.60
	ReAct	85.48	90.01	81.50	79.28
	ODIN	67.70	71.13	57.84	54.57
	FDE (ours)	92.94	92.15	81.96	77.02

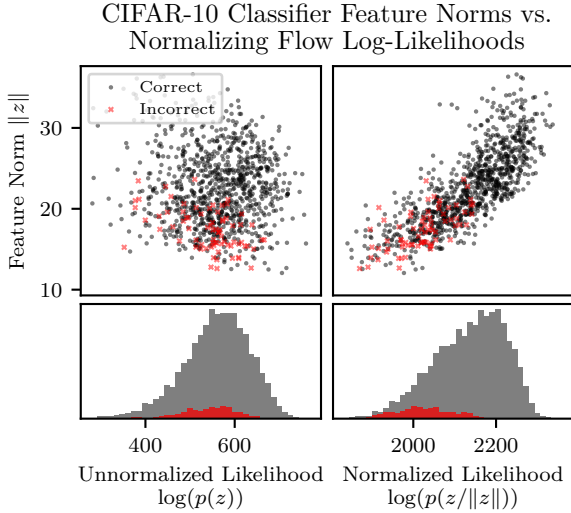


Figure 5: Visualization of feature vector norms vs. log-likelihood for a flow model trained with normalized (left) and unnormalized (right) feature vectors. For a flow model trained on unnormalized features, there is no correlation between feature norm, classification accuracy, and flow likelihood. For a flow model trained on normalized features, a correlation is observed between feature norm, classification accuracy, and flow likelihood.

feature vectors confounds the interpretation of the log-likelihood of the features. Features with very large norms may be modelled as having a low likelihood (see Figure 5), which is counterproductive. The OOD detection task is concerned with OOD data that are outside the expected semantic content, not outside the expected data likelihood (a perfect image of a cat may be semantically ID, but could be considered OOD due to an unusually high feature

vector norm). We resolve this by training the normalizing flow model on normalized feature vectors: density estimation is performed on the semantic content of the feature space disentangled from the classifier’s confidence. As shown in Figure 5, training our flow on normalized features yields a strong correlation between feature norms and feature log-likelihood: this can be interpreted as a desirable correlation between classifier confidence and likelihood of the semantic content of the feature. Further, training on normalized features introduces a correlation between feature likelihood and the probability of correct classification (despite this being an unsupervised method with no access to classification labels); no such correlation exists when training on unnormalized features. Experiments show that OOD detection performance is substantially improved on far-OOD data when performing density estimation on normalized features (Tables III, IV).

C. Flow Architecture

Normalizing flow architecture is an active area of research, with different flow designs having their own pros and cons. RealNVP [27] is fast and simple flow architecture but performs poorly compared to more modern methods. Glow [18] demonstrates good performance as a generative model but is not state of the art for density estimation, and Residual Flows [28] are excellent density estimators but are slower to train than alternatives.

Surprisingly, our experiments show that the performance of OOD detection is relatively insensitive to flow architecture. We believe this is due to the fact that discrimination between two distributions (the validation dataset and OOD dataset) is the key task, rather than high quality modeling of the training distribution. Maximum OOD detection performance is often seen after only a few training epochs of the flow model, far before the training distribution is adequately modelled.

Table III: Flow AUROC results, CIFAR-10: normalized vs. unnormalized features

Backbone Model	Feature Scaling	Out-of-Distribution Dataset (AUROC \uparrow)	
		SVHN	CIFAR-100
ResNet18 + CIFAR-10	Unnormalized	85.63	57.68
	Normalized	96.48	85.93

Table IV: Flow AUROC results, ImageNet-1k: normalized vs. unnormalized features

Backbone Model	Feature Scaling	Out-of-Distribution Dataset (AUROC \uparrow)			
		Textures	iNaturalist	SUN	Places
ResNet50 + ImageNet-1k	Unnormalized	86.95	49.78	60.21	58.89
	Normalized	98.20	81.75	78.61	72.46
Swin-T + ImageNet-1k	Unnormalized	87.87	89.27	82.27	79.56
	Normalized	92.94	92.15	81.96	77.02

Table V: Normalizing flow OOD detection vs. architecture comparison. AUROC results are generally comparable, and flow models with superior density estimation do not equate to improved OOD detection performance.

