

VaB-AL

Choi, Jongwon; Yi, Kwang Moo; Kim, Jihoon; Choo, Jinho; Kim, Byoungjip; Chang, Jinyeop; Gwon, Youngjune; Chang, Hyung Jin

DOI:

[10.1109/CVPR46437.2021.00668](https://doi.org/10.1109/CVPR46437.2021.00668)

License:

Other (please specify with Rights Statement)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Choi, J, Yi, KM, Kim, J, Choo, J, Kim, B, Chang, J, Gwon, Y & Chang, HJ 2021, VaB-AL: incorporating class imbalance and difficulty with variational Bayes for active learning. in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9577648, Proceedings. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE, pp. 6745-6754, 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, Tennessee, United States, 20/06/21.
<https://doi.org/10.1109/CVPR46437.2021.00668>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is the accepted manuscript of J. Choi et al., "VaB-AL: Incorporating Class Imbalance and Difficulty with Variational Bayes for Active Learning," 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2021, pp. 6745-6754, doi: 10.1109/CVPR46437.2021.00668. in 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

© 2021IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

VaB-AL: Incorporating Class Imbalance and Difficulty with Variational Bayes for Active Learning

Jongwon Choi^{1*}, Kwang Moo Yi^{2*}, Jihoon Kim³, Jinho Choo³,
Byoungjip Kim³, Jinyeop Chang³, Youngjune Gwon³, Hyung Jin Chang⁴

¹Chung-Ang University, South Korea

³Samsung SDS, South Korea

²University of British Columbia, Canada

⁴University of Birmingham, United Kingdom

choijw@cau.ac.kr, kmyi@cs.ubc.ca, h.j.chang@bham.ac.uk

{jihoon00.kim, jinhol2.choo, bjip.kim, jinyeop.chang, gyj.gwon}@samsung.com

Abstract

Active Learning for discriminative models has largely been studied with the focus on individual samples, with less emphasis on how classes are distributed or which classes are hard to deal with. In this work, we show that this is harmful. We propose a method based on the Bayes' rule, that can naturally incorporate class imbalance into the Active Learning framework. We derive that three terms should be considered together when estimating the probability of a classifier making a mistake for a given sample; i) probability of mislabelling a class, ii) likelihood of the data given a predicted class, and iii) the prior probability on the abundance of a predicted class. Implementing these terms requires a generative model and an intractable likelihood estimation. Therefore, we train a Variational Auto Encoder (VAE) for this purpose. To further tie the VAE with the classifier and facilitate VAE training, we use the classifiers' deep feature representations as input to the VAE. By considering all three probabilities, among them, especially the data imbalance, we can substantially improve the potential of existing methods under limited data budget. We show that our method can be applied to classification tasks on multiple different datasets – including one that is a real-world dataset with heavy data imbalance – significantly outperforming the state of the art.

1. Introduction

Active learning focuses on efficient labelling of data and has drawn much interest lately [38, 41, 48], due to deep learning being attempted at new domains, such as biomedical imaging [3, 16] and industrial imaging [27, 49], where ac-

* These authors contributed equally to this work.

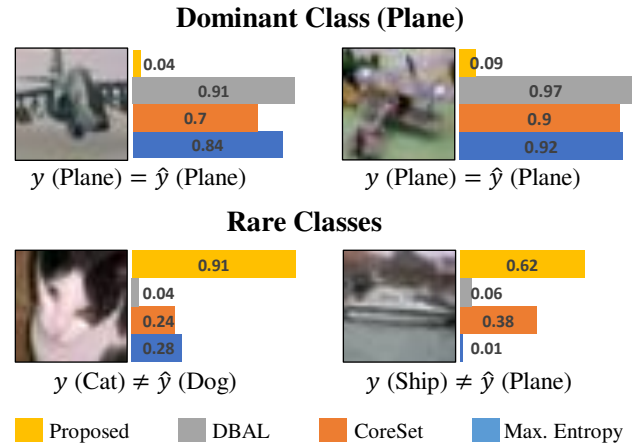


Figure 1. **Class matters** – We show an example where existing methods fail to identify important samples to put into the updated training set. We show the importance metrics given by various methods – DBAL [10], CoreSet [38], Max. Entropy [45], and our method – for selected images when the dataset is dominated by plane images. We further show below each image, the ground-truth class y and the predicted class \hat{y} . Because of the data imbalance, existing methods heavily favour the plane class, even when the results are correct, as seen on the left. And for rare classes in the right column, they fail to recognise that the samples are important, which leads to the improved performance. On the other hand, our method correctly identifies them as important, by considering also the class difficulty and the class prior.

quiring data can be costly [39]. Even for cases where data is not scarce, the effective usage of data may reduce training time, therefore the computational cost, including carbon foot-prints required to train each model. There have been various studies based on semi-supervised [7, 21] and unsupervised [37, 46] learning schemes to improve the training efficiency of data. However, with limited labelling budget, the performance of the studies are significantly worse to the

supervised learning with the additionally labelled data [35]. In other words, their label efficiency could be improved.

Existing methods [4, 10, 38, 41, 48], regardless of how they are formulated, have a common underlying assumption that all classes are equal – they do not consider that some classes might just be harder to learn compared to others, or some classes might be more prevalent in the dataset than others. Instead, they focus on, given a data sample, how much error a trained model is expected to make, or the estimated uncertainties [10, 48]. These assumptions could be harmful, as in practice, since data is often imbalanced and not all classes are of the same difficulty [1, 51]. This can create a bias in the labelled data pool, leading to the trained classifier and active learning methods also being biased in deciding which samples to label. As we show in Fig. 1, this can damage the capabilities of an active learning method significantly, even for typical benchmark datasets [6, 22].

In this work, we present a novel formulation for active learning, based on the classical Bayes’ rule that allows us to incorporate multiple factors of a classification network together. Through derivation, we show that the probability of a classifier making mistakes can be decomposed into three terms; i) the probability of misclassification for a given predicted class, ii) the likelihood of a sample given predicted class, and iii) the prior probability of a class being predicted. In other words, one needs to take into account i) the difficulty of a class, ii) the performance of the classifier and iii) the abundance of a certain class of data holistically when determining the potential of a classification error. We take all of them into account and choose samples to be labelled by selecting those that have the highest misclassification probability.

While the task is discriminative, our method requires the estimation of likelihood, which could be intractable. We, therefore, propose to use a Variational Auto Encoder (VAE) [20] to model the lower bound of the likelihood of a sample. To make VAE conditioned on a predicted label, a naive way would be applied to train multiple VAEs for each predicted class. However, this quickly becomes impractical with a large number of classes. We thus propose to train a single VAE, with regularisation that acts as conditioning based on the predicted label. To further tie the VAE with the classifier, and for quick training of the VAE, we use the deep feature representations of the classifier as inputs to the VAE. Being generative, this training of VAE does not involve any labels, and we utilise the vast amount of unlabelled data with their predicted labels while inferring probabilities that are independent of data samples on the labelled ones.

We show empirically in Section 4 that our method allows a significant leap in performance for the benchmark dataset and the real application dataset, especially when the labelling budget is highly limited. Furthermore, we posit that considering the prior, that is, the distribution of labels

during training, is critical when analysing the uncertainty of estimates.

In summary, our contributions are four-fold:

- we derive a novel formulation for active learning based on the Bayes’ rule and posterior probability;
- we propose a framework based on VAE that realises this formulation;
- we reveal that considering the differences between classes – abundance and difficulty – is important;
- we outperform the state of the art in the various experiments.

2. Related Works

Traditional Active Learning. Increasing the label efficiency, thus reducing the cost associated with obtaining labels, has been of interest for decades. Even before deep learning became popular, various methods were suggested towards this goal [39]. Methods were proposed to estimate the uncertainty of the unlabelled samples through the probability of prediction [24], the difference between the best prediction and the second one [25, 30, 36], or the entropy covering all the possible classes [17, 40]. For support vector machine classifiers the methods were suggested to utilise the distance from the decision boundary, for both the classification task [26, 43] and the detection task [44]. The algorithms clustering and finding representative samples were also suggested as another way [2, 12, 32].

Discrete optimisation algorithms have been proposed to consider the relationship between the sampling result and the model performance [8, 11, 47]. In a voting scheme-based algorithm [31], multiple different models are trained by the labelled pool, which determines the next queries according to their disagreement. Despite these efforts, the classical approaches are geared towards simple features and may hold limitations when applying to a large deep network with many nonlinear estimations.

Active Learning for deep networks. Active learning algorithms for deep networks can be categorised into uncertainty-based methods and representation-based methods. The uncertainty-based methods aim to select the uncertain samples from the unlabelled data pool and annotate them to increase the labelled pool [4, 10, 48]. Yoo and Kwon [48] proposed to use a small auxiliary “module” network predicting the training loss of the baseline network that is being trained with the active learning scheme. They then select the samples that are expected to give high losses. While their method is similar to ours in that an additional network is trained, as they require a ground-truth loss value while training, the auxiliary network can only be trained with labelled data, creating yet another network that the performance depends on how data is sampled. In contrast, we train our VAE with unlabelled data since we only rely on the pre-

dicted labels from the baseline network during training, and result in stable performance even with few labelled data. Gal *et al.* [10] proposed a method based on the estimation of sample-wise posterior probability through a Bayesian deep learning [9] framework. The method can be implemented simply by locating several dropout layers in a deep network, but this increases training time significantly until the convergence. Beluch *et al.* [4] suggest an active sampling method that estimates the disagreement of the prediction by using multiple deep networks. The downside of their method is that, as they use multiple networks, the memory and the computational requirement increases proportionally.

The representation-based methods target on finding representative samples within the high-dimensional space that deep networks learn [38, 41]. Sener and Savarese [38] proposed the Coreset algorithm that determines representative samples by using the feature maps of the intermediate layers of a deep network, rather than the last layer. However, the optimisation method in the Coreset algorithm does not scale well as the number of classes, and the number of unlabelled samples grows. To improve the scalability, Sinha *et al.* [41] proposed to map the high dimensional feature maps into a lower dimension through adversarial training. Unfortunately, being based on adversarial training, the method requires a large amount of training data for the mapping to be reliable.

Beyond them, hybrid approaches combine the best of both uncertainty-based and representation-based methods [34, 50]. Some works focused on a specific task: for example person re-identification [29] and a human pose estimation [28].

While our work is most similar to uncertainty-based methods, it falls into neither uncertainty-based nor representation-based methods. Contrary to the previous uncertainty-based works, we take into account characteristics that are not restricted to a single sample – we consider the class difficulty and class imbalance. Also, unlike the representation-based methods, we are not aiming to find representative samples, but a global trend of samples that are predicted to belong to a certain class.

Semi-supervised learning with VAEs. As we utilise VAEs [20], we also briefly review works related to VAEs that associate them with labelled data. Since VAEs model the likelihood of data, Lee *et al.* [23] used them to identify out-of-distribution samples for each class. We are loosely inspired by them, as we also use conditioned VAEs. However, unlike them, we estimate one portion of our conditional probabilities in estimating the label correctness. M2 VAE models [19] and Conditional VAEs [42] have been proposed to model conditional distributions. They directly add the condition as an additional latent dimension that is trained independently with the other latent dimensions for the reconstruction. In contrast, we apply conditioning implicitly during training to represent the class information and the feature distribution in the same latent dimensions. In our early

attempts, we were not able to obtain successful modelling for our application with the former.

3. Methodology

We first formally describe the active learning problem and detail how we select new samples to be labelled with the estimated correctness of predictions. We then derive our method via Bayes' rule and describe how our derivation can be implemented in practice through a VAE. Afterwards, we provide a summary of how we tie every component together as an active learning algorithm and implementation details.

3.1. Problem formulation

Active learning can be formulated as a problem of increasing the pool of labelled data at every round by labelling subsets of the unlabelled pool. Formally, given a pool of data \mathcal{P} , we keep a pool of labelled data \mathcal{P}_L and unlabelled data \mathcal{P}_U , such that $\mathcal{P} = \mathcal{P}_L \cup \mathcal{P}_U$ and $\emptyset = \mathcal{P}_L \cap \mathcal{P}_U$. Then, for each active learning round r , we select $\mathcal{P}_S^{(r)}$ that is a subset of \mathcal{P}_U with N_r samples according to a criterion that defines the active learning method, which is moved from \mathcal{P}_U to \mathcal{P}_L . Thus, $\mathcal{P}_U^{(r+1)} = \mathcal{P}_U^{(r)} - \mathcal{P}_S^{(r)}$ and $\mathcal{P}_L^{(r+1)} = \mathcal{P}_L^{(r)} + \mathcal{P}_S^{(r)}$.

The core of active learning is how $\mathcal{P}_S^{(r)}$ is selected, which in our case is based on the probability of a model providing a wrong answer for the given data.

Active learning based on the probability of wrong predictions. The underlying idea of our method is that when acquiring data with a limited labelling budget, one should acquire those that have the highest probability of making wrong predictions [10, 25, 29, 48]. Formally, if we let y denote the real label and \hat{y} the label predicted by a model, we find samples \mathbf{x} with large

$$p(y \neq \hat{y}|\mathbf{x}). \quad (1)$$

Unfortunately, this is not a probability distribution that can be modelled directly, and we, therefore, estimate this probability via Bayes' rule and approximations.

3.2. Bayesian active learning

We now derive our Bayesian formulation. To estimate $p(y \neq \hat{y}|\mathbf{x})$, we take an alternative route through Bayes' rule, instead of the direct estimation that could be given by a discriminative model. We first represent this probability with its complement, which can then be written as the sum of joint probabilities. We write

$$p(y \neq \hat{y}|\mathbf{x}) = 1 - p(y = \hat{y}|\mathbf{x}) = 1 - \sum_{n=1}^{N_c} p(y_n, \hat{y}_n|\mathbf{x}), \quad (2)$$

where N_c is the number of classes and $p(y_n)$ and $p(\hat{y}_n)$ is a shorthand for $p(y = n)$ and $p(\hat{y} = n)$, respectively.

Then, each $p(y_n, \hat{y}_n | \mathbf{x})$ within the summation can be written through Bayes' rule as following:

$$p(y_n, \hat{y}_n | \mathbf{x}) = \frac{p(y_n | \hat{y}_n, \mathbf{x}) p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}{\sum_{n=1}^{N_c} p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}. \quad (3)$$

Here, $p(y_n | \hat{y}_n, \mathbf{x})$ corresponds to the probability of the real label being n , given that the predicted label is n and the data being \mathbf{x} . However, this is a probability that cannot be evaluated unless we have the paired true label. Thus, we instead do an informed guess, by ignoring \mathbf{x} and approximating with $p(y_n | \hat{y}_n)$. In other words, we assume that the probability of a model making a mistake is highly related to the label. Thus we approximate by writing

$$p(y_n, \hat{y}_n | \mathbf{x}) \approx \frac{p(y_n | \hat{y}_n) p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}{\sum_{n=1}^{N_c} p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}. \quad (4)$$

Finally, with Eq. (2), we have

$$p(y \neq \hat{y} | \mathbf{x}) \approx 1 - \sum_{n=1}^{N_c} \frac{p(y_n | \hat{y}_n) p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}{\sum_{n=1}^{N_c} p(\mathbf{x} | \hat{y}_n) p(\hat{y}_n)}. \quad (5)$$

Note that here, i) $p(y_n | \hat{y}_n)$ is the probability of a model making mistake based on label, ii) $p(\mathbf{x} | \hat{y}_n)$ is the *likelihood* of a sample given predicted label, iii) $p(\hat{y}_n)$ is the *prior* on the distribution of predicted labels, which represents how imbalanced the predictions of a model are. As mentioned earlier in Section 1, the likelihood estimation is non-trivial and requires a generative model. Furthermore, through our Bayesian formulation, the confirmation bias of pseudo-labels can be also alleviated. The prior terms would encode whether such a bias exists, and compensate for it. As our formulation is probabilistic, the predictions would not be affected as drastically as other heuristic methods. We now detail how we model these three probabilities.

3.3. Estimating probabilities with regularized VAE

To estimate the probabilities, we use a VAE [20]. Before we discuss the details, let us first clarify that we train this VAE exclusively in the unlabelled data pool \mathcal{P}_U , and associate it with the behaviour of the discriminative model on unseen data. We do not use the labelled pool \mathcal{P}_L , as it is likely that the classifier has overfitted to the data. Note also that while we explain in terms of \mathbf{x} , in fact, we use the deep feature representations given by the classifier to tie the VAE more with the classifier and to facilitate training by removing the need of learning the deep features. See Fig. 2 for an illustration of our framework.

We first detail why and how we estimate the likelihood with a VAE and then discuss the other two probabilities.

Likelihood of a sample – $p(\mathbf{x} | \hat{y}_n)$. Estimating $p(\mathbf{x} | \hat{y}_n)$ is not straightforward. A naive idea would be to implement




Figure 2. **Framework** – An illustration of our framework. We train a Variational Auto Encoder (VAE) that models the deep feature representation of the classifier with unlabelled data. We then use the VAE to estimate $p(\mathbf{x} | \hat{y}_n)$, $p(y_n | \hat{y}_n)$, and $p(\hat{y}_n)$, which in turn is used to estimate the probability of labelling error, $p(y \neq \hat{y} | \mathbf{x})$. See Section 3.3 for details.

multiple generative models that model $p(\mathbf{x})$, each trained with labels predicted to be of a certain class. However, this becomes quickly impractical as the number of classes grows. Moreover, estimating $p(\mathbf{x})$ can be intractable in practice [20].

In our work, we use a *single* VAE to estimate the lower bound of $p(\mathbf{x} | \hat{y}_n)$ for all \hat{y}_n . We use a VAE, as it learns to reconstruct the data sample using the tractable lower bound for $p(\mathbf{x})$. Distinct from existing work, to condition the $p(\mathbf{x})$ based on predicted labels, we propose to learn a latent space where the *absence* of parts of the latent embeddings are related to the predicted label. In other words, this is as if we are training multiple VAEs with shared weights and overlapping latent spaces. Once such latent embeddings are learned, we compute the lower bound of $p(\mathbf{x} | \hat{y}_n)$, by simply enforcing the absence manually via masking – thus selecting a VAE dedicated to a certain predicted class among the multiple virtual VAEs – and computing $p(\mathbf{x})$. This strategy allows us to have a manageable latent space while still being able to deal with many classes.

In more detail, if we denote the j -th embedding dimension of VAE as z_j and write $j \in C_n$ to denote dimension j is related to class n , we write this absence condition as

$$\hat{y} = \underset{n}{\operatorname{argmin}} \left[\sum_{j \in C_1} z_j^2, \sum_{j \in C_2} z_j^2, \dots, \sum_{j \in C_{N_c}} z_j^2 \right]^T. \quad (6)$$

Notice how this condition is conceptually similar to disentangled representations [5, 42]. In our earlier attempts, we have also tried forming this condition as disentangled representations or enforcing $\sum_{j \in C_n} z_j^2$ to be zero if \hat{y}_n , which neither was successful. We suspect that enforcing such constraints limit the capacity of the latent space and interferes with the training of the VAE too much. We have also tried other ways of enforcing absence – using the $\ell - 1$ norm or the sigmoid – but using the square worked best. We provide empirical results in Section 4.4.

We enforce this constraint as a form of regularisation. Let $\mathbf{w} = [w_1, w_2, \dots, w_n]^T$, where $w_n = \sum_{j \in C_n} z_j^2$, then we form an additional regularisation loss $\mathcal{L}_{\text{Class}}$ to be used

during training of VAE as

$$\mathcal{L}_{\text{Class}} = \mathcal{H}(\text{softmax}(-\mathbf{w}), \hat{\mathbf{y}}), \quad (7)$$

where \mathcal{H} denotes the cross entropy, softmax is the softmax, and $\hat{\mathbf{y}}$ is the one-hot encoded vector representation of \hat{y} . With this regularisation term, recall from [20], that the training loss for VAEs \mathcal{L}_{VAE} is the inverse of the empirical lower bound (ELBO) of the likelihood, which is defined as

$$\mathcal{L}_{\text{VAE}} = -\mathbb{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}|\mathbf{z})] + D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})), \quad (8)$$

where $p_{\theta}(\mathbf{x}|\mathbf{z})$ is the decoder with parameters θ and $q_{\phi}(\mathbf{z}|\mathbf{x})$ is the encoder with parameters ϕ . Therefore, the total loss to train our VAE is

$$\mathcal{L} = \mathcal{L}_{\text{VAE}} + \lambda \mathcal{L}_{\text{Class}}, \quad (9)$$

where λ is a hyperparameter that controls the regularisation strength. Note that with this loss, the VAE now also tries to mimic the behaviour of the classifier.

Once the VAE is trained, this VAE – without any conditioning – is now able to estimate the lower bound of the likelihood of a given data $p(\mathbf{x})$ [20] by simply computing the inversed value of \mathcal{L}_{VAE} . Furthermore, by masking the embedding space associated with $\hat{\mathbf{y}}$ with zero, we can compute the lower bound of $p(\mathbf{x}|\hat{\mathbf{y}})$. To avoid the decoder going too far from the latent embeddings, it was trained for; we use Eq. (6) to obtain $\hat{\mathbf{y}}$, instead of the one from the classifier.

Probability of labelling error – $p(y_n|\hat{y}_n)$. While the labelled pool \mathcal{P}_L is the only set of data that we have labels for, as the trained classifier is likely to have overfitted to \mathcal{P}_L , we cannot use it for modelling this probability. We, therefore, use the labels given by the VAE for the data samples in the labelled pool \mathcal{P}_L . Note that the VAE has never seen these data points during training. Mathematically, if we denote the i -th sample as $\mathbf{x}^{(i)}$, its label as $y^{(i)}$, the predicted label from Eq. (6) as $\hat{y}^{(i)}$, and introduce an indicator function δ that is 1 if all inputs are equal and 0 otherwise, we write

$$p(y_n|\hat{y}_n) \approx \frac{\mathbb{E}_{\mathbf{x} \in \mathcal{P}_L, \mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\delta(y^{(i)}, \hat{y}^{(i)}, n)]}{\mathbb{E}_{\mathbf{x} \in \mathcal{P}_L, \mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\delta(\hat{y}^{(i)}, n)]}. \quad (10)$$

Here, we approximate the expectations with Monte Carlo estimates. Note that we also take the expectation over \mathbf{z} , to take into account the stochastic nature of VAEs.

Prior – $p(\hat{y}_n)$. The prior is also acquired by using the labelled samples included in \mathcal{P}_L as this probability should be related to how the classifier is trained. Same as in the case of $p(y_n|\hat{y}_n)$, we cannot use the classifier predictions for the labelled pool, and we use the predictions from the VAE. Thus, sharing the notations as in Eq. (10), we write

$$p(\hat{y}_n) \approx \mathbb{E}_{\mathbf{x} \in \mathcal{P}_L, \mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} [\delta(\hat{y}^{(i)}, n)]. \quad (11)$$

Algorithm 1: Proposed Method

```

Input:  $\mathcal{P}_U^{(0)}, \mathcal{P}_L^{(0)}, N_r, \mathcal{M}$ 
 $r = 0$  while not at the maximum round do
    Train  $\mathcal{M}$  using  $\mathcal{P}_L^{(r)}$ 
    Freeze  $\mathcal{M}$ 
    Train the VAE module using  $\mathcal{P}_U^{(0)}$  by Eq. (9)
    Estimate  $p(\mathbf{x}|\hat{y}_n)$  by Eq. (8)
    Estimate  $p(y_n|\hat{y}_n)$  by Eq. (10)
    Estimate  $p(\hat{y}_n)$  by Eq. (11)
    Estimate  $p(y \neq \hat{y}|\mathbf{x})$  by Eq. (2)
     $\mathcal{X} \leftarrow N_r$  samples with the highest uncertainty
     $\mathcal{P}_U^{(r+1)} \leftarrow \mathcal{P}_U^{(r)} - \mathcal{X}$ 
     $\mathcal{P}_L^{(r+1)} \leftarrow \mathcal{P}_L^{(r)} \cup \mathcal{X}$ 
     $r \leftarrow r + 1$ 
end

```

3.4. Summary and implementation details

We summarise our method in Algorithm 1. For each Active Learning round r , we train a classifier $\mathcal{M}^{(r)}$ with the labelled pool $\mathcal{P}_L^{(r)}$. We then freeze the weights of the classifier, and train our VAE – $p_{\theta}^{(r)}(\mathbf{x}|\mathbf{z})$ and $q_{\phi}^{(r)}(\mathbf{z}|\mathbf{x})$ – with the unlabelled pool $\mathcal{P}_U^{(r)}$. We then estimate the three probabilities, which we use to construct the final estimate of $p(y \neq \hat{y}|\mathbf{x})$, and sample those from $\mathcal{P}_U^{(r)}$ that have the highest value.

As noted earlier, we use the deep features of the baseline network (the classifier) as input to the VAE. Specifically, we extract deep features from the specific four layers of baseline. If the feature representations are from fully-connected layers, we use it as is. If it is from a convolutional layer, thus a feature map, we first apply global average pooling. We then apply batch normalisation to each feature, so that they are roughly in the same range, and feed it to a fully connected layer with 128 neurons and Sigmoid activation. Finally, the outputs from these fully-connected layers are concatenated to form a 512×1 vector and given as input to the VAE.

For the architecture of VAEs, we opt for a simple one as the deep features already have abstracted context. We apply four layers of fully connected layers with again 128 neurons and ReLU activations, with the exception of the layer that outputs the latent embeddings. For this layer, we use $10 \times N_c$ weights and dedicate 10 dimensions per each class.

We implement our method with PyTorch [33]. To train our VAE, we set $\lambda=0.005$ in Eq. (9), for all tasks. We use Adam optimiser [18] with a learning rate of 10^{-4} and default parameters. We train for 20 epochs. To remove randomness from experiments, we perform our training multiple times with different initial conditions. We report both the average and the standard deviation of our experiments. For inference, to perform the Monte Carlo estimation in Eqns. (8), (10), and (11), we iterate the inference 100 times for every sample.

4. Experimental Results

We first introduce the dataset, then the experimental setup, including the classification network and the compared baselines. We then report our results and discuss the effect of the individual probability terms through an ablation study.

4.1. Dataset

To validate our experiments we apply our method to the standard *CIFAR-10* and *CIFAR-100* [22] datasets as well as the Northeastern University surface defect classification dataset (NEU) [14]. We use the CIFAR datasets to be comparable with recent Active Learning studies [41], and NEU to demonstrate that class imbalance is important in real-world applications. In more detail, the *CIFAR-10* dataset contains 60,000 images sized by $32 \times 32 \times 3$ with 10 classes. The *CIFAR-100* dataset also consists of 60,000 images of the $32 \times 32 \times 3$ size with 100 classes. Both the *CIFAR-10* and the *CIFAR-100* datasets are well balanced – all included images are assigned exactly one class among the 10 or the 100, and the number of instances per each class is equal. The datasets come with a pre-determined training and evaluation split of 50,000 and 10,000 images, respectively. The goal of the NEU dataset is to classify the 9 defect types on steels. The main feature of this dataset is that it is heavily imbalanced – the number of instances per class varies from 200 to 1589. In total, the NEU dataset contains 7226 defect images of size 64×64 pixel. This dataset does not provide pre-determined splits, and we thus randomly reserve 20% of the dataset for evaluation. This results in 5778 training samples and 1448 validation samples.

Introducing synthetic class imbalance. To demonstrate the importance of considering class-wise probabilities, we build additional variants of *CIFAR* datasets by removing a part of the samples in *CIFAR-10* and *CIFAR-100*. First, we build four datasets that have dominant classes within them – we hereafter refer to them as *dominant* datasets. A real-world example would be when scraping data from the internet – there will be many images of people, cats, and dogs, but not so many of platypus. The first three dominant datasets are built from *CIFAR-10* by randomly removing 90% samples of every category except for the $\{1, 5, 10\}^{th}$ classes, respectively. For *CIFAR-100*, as there are many classes, there are only a few instances each – 500 for each class in the training set. We, therefore, build the last dominant dataset by removing 40% of the original samples for all categories other than the middle ones from 45^{th} class to 55^{th} class, from the *CIFAR-100* dataset. We denote these dominant datasets as *CIFAR-10*^{+[1]}, *CIFAR-10*^{+[5]}, *CIFAR-10*^{+[10]}, and *CIFAR-100*^{+[45:55]}, respectively.

We further consider the case when some samples are rare. This would be, for example, cases where there are rare events, such as accidents or defects that need to be

discovered. Similar to the *dominant* datasets, we build *rare* datasets by taking certain classes away from *CIFAR* datasets. Specifically, we use three variants, where we remove 90% of the samples that correspond to the $\{1, 5, 10\}^{th}$ classes. We denote these as *CIFAR-10*^{-[1]}, *CIFAR-10*^{-[5]}, and *CIFAR-10*^{-[10]}, respectively.

With these datasets, we run different active learning setups, based on the size of the dataset. For NEU, we run five active learning rounds, where each round samples 250 samples. For *CIFAR-10* based ones, we run six active learning round, which 500 samples each. For *CIFAR-100*, we run five rounds with 1000 samples each.

4.2. Experimental setup

The baseline classification network. As the baseline classifier, we utilise ResNet-18 [13]. This network is composed of an initial simple convolution layer, followed by four basic residual blocks. As discussed previously in Section 3.4, the output feature maps from these four residual blocks are used to form the input to our VAE. At every Active Learning round, we train the network for 200 epochs, with a batch size of 128. We train with a typical setup for classification on *CIFAR-10*: SGD with the momentum of 0.9, weight decay of 0.0005, the learning rate is also initialised as 0.1 and decreased to 0.01 after 160 epochs. We repeat the same experiments with the different random seeds three times to remove the randomness from our experiments.

The competitors. To demonstrate the effectiveness of the proposed algorithm, we compare our method with the state-of-the-art for active learning. We consider VAAL [41], DBAL [10], and Core-set [38]. We utilise the author’s implementation for VAAL and Core-set. For DBAL we integrate the author’s implementation to our framework to utilise our baseline classification network – this method is architecture dependant. In more detail, we utilise the authors’ code for uncertainty estimation and embed it into ours. Among the multiple methods to estimate the uncertainty in DBAL, we use the Bayesian Active Learning by Disagreement (BALD) [15] as the estimation method – it showed the best generality in our experiments. For the dropout layers of DBAL, we place them before each the residual blocks of the ResNet and fix the dropout ratio to 0.25. As suggested by the authors [10], we further use 100 samples to estimate the uncertainty. In addition, we also consider uniform random sampling as a baseline. Finally, for a fair comparison, after new samples are extracted by these methods, we train the networks with the same random seed to exclude any random surprises.

4.3. Comparison results

NEU dataset – a real-world imbalanced dataset. We first compare different methods on the NEU dataset, which is a typical case of an imbalanced real-world dataset. We report




Figure 3. Results for the dominant, rare, and full datasets.

the results in Fig. 3(a). As shown, the proposed method delivers improved performance, especially for the early iterations. Methods here mostly converge after 5 rounds, but the small difference in the later rounds are simply the results of using most of the available data. Hence, focus should be in the earlier stages, where our method dominates. Specifically, on the second round, our method performs **8.3%** better than the second best method that is *CoreSet*. Interestingly, for this small dataset, VAAL performs worse than simple random selection because VAAL requests a lot of labels to *bake in*.

Dominant datasets. We now consider the dominant variants. In the dominant situation, a large part of the entire training samples is dedicated to a few specific categories – in the case of *CIFAR-10*, just one. In Fig. 3(b), we average the results from *CIFAR-10*^[1], *CIFAR-10*^[5], and *CIFAR-10*^[10] that we run three times each – a total of nine experiments. As shown, the proposed method outperforms the compared methods by a significant margin. Especially, we perform **4.0%** better than the second best in the third round. This gap comes from the fact that the dataset is biased, and this causes the active learning methods also to become biased. However, our method successfully mitigates this class imbalance and provides improved performance.

A similar phenomenon happens when there are more classes. In Fig. 3(c), we report results for *CIFAR-100*^[45:55]. Here, it is worth noting that all compared methods perform worse than *Random*. This is because the dataset has 100 classes, and it is easy to start ignoring a certain class. Nonetheless, our method is the only method that provides improved label efficiency. To mitigate this, existing works have only studied when there is a sufficient number of labels from the beginning – for example in VAAL [41], 10% of the entire set. However, this greatly limits the applicability of active learning methods, as 10% is sometimes already too much to afford. In appendix A.1, we present the additional results of our method with the different classifier.

Rare datasets. We further test the case when some classes are rare. In Fig. 3(d), we show the average performance of each method on the three rare datasets – *CIFAR-10*^[1], *CIFAR-10*^[5], and *CIFAR-10*^[10]. Similar to the dominant case, our method shows the best performance among the compared methods. Interestingly, in this case, similar to the results with *CIFAR-100*^[45:55], existing methods perform worse than the random baseline. Again, this is because existing methods can break down easily without enough labelled samples from the beginning. This is a limitation our method does not have.

On the full dataset. For completeness, we further study how methods perform with the full *CIFAR-10* and *CIFAR-100* datasets. However note that, as we demonstrated previously, these datasets are with perfect data balance, which is not the case for real-world datasets. We report a summary of these results in Figures Fig. 3(e) and (f). Our method performs comparable to existing methods for *CIFAR-10* and outperforms existing methods for *CIFAR-100*. However, experiment for *CIFAR-100* shows a limitation of active learning methods including ours, that when there are too many classes and very few examples, their performances are close to random, and sometimes even worse. Nonetheless, compared to existing methods, our method performs best, even in this extreme situation.

Additional results. In the supplementary document, we provide the original plots for all experiments before averaging. From them, we can confirm that the proposed method outperforms existing methods consistently for various situations.

4.4. Ablation study

4.4.1 Effectiveness of $p(\hat{y})$ and $p(y_n|\hat{y})$.

To validate their effectiveness, we perform an ablation study by excluding one or both of the prior ($p(\hat{y})$) and the label difficulty ($p(y|\hat{y})$) – we artificially set either to 1. We use

$p(\hat{y})$	$p(y \hat{y})$	<i>CIFAR-10</i>		<i>CIFAR-10</i> ^{+[1]}	
		avg.	final	avg.	final
-	-	82.56%	88.66%	48.02%	60.13%
✓	-	82.91%	90.68%	54.23%	70.25%
-	✓	82.71%	90.51%	51.86%	65.98%
✓	✓	83.36%	91.12%	54.16%	70.35%

Table 1. Validity of the prior $p(\hat{y})$ and label difficulty $p(y|\hat{y})$.




Figure 4. **Top samples for *CIFAR-10*^{+[1]}** – DBAL (Top), CoreSet (Second row), and ours (Third row). The dominant class *plane* is shown in red. Our method creates a balanced set by considering the class imbalance, whereas other methods are susceptible to it.

CIFAR-10 and *CIFAR-10*^{+[1]} for these experiments to remove randomness. To avoid tuning on the test set, we use the validation splits for these experiments. We take the average performance over all active learning rounds and also the average final performance. We summarise the results in Table 1. We report both the average performance over all six active learning rounds (avg.) and the performance of the final round (final). We run each experiment three times and report the average. As shown, all terms contribute to performance improvement. We observe that all terms contribute to providing the best performance. Among the two terms, