

BIGEMINAL PRIORS VARIATIONAL AUTO-ENCODER

Anonymous authors

Paper under double-blind review

ABSTRACT

Variational auto-encoders (VAEs) are an influential and generally-used class of likelihood-based generative models in unsupervised learning. The likelihood-based generative models have been reported to be highly robust to the out-of-distribution (OOD) inputs and can be a detector by assuming that the model assigns higher likelihoods to the samples from the in-distribution (ID) dataset than an OOD dataset. However, recent works reported a phenomenon that VAE recognizes some OOD samples as ID by assigning a higher likelihood to the OOD inputs compared to the one from ID. In this work, we introduce a new model, namely *Bigeminal Priors Variational auto-encoder (BPVAE)*, to address this phenomenon. The BPVAE aims to enhance the robustness of the VAEs by combining the power of VAE with the two independent priors that belong to the training dataset and simple dataset, which complexity is lower than the training dataset, respectively. BPVAE learns two datasets' features, assigning a higher likelihood for the training dataset than the simple dataset. In this way, we can use BPVAE's density estimate for detecting the OOD samples. Quantitative experimental results suggest that our model has better generalization capability and stronger robustness than the standard VAEs, proving the effectiveness of the proposed approach of hybrid learning by collaborative priors. Overall, this work paves a new avenue to potentially overcome the OOD problem via multiple latent priors modeling.

1 INTRODUCTION

Out-of-distribution (OOD) detection is a crucial issue for machine learning (ML) security, which usually arises in many application scenarios, such as medical diagnosis and credit card fraud detection. There is a widely held view that likelihood-based generative models have strong robustness to the OOD inputs Bishop (1994); Blei et al. (2017). Based on this opinion, a well-calibrated generative model can be a good detector by assigning higher likelihoods to the samples from the in-distribution (ID) dataset than OOD dataset. Hence, the deep generative models are generally considered as reliable for anomaly detection tasks (Chalapathy et al., 2018; Xu et al., 2018; Ostrovski et al., 2017). However, recent works Nalisnick et al. (2019a); Hendrycks et al. (2019); Choi et al. (2018); Lee et al. (2017); Nalisnick et al. (2019b); Huang et al. (2019); Maaløe et al. (2019) have reported the phenomenon that the density estimate of the deep generative model, in some cases, is not able to detect OOD inputs correctly. For instance, VAEs Kingma and Welling (2014); Rezende et al. (2014) cannot identify images of common objects such as airplane, bird, cat, and dog (i.e., CIFAR10) from the OOD datasets (i.e., MNIST, FashionMNIST, KMNIST, and GTSRB), assigning higher likelihoods to the OOD samples when VAE is trained on CIFAR10 (Shown in **Figure 1a**). These findings conflict with the previous OOD detection method proposed by Bishop (1994). To alleviate and resolve this issue, deep generative models are expected to understand the attributes of OOD data deeply and fully when utilizing density estimate detecting OOD samples.

A variety of works have emerged attempting to solve this problem. For instance, Serrà et al. (2020) demonstrated that the input complexity affects greatly the density estimate of deep generative models by designing controlled experiments on Glow model with different levels of image complexity Kingma and Dhariwal (2018). Similar qualitative results of VAEs are obtained in our experiments (**Figure 1**). Also, we computed the likelihoods of training samples from Cifar10, FashionMnist, GTSRB, IMGAENET, KMNIST, OMNIGLOT, and SVHN (Shown in **Figure 2a**). We find that the simple samples (KMNIST, OMNIGLOT, and MNIST) trained by VAEs with high likelihoods can assign lower likelihoods to the complex test samples (CIFAR10, SVHN, and IMGAENET) to

detect OOD samples because the likelihood of complex samples trained on VAEs is smaller than the likelihood of simple samples trained on VAEs. In contrast, the VAEs trained on complex samples with a low likelihood usually give a higher likelihood for the simple test samples but identify them as ID samples (**Figure 2b**). Inspired by this intriguing finding, here we propose a method that feeds the external dataset (called the simple dataset) as inputs while training VAEs on the training dataset (called the basic dataset), which is more straightforward than training VAE on the basic dataset. In this manner, VAEs can learn the features from two data distributions, assigning a higher likelihood for the basic dataset than the simple dataset. And the density estimate of VAEs can be used for detecting OOD samples.

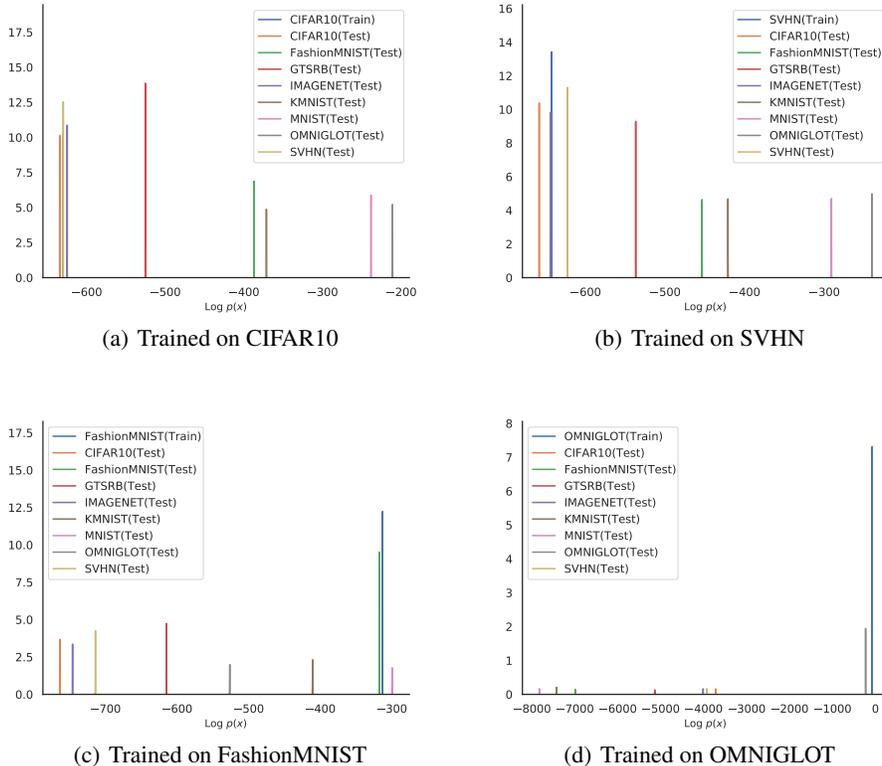


Figure 1: Histogram of log-likelihoods from a VAE model trained on CIFAR10, SVHN, FashionMNIST, and OMNIGLOT. (see similar results in Nalisnick et al. (2019a); Choi et al. (2018); Serrà et al. (2020)). Other added results are shown in Figure 8 in Appendix.

In this work, we introduce the Bigeminal Priors Variational auto-encoder (BPVAE) (**Figure 3**), an advanced extension of VAEs with two independent latent priors that belong to the basic and simple datasets, respectively. To build this hybrid model with an effective synergetic mode, two tricky problems arise. The first one is that how to choose the simple dataset and two priors for BPVAE. Due to a lot of other candidate datasets different from the basic dataset, it is hard to find the most appropriate candidate. here we firstly select a dataset randomly as a candidate and then train VAEs on the basic and candidate datasets, respectively. By comparing the likelihoods of samples from basic dataset with that of other candidate datasets, we can choose the right candidate dataset with higher likelihoods than the basic dataset as the simple dataset (**Figure 2a**). For example, take GTSRB as the basic dataset, then FashionMNIST, MNIST, and KMNIST can be regarded as the simple datasets, while CIFAR10, IMAGENET, and SVHN cannot be the simple datasets. On the other hand, there are plenty of candidate priors (e.g., standard normal distribution prior, Gaussian mixture prior Dilokthanakul et al. (2016), Vamp Prior Tomczak and Welling (2018), Resampled Prior Bauer and Mnih (2018), Reference prior Bernardo (1979); Berger et al. (2009); Nalisnick and Smyth (2017))

for BPVAE. How to combine adaptive priors remains a vital step. Note that the BPVAE has two priors in the latent space. A good prior to basic dataset is expected to carry the pivotal features of the basic dataset, which is called the basic prior (b-prior for short). And a good prior to the simple dataset (s-prior for short) should cover the core features of the simple dataset and follow a distribution different from that of the basic dataset. Therefore, BPVAE with b-prior assigns a lower likelihood for the simple dataset. Overall, the uncertainty of b-prior is higher than s-prior, for the complexity and uncertainty of the datasets are positively correlated.

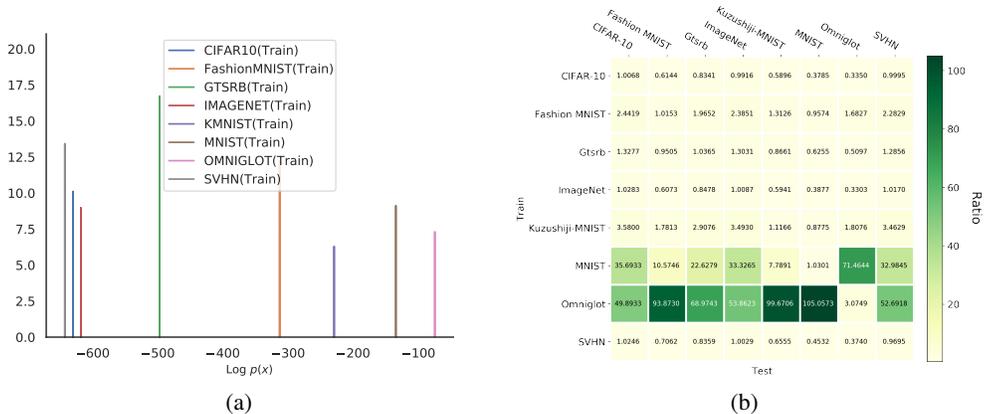


Figure 2: (a) Histogram of log-likelihoods of VAE trained on Cifar10, FashionMnist, GTSRB, IMGAENET, KMNIST(Kuzushiji-MNIST), OMNIGLOT, and SVHN, respectively. (b) Likelihood Ratios of training and testing samples (*the higher is better*). The Likelihood Ratios < 1 is represented by the likelihood of testing simple higher than training samples

2 RELATED WORK

Various works Nalisnick et al. (2019a); Choi et al. (2018); Hendrycks et al. (2019); Lee et al. (2017) have reported that the deep generative models are not able to correctly detect OOD samples until the models have an excellent understanding of OOD inputs. Maaløe et al. (2019) indicated that Bidirectional-Inference Variational Auto-encoder (BIVA) with the multiple latent layers could capture the high-level semantic information of the data, with a better understanding of OOD representation. However, the standard VAEs with one latent layer have a poor performance for anomaly detection. Choi et al. (2018) proposed an algorithm which takes the ensembles of generative models to estimate the Watanabe-Akaike Information Criterion (WAIC) Watanabe (2010) as a metric for outliers. Hendrycks et al. (2019) showed that robustness and uncertainty Malinin and Gales (2018); Hafner et al. (2018) concerning to the outlier exposure (OE) can be improved by training the model with OOD data, which can improve model calibration and several anomaly detection techniques were proposed accordingly. Ran et al. (2020) proposed an improved noise contrastive prior (INCP) to acquire a reliable uncertainty estimate for the standard VAEs. Patterns between OOD and ID inputs can be well captured and distinguished via VAEs with reliable uncertainty estimate. Nalisnick et al. (2019b) proposed a statistical method to test the OOD inputs using a Monte Carlo estimate of empirical entropy, but their approach is limited in the batches of inputs with the same type. Huang et al. (2019) tried to use other generative models (i.e., neural rendering model (NRM)) for OOD detection, and they found the joint likelihoods of latent variables to be the most effective one for OOD detection. Song et al. (2019) demonstrate that OOD detection failure can induce sophisticated statistics based on the likelihoods of individual samples; they proposed a method that is in-batch dependencies for OOD detection.

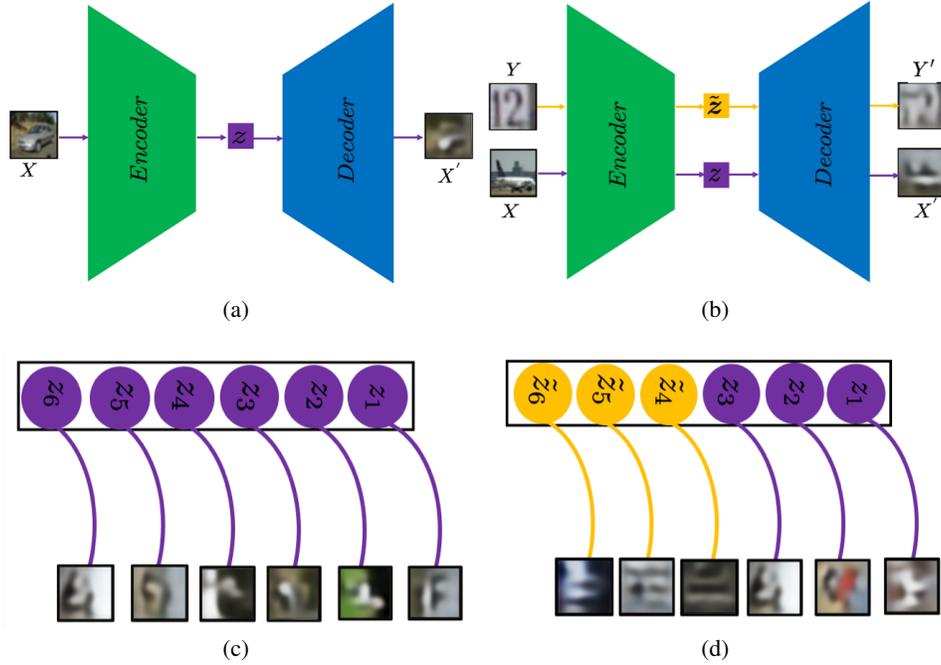


Figure 3: Overview framework of standard VAE (a) and BPVAE(b). The encoder and decoder are represented by a green and blue trapezoid, respectively. The purple and yellow squares denote the latent space of in-distribution(ID) and out-of-distribution(OOD) data, respectively. Compared the generator of standard VAE (c) and BPVAE(d). The latent space of standard VAE only learn and capture ID-distribution data, while the latent space of BPVAE cover the features of ID and OOD data concurrently.

3 APPROACH

3.1 VARIATIONAL AUTOENCODER

VAEs Rezende et al. (2014); Kingma and Welling (2014) are a variety of latent variable models optimized by the maximum marginal likelihood of an observation variable $p(\mathbf{x})$ **Figure 3a**. The marginal likelihood can be written as follows:

$$\begin{aligned} \log p(\mathbf{x}) = & \mathbf{E}_{z \sim q_{\theta}(z|\mathbf{x})} [\log p_{\phi}(\mathbf{x} | z)] - \mathbf{D}_{KL}[q_{\theta}(z | \mathbf{x}) \| p(z)] \\ & + \mathbf{D}_{KL}[q_{\theta}(z | \mathbf{x}) \| p(z | \mathbf{x})], \end{aligned} \quad (1)$$

where $p(z)$ and $p(z | \mathbf{x})$ are the prior by using a standard normal distribution and the true posterior respectively. $q_{\theta}(z | \mathbf{x})$ is the variational posterior (encoder) by employing a Gaussian distribution, and $p_{\phi}(\mathbf{x} | z)$ is the generative model (decoder) by using a Bernoulli distribution. Both are modeled by a neural network with their parameter θ, ϕ , respectively. Thus, we train VAE with training samples to maximize the following objective variational evidence lower bound (ELBO):

$$\mathcal{L}(\phi, \theta) = \mathbf{E}_{z \sim q_{\theta}(z|\mathbf{x})} [\log p_{\phi}(\mathbf{x} | z)] - \mathbf{D}_{KL}[q_{\theta}(z | \mathbf{x}) \| p(z)] \quad (2)$$

where $q_{\theta}(z | \mathbf{x})$ and $q_{\theta}(\tilde{z} | \tilde{\mathbf{x}})$ are variational posteriors for matching the true posteriors ($p(z | \mathbf{x})$ and $p(\tilde{z} | \tilde{\mathbf{x}})$) which are given by $\tilde{\mathbf{x}}$ and \mathbf{x} respectively. For a given dataset, the marginal likelihood is a constant. From Eq. 2 and Eq. 1, we get

$$\log p(\mathbf{x}) \geq \mathcal{L}(\phi, \theta) \quad (3)$$

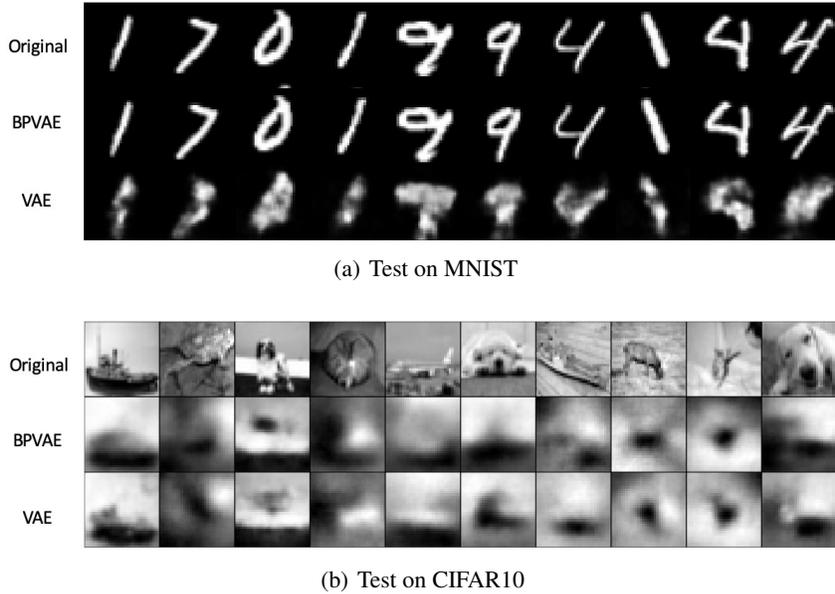


Figure 4: Reconstruction performance for MNIST and CIFAR10 by VAEs and BPVAEs. Here CIFAR10 is used as basic dataset and MNIST is used as simple dataset.

Assuming variational posterior has arbitrarily high-capacity for modeling, $q_\theta(z | \mathbf{x})$ approximates intractably $p(z | \mathbf{x})$ and the KL-divergence between $q_\theta(z | \mathbf{x})$ and $p(z | \mathbf{x})$ will be zero. The $\mathcal{L}(\phi, \theta)$ can be replaced by $\log p(\mathbf{x})$.

3.2 BIGEMINAL PRIORS VARIATIONAL AUTOENCODER

The BPVAE consists of an encoder, a decoder, and two priors (b-prior and s-prior) **Figure 3b**, which is trained on both the basic dataset learned by b-prior and the simple dataset learned by s-prior. Specifically, the uncertainty of b-prior is higher than s-prior due to the positive correlation between the complexity and uncertainty of the dataset. We assume that both the b-prior and s-prior belong to normal distribution. And we use the variance of a normal distribution to represent the uncertainty level. The priors are formulated as followings,

$$\begin{aligned} p_b(z) &\sim \mathcal{N}(z | \mu_z, \sigma_z^2 I) \\ p_s(\tilde{z}) &\sim \mathcal{N}(\tilde{z} | \mu_{\tilde{z}}, \sigma_{\tilde{z}}^2 I) \end{aligned} \quad (4)$$

where the mean value $\mu_z = \mu_{\tilde{z}} = 0$. And the variances σ_z and $\sigma_{\tilde{z}}$ are hyper-parameters determining the uncertainty of b-prior and s-prior. σ_z is always set to be greater than $\sigma_{\tilde{z}}$ so that b-prior has enough capacity to capture the basic dataset features. In this manner, BPVAE can capture the features of basic dataset, as well as of simple dataset. To facilitate the training implementation, we modified the loss function as follow:

$$\begin{aligned} \log p(\mathbf{x}) + \log p(\mathbf{y}) &= \mathbf{E}_{\mathbf{z} \sim q_\theta(\mathbf{z} | \mathbf{x})} [\log p_\phi(\mathbf{x} | \mathbf{z})] - \mathbf{D}_{KL} [q_\theta(\mathbf{z} | \mathbf{x}) \| p_b(\mathbf{z})] \\ &\quad + \mathbf{E}_{\mathbf{z} \sim q_\theta(\mathbf{z} | \mathbf{y})} [\log p_\phi(\mathbf{y} | \mathbf{z})] - \mathbf{D}_{KL} [q_\theta(\mathbf{z} | \mathbf{y}) \| p_b(\mathbf{z})] \end{aligned} \quad (5)$$

where $q_\theta(\tilde{z} | \mathbf{y})$ and $q_\theta(\mathbf{z} | \mathbf{x})$ are the variational posterior for the simple and basic dataset, $q_\theta(\tilde{z} | \mathbf{y})$ and $q_\theta(\mathbf{z} | \mathbf{x})$ are the decoder for the simple and basic data, and which are modeled by a neural network with their parameters θ and ϕ , respectively.

Table 1: Evaluation on the basic dataset and the simple dataset

	Method	MSE	PSNR	SSIM
Evaluation on basic dataset	BPVAE	0.017	18.250	0.544
	VAE	0.016	18.282	0.543
Evaluation on simple dataset	BPVAE	0.007	22.392	0.909
	VAE	0.0346	14.831	0.601

4 RESULTS

4.1 DOES BPVAE KNOW WHAT IT DOESN'T KNOW?

To investigate whether VAEs have a good understanding of the distribution of training data, we carry out reconstruction experiments under multiple conditions. Despite the variety of choices for settings of reconstruction experiment, here we take the following setting as an example: CIFAR10 as the basic dataset for training and MNIST as the simple dataset. After training VAEs and BPVAEs separately, we generated reconstructed images during the inference process. As shown in **Figure 4**, we visualized the results of standard VAEs and our proposed BPVAEs in comparison. It is evident that BPVAEs obtain much better performance than standard VAEs on MNIST, while these two models achieve comparable results on CIFAR10. The great performance of BPVAEs on MNIST can be attributed to the effective capacity of the extra introduced s-priors, which can assist BPVAEs of capturing external feature representation for the data from simple dataset, in which case VAEs failed due to lack of various latent priors with strong capacity.

Besides, we evaluate the reconstruction effects quantitatively by using MSE (Mean Squared Error), PSNR (Peak Signal to Noise Ratio) and SSIM (Structural Similarity) Horé and Ziou (2010). Note that as for PSNR, the value lower is better, and as for PSNR and SSIM, the value higher is better.

Results for comprehensive comparisons between two models are presented in **Table 1**. The tables demonstrate that BPVAEs can obtain much better performance than standard VAEs no matter it is evaluated by MSE, PSNR or SSIM, meanwhile retaining the comparable capacity to capture and reconstruct data from basic dataset, which is consistent with the aforementioned descriptions from qualitative observation.

4.2 ANALYSIS

To explore the internal mechanism of BPVAEs, we perform some comparison experiments with multiple OOD testing samples from the different data distribution. As described in **Figure 5**, we train the BPVAE model on CIAFR10 (basic dataset), meanwhile with other datasets as simple dataset. Take **Figure 5a** as an example, after the training process on CIFAR10 and FashionMNIST, BPVAEs can output lower likelihoods for most low-complexity datasets (such as MNIST, SVHN, etc.), avoiding the excessive-high likelihood problem for the simple data. This is because VAEs with the hybrid prior mode are equipped with a higher capacity of representation and therefore can shift the input distribution towards the lower-likelihood direction. On the contrary, for GTSRB and KMNIST in this case, BPVAEs fail to transform their distribution into representation in lower space of data distribution. **Figure 5b,c,d** present similar phenomenon. This illustrates that although adding extra prior can indeed facilitate the VAEs' robustness and representation capacity, alleviating OOD problem by shifting data distribution of low-complexity dataset, but the distribution scale where it can cover is not infinite, which usually lies in the nearby neighborhood from the data distribution captured by b-prior and s-prior.

In order to alleviate the aforementioned limitation, we propose a more comprehensive approach to broaden its applicable scale. As **Figure 6** presents, our model can cover all key representation and shift all the data distribution toward the lower-likelihood area, via combining multiple priors and training BPVAEs on a variety of selected datasets. This is presumably helpful for OOD detection as well, and we will show its performance on OOD detection tasks in the next section.

4.3 OOD DETECTION

We perform OOD detection experiments on FashionMNIST and CIFAR10 datasets. For gray images, we train VAEs on FashionMNIST, train BPVAEs on FashionMNIST (basic dataset) and OMNIGLOT (simple dataset). And then we conduct OOD test with MNIST data as inputs. For RGB images, we train VAEs on CIFAR10, train BPVAEs on CIFAR10 (basic dataset) and GTSRB (simple dataset). And then we conduct OOD test with SVHN data as inputs. As depicted in Table 2 and 3, our BPVAEs can achieve higher AUROC and AUPRC values than Standard VAEs, meanwhile surpassing other classical baselines. Overall, these comprehensive comparisons suggest that our proposed model is equipped with strong robustness and detection capability.

Table 2: AUROC and AUPRC for detecting OOD inputs using likelihoods of BPVAE, likelihood of VAE, and other baselines on FashionMNIST vs. MNIST datasets.

Model	AUROC	AUPRC
BPVAE(ours)	1.000	1.000
Standard VAE	0.012	0.113
Likelihood Ratio(μ, λ) Ren et al. (2019)	0.994	0.993
ODIN Liang et al. (2018)	0.752	0.763
Mahalanobis distance Lee et al. (2018)	0.942	0.928
Ensemble, 20 classifiers Lakshminarayanan et al. (2017)	0.857	0.849
WAIC,5 models Choi et al. (2018)	0.221	0.401

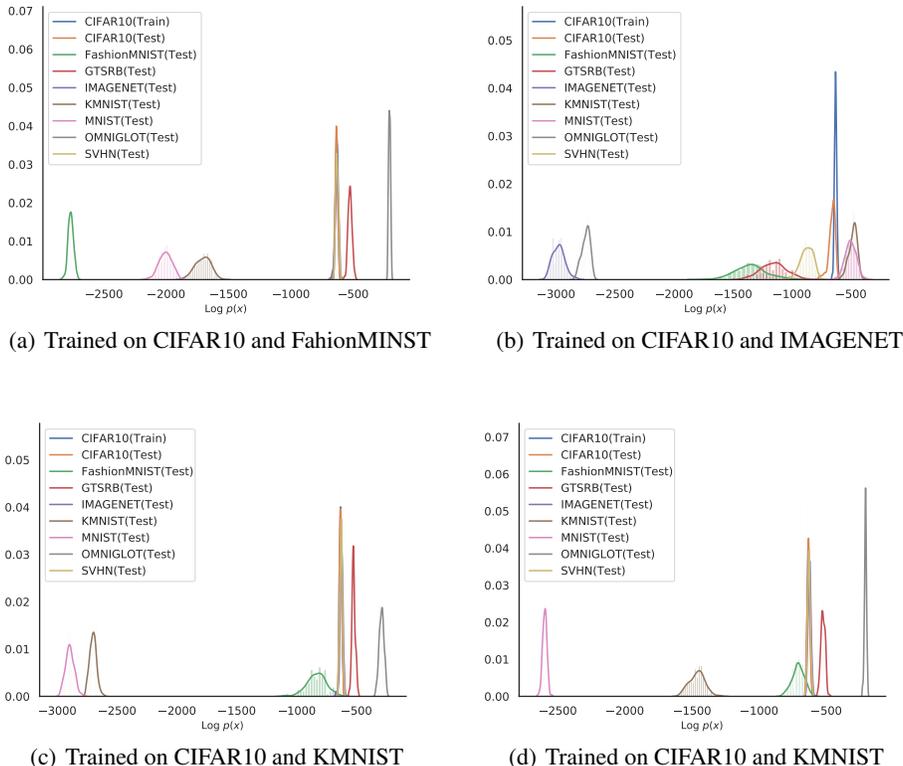


Figure 5: Histogram of log-likelihoods from VAEs model, which are trained on different groups of datasets.

Table 3: AUROC and AUPRC for detecting OOD inputs using likelihoods of BPVAE and VAE, and other baselines on CIFAR10 vs. SVHN datasets.

Model	AUROC	AUPRC
BPVAE(ours)	1.000	1.000
Standard VAE	0.037	0.214
Likelihood Ratio(μ, λ) Ren et al. (2019)	0.930	0.881
ODIN Liang et al. (2018)	0.938	0.926
Mahalanobis distance Lee et al. (2018)	0.728	0.711
Ensemble, 20 classifiers Lakshminarayanan et al. (2017)	0.946	0.916
WAIC,5 models Choi et al. (2018)	0.146	0.363

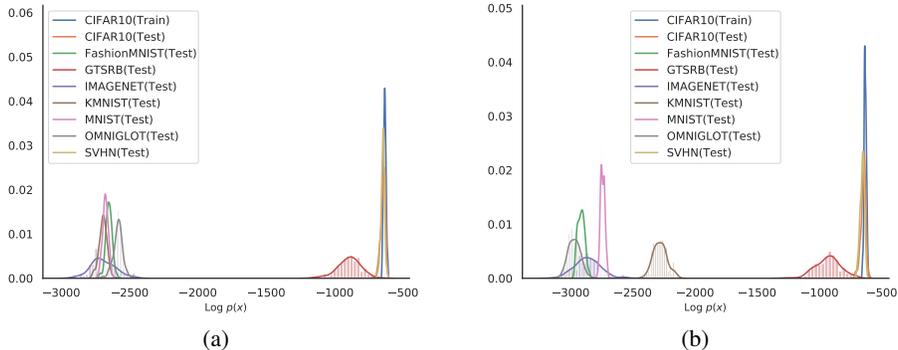


Figure 6: Histogram of log-likelihoods from VAEs model, which are trained on different groups of datasets. (a) Trained on CIFAR10(Basic), FashionMNIST(simple), and KMNIST(simple); (b) Trained on CIFAR10(Basic), FashionMNIST(simple), and MNIST(simple);

5 DISCUSSION

OOD problem has been increasingly gaining attention and interests, which remains an intriguing property and a challenging issue for likelihood-based generative models. In this work, we introduced external latent priors to assist VAEs in capturing more abstract representations for data which not belong to in-distribution. Through building an effective synergistic mode, VAEs can obtain powerful representation ability for different data from various datasets. In this manner, VAEs can be well-calibrated by shifting the likelihood distribution of data with simpler complexity to lower-likelihood intervals compared to basic dataset, in which way the high-likelihoods problem of OOD can be overcome to a large extent. Interestingly, we find there is a trivial trade-off when employing detection tasks, that is, even this method can alleviate OOD problem to a great extent, the likelihood interval scale which can be covered by bridging two latent priors is a little limited. Hence we introduce a hybrid VAE version with multiply latent priors, which can alleviate the trade-off greatly. Besides, we only impose the proposed approach on VAE model, designing the hybrid latent priors for other models like Glow, PixelCNN Van den Oord et al. (2016) will be an interesting research topic. And we are expected to continue related exploration further. Overall, from a brand-new perspective, this work provides a potential way to tackle the OOD problem intertwined with VAEs.

REFERENCES

M. Bauer and A. Mnih. Resampled priors for variational autoencoders. *arXiv preprint arXiv:1810.11428*, 2018.

J. O. Berger, J. M. Bernardo, and D. Sun. The formal definition of reference priors. *Annals of Statistics*, 37: 905–938, 2009.

- J. M. Bernardo. Reference posterior distributions for bayesian inference. 1979.
- C. M. Bishop. Novelty detection and neural network validation. *IEE Proceedings-Vision, Image and Signal processing*, 141(4):217–222, 1994.
- D. Blei, K. Heller, T. Salimans, M. Welling, and Z. Ghahramani. Panel discussion. advances in approximate bayesian inference. *NeurIPS Workshop*, 2017.
- R. Chalapathy, E. Toth, and S. Chawla. Group anomaly detection using deep generative models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 173–189. Springer, 2018.
- H. Choi, E. Jang, and A. A. Alemi. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018.
- N. Dilokthanakul, P. A. M. Mediano, M. Garnelo, M. J. Lee, H. Salimbeni, K. Arulkumaran, and M. Shanahan. Deep unsupervised clustering with gaussian mixture variational autoencoders. *ArXiv*, abs/1611.02648, 2016.
- D. Hafner, D. Tran, A. Irpan, T. Lillicrap, and J. Davidson. Reliable uncertainty estimates in deep neural networks using noise contrastive priors. *arXiv preprint arXiv:1807.09289*, 2018.
- D. Hendrycks, M. Mazeika, and T. G. Dietterich. Deep anomaly detection with outlier exposure. *International Conference on Learning Representations (ICLR)*, 2019.
- A. Horé and D. Ziou. Image quality metrics: Psnr vs. ssim. In *International Conference on Pattern Recognition*, 2010.
- Y. Huang, S. Dai, T. Nguyen, R. G. Baraniuk, and A. Anandkumar. Out-of-distribution detection using neural rendering generative models. *arXiv preprint arXiv:1907.04572*, 2019.
- D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *International Conference on Learning Representations (ICLR)*, 2014.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems (NeurIPS)*, 2017.
- K. Lee, H. Lee, K. Lee, and J. Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. *arXiv preprint arXiv:1711.09325*, 2017.
- 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 (NeurIPS)*, 2018.
- S. Liang, Y. Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *International Conference on Learning Representations (ICLR)*, 2018.
- L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther. Biva: A very deep hierarchy of latent variables for generative modeling. *arXiv preprint arXiv:1902.02102*, 2019.
- A. Malinin and M. Gales. Predictive uncertainty estimation via prior networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7047–7058, 2018.
- E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan. Do deep generative models know what they don’t know? *International Conference on Learning Representations (ICLR)*, 2019a.
- E. Nalisnick, A. Matsukawa, Y. W. Teh, and B. Lakshminarayanan. Detecting out-of-distribution inputs to deep generative models using a test for typicality. *arXiv preprint arXiv:1906.02994*, 2019b.
- E. T. Nalisnick and P. Smyth. Learning approximately objective priors. *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence*, 2017.
- G. Ostrovski, M. G. Bellemare, A. van den Oord, and R. Munos. Count-based exploration with neural density models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2721–2730. JMLR. org, 2017.
- X. Ran, M. Xu, L. Mei, Q. Xu, and Q. Liu. Detecting out-of-distribution samples via variational auto-encoder with reliable uncertainty estimation. *ArXiv*, abs/2007.08128, 2020.

- J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, M. Depristo, J. Dillon, and B. Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286, 2014.
- J. Serrà, D. Álvarez, V. Gómez, O. Slizovskaia, J. F. Núñez, and J. Luque. Input complexity and out-of-distribution detection with likelihood-based generative models. *International Conference on Learning Representations (ICLR)*, 2020.
- J. Song, Y. Song, and S. Ermon. Unsupervised out-of-distribution detection with batch normalization. *arXiv preprint arXiv:1910.09115*, 2019.
- J. Tomczak and M. Welling. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pages 1214–1223, 2018.
- A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016.
- S. Watanabe. Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(Dec):3571–3594, 2010.
- H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei, Y. Feng, et al. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. In *Proceedings of the 2018 World Wide Web Conference*, pages 187–196. International World Wide Web Conferences Steering Committee, 2018.

A DATA PREPROCESSING FOR DIFFERENT DATASETS

We conducted some preprocessing operations in order to unify the training condition for different datasets. For datasets with RGB images, such as CIFAR10, GTSRB, IMAGENET,SVHN, we transformed these data into gray images with single channel. Then we reshaped all the images into the size (32, 32, 1). We visualized the preprocessed data from eight datasets in **Figure 7**. It is evident that the data from CIFAR10, GTSRB,IMAGENET and SVHN have the higher complexity then that from simpler datasets such as FashionMNIST, OMNIGLOT, KMNIST and MNIST.

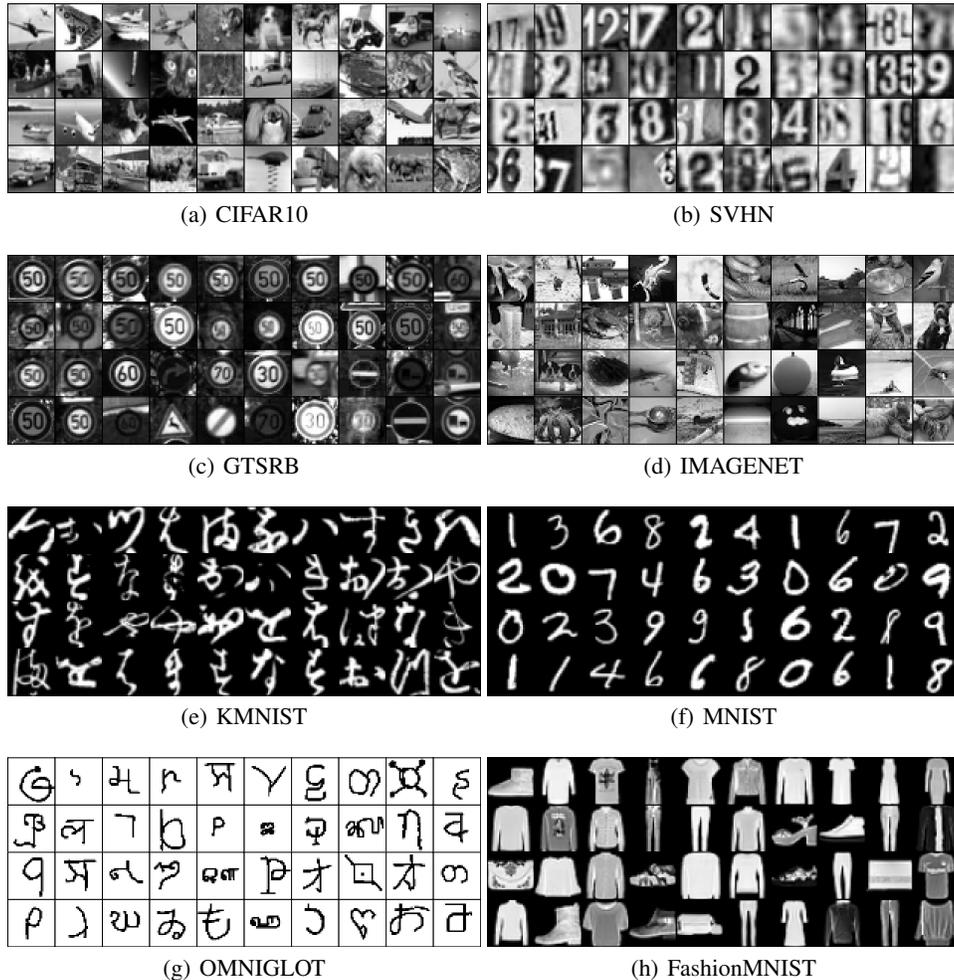


Figure 7: Presentation of preprocessed data from eight datasets, which are CIFAR10, SVHN, GTSRB, IMAGENET, KMNIST, MNIST, OMNIGLOT and FashionMNIST.

B LOG-LIKELIHOODS OF OTHER DATASETS

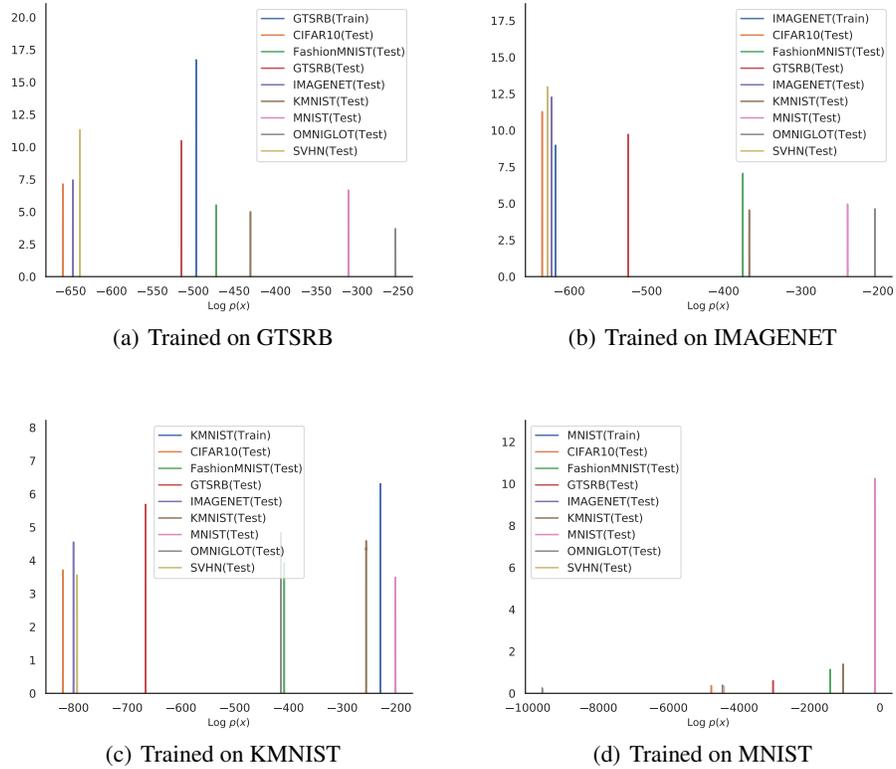


Figure 8: Histogram of log-likelihoods from a VAE model trained on GTSRB, IMAGENET, KMNIST and MNIST.

C SETTINGS FOR IMPLEMENTATION DETAIL

In the experiments, VAE and BPVAE are trained with images normalized to $[0, 1]$ on $1 \times$ NVIDIA TITAN RTX GPU. In all experiments, VAE and BPVAE consist of an encoder with the architecture given in Ran et al. (2020) and a decoder shown in Ran et al. (2020). Both VAE and BPVAE use Leaky Relu activation function. We train the VAE and BPVAE for 200 epochs with a constant learning rate $1e^{-4}$, meanwhile using Adam optimizer and batch size 64 in each experiment.