Backbone Model	Flow Architecture	Out-of-Distribution Dataset (AUROC \uparrow)	
		SVHN	CIFAR-100
ResNet18 + CIFAR-10	RealNVP	92.13	86.33
	Glow	96.99	85.87
	Residual Flow	96.89	80.32

As discussed in Section VI-A, extensively training the flow model to maximize the likelihood of the ID data is unimportant for OOD detection. Instead, the difference in likelihood between the ID and OOD distributions is more important. As such, more sophisticated flow models which advance the state of the art in density estimation and offer improved likelihood of ID data are not necessarily advantageous for OOD detection (see Table V). In our experiments Glow [18] was used, as it performed well while being faster than more complex methods, such as Residual flows [28] and FFIORD [39], and was stable to train.

D. Backbone Feature Distribution

OOD detection performance is independent of classification accuracy, and is strongly affected by the distribution of feature representations produced by the backbone model. OOD detection performance may vary wildly between different classifiers of similar accuracy, and understanding which innate properties of neural networks (including model architectures, training hyperparameters, pretraining data distributions, and pretraining loss) improve the OOD detection task is understudied compared to the primary task of improving classification accuracy.

To investigate the factors that influence a backbone model’s OOD detection performance, we examine the feature space distribution using two metrics: uniformity and tolerance [40].

$$\text{Uniformity}(z; t) = \log \mathbb{E}_{x, y \sim p_{\text{data}}} \left[e^{-t \|z(x) - z(y)\|_2^2} \right]$$

$$\text{Tolerance}(z) = \mathbb{E}_{x, y \sim p_{\text{data}}} \left[(z(x)^T z(y)) \cdot \mathbf{1}_{\{l(x) \neq l(y)\}} \right]$$

Here, $l(x)$ is the supervised label of datapoint x . Uniformity uses a Gaussian similarity kernel to measure the spread

of features over the feature space hypersphere. We use a weight parameter of $t = 2$, and report negative uniformity consistent with [40]. Higher uniformity means more of the hypersphere is occupied by features produced by the model. Tolerance measures the cosine similarity of intra-class feature representations: tolerance is higher when class representations form tight clusters in feature space, and is lower when class representations are more diffuse in feature space. To investigate the connection between feature space distribution and the ability to detect OOD data using feature density estimation via a normalizing flow, we apply our method to 80 pretrained PyTorch ImageNet-1k classification models [38]. We evaluate the uniformity and tolerance of each classifier’s feature space for ID data, train a normalizing flow on this feature space, and evaluate the performance of our OOD detection method.

Visualizing the performance of our method (evaluated as ImageNet-1k vs. Textures AUROC) versus the uniformity and tolerance of 80 unique feature spaces (Figure 6) shows two strong correlations: AUROC is positively correlated with tolerance (tight class clustering, Pearson correlation coefficient $r = 0.72$), and negatively correlated with uniformity (volume of feature space occupied, $r = -0.68$). Our method is thus best applied to models with compact ID class representations occupying a lower volume of feature space. It is easier to fit a normalizing flow density model to these distributions, and they have an increased likelihood of OOD samples falling in the low density regions.

Additionally, we see that uniformity is correlated with classifier top-1 validation accuracy ($r = 0.50$), while no strong correlation exists between tolerance and top-1 validation accuracy ($r = -0.02$, Figure 6). A tradeoff is thus apparent for uniformity: high uniformity correlates with im-

Feature Space Uniformity and Tolerance vs.
ImageNet-1k Accuracy and Flow AUROC




Figure 6: Evaluating 80 pretrained PyTorch ImageNet-1k classifiers for feature space uniformity and tolerance, validation accuracy, and flow AUROC (ImageNet-1k vs. Textures). The Pearson correlation coefficient is annotated.

proved classification accuracy, but low uniformity improves OOD detection using feature density estimation. An ideal model may balance this by exhibiting high uniformity and high tolerance, encouraging both validation accuracy and OOD detection. Future work will investigate the role of training regularizations, model architectures, and pretraining loss and task in creating feature spaces amenable to both strong OOD detection and classification accuracy.

VII. CONCLUSION

For machine learning systems to be safely deployed in the open world, it is essential that out-of-distribution data can be accurately identified to safeguard against unintended model behavior. We investigate a method for out-of-distribution detection by performing density estimation in the feature space of pretrained image classification models using normalizing flows. In contrast with prior work in this space, our experiments show that feature density estimation via normalizing flows can achieve strong OOD detection performance on a variety of common benchmarks on large scale datasets. Our method outperforms all existing methods for detecting far-OOD data, as demonstrated by the results on CIFAR-10 vs. SVHN, and ImageNet-1k vs. Textures.

Performing density estimation on normalized feature vectors and under-training the normalizing flow are shown to be particularly important, and we observe the surprising behavior that OOD detection performance peaks very early

in flow training. We further show that OOD detection performance is not dependent on the flow model’s ability to perform high quality density estimation, but is strongly dependent on the distribution of feature representations of the backbone model. Specifically, evaluations of 80 pre-trained ImageNet-1k classifiers show that performance of our method is strongly correlated with the tolerance of the classifier’s feature space. Using the discussed techniques, we demonstrate that normalizing flows are effective tools for OOD detection, blazing a trail towards the safe deployment of machine learning and robotic systems in challenging open world environments.

REFERENCES

- [1] J. Yang, K. Zhou, Y. Li, and Z. Liu, “Generalized out-of-distribution detection: A survey,” *arXiv preprint arXiv:2110.11334*, 2021.
- [2] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan, “Do deep generative models know what they don’t know?” in *International Conference on Learning Representations*, 2019.
- [3] P. Kirichenko, P. Izmailov, and A. G. Wilson, “Why normalizing flows fail to detect out-of-distribution data,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 20 578–20 589.
- [4] T. Fawcett, “An introduction to roc analysis,” *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006, rOC Analysis in Pattern Recognition. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016786550500303X>
- [5] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” *Proceedings of International Conference on Learning Representations*, 2017.
- [6] S. Liang, Y. Li, and R. Srikant, “Enhancing the reliability of out-of-distribution image detection in neural networks,” in *International Conference on Learning Representations*, 2018.
- [7] W. Liu, X. Wang, J. Owens, and Y. Li, “Energy-based out-of-distribution detection,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 21 464–21 475.
- [8] A. Djurisic, N. Bozanic, A. Ashok, and R. Liu, “Extremely simple activation shaping for out-of-distribution detection,” in *The Eleventh International Conference on Learning Representations*, 2023.
- [9] Y. Sun, C. Guo, and Y. Li, “React: Out-of-distribution detection with rectified activations,” in *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021.
- [10] K. Lee, K. Lee, H. Lee, and J. Shin, “A simple unified framework for detecting out-of-distribution samples and adversarial attacks,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.
- [11] M. Salehi, H. Mirzaei, D. Hendrycks, Y. Li, M. H. Rohban, and M. Sabokrou, “A unified survey on anomaly, novelty, open-set, and out of-distribution detection: Solutions and future challenges,” *Transactions on Machine Learning Research*, 2022.

- [12] G. Floto, S. Kremer, and M. Nica, "The tilted variational autoencoder: Improving out-of-distribution detection," in *The Eleventh International Conference on Learning Representations*, 2023.
- [13] J. D. Havtorn, J. Frellsen, S. Hauberg, and L. Maaløe, "Hierarchical vaes know what they don't know," *CoRR*, vol. abs/2102.08248, 2021. [Online]. Available: <https://arxiv.org/abs/2102.08248>
- [14] M. S. Graham, W. H. Pinaya, P.-D. Tudosiu, P. Nachev, S. Ourselin, and J. Cardoso, "Denoising diffusion models for out-of-distribution detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2023, pp. 2947–2956.
- [15] Z. Liu, J. P. Zhou, Y. Wang, and K. Q. Weinberger, "Unsupervised out-of-distribution detection with diffusion inpainting," *arXiv preprint arXiv:2302.10326*, 2023.
- [16] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods," *Machine Learning*, vol. 110, no. 3, pp. 457–506, mar 2021.
- [17] I. Kobyzev, S. J. Prince, and M. A. Brubaker, "Normalizing flows: An introduction and review of current methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 11, pp. 3964–3979, 2021.
- [18] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.
- [19] J. Ho, X. Chen, A. Srinivas, Y. Duan, and P. Abbeel, "Flow++: Improving flow-based generative models with variational dequantization and architecture design," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 2722–2730. [Online]. Available: <https://proceedings.mlr.press/v97/ho19a.html>
- [20] L. H. Zhang, M. Goldstein, and R. Ranganath, "Understanding failures in out-of-distribution detection with deep generative models," *Proceedings of machine learning research*, vol. 139, pp. 12 427–12 436, 2021.
- [21] D. Gudovskiy, S. Ishizaka, and K. Kozuka, "Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2022, pp. 98–107.
- [22] M. Rudolph, B. Wandt, and B. Rosenhahn, "Same same but different: Semi-supervised defect detection with normalizing flows," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2021, pp. 1907–1916.
- [23] M. Rudolph, T. Wehrbein, B. Rosenhahn, and B. Wandt, "Fully convolutional cross-scale-flows for image-based defect detection," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2022, pp. 1088–1097.
- [24] M. Cho, T. Kim, W. J. Kim, S. Cho, and S. Lee, "Unsupervised video anomaly detection via normalizing flows with implicit latent features," *Pattern Recognition*, vol. 129, p. 108703, 2022.
- [25] C. M. Jiang, M. Najibi, C. R. Qi, Y. Zhou, and D. Anguelov, "Improving the intra-class long-tail in 3d detection via rare example mining," in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 158–175.
- [26] H. Zhang, A. Li, J. Guo, and Y. Guo, "Hybrid models for open set recognition," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 102–117.
- [27] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real NVP," in *International Conference on Learning Representations*, 2017.
- [28] R. T. Q. Chen, J. Behrmann, D. K. Duvenaud, and J.-H. Jacobsen, "Residual flows for invertible generative modeling," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.
- [29] A. Krizhevsky, "Learning multiple layers of features from tiny images," *University of Toronto*, 05 2012.
- [30] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and F.-F. Li, "Imagenet: a large-scale hierarchical image database," 06 2009, pp. 248–255.
- [31] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Ng, "Reading digits in natural images with unsupervised feature learning," *NIPS*, 01 2011.
- [32] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 6, pp. 1452–1464, 2018.
- [33] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [34] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi, "Describing textures in the wild," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3606–3613.
- [35] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie, "The inaturalist species classification and detection dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [36] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010, pp. 3485–3492.
- [37] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [38] T. P. Foundation, "Models and pre-trained weights - torchvision 0.17 documentation," <https://pytorch.org/vision/stable/models.html>, accessed 2024-01-31.
- [39] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, and D. Duvenaud, "Scalable reversible generative models with free-form continuous dynamics," in *International Conference on Learning Representations*, 2019.
- [40] F. Wang and H. Liu, "Understanding the behaviour of contrastive loss," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 2495–2504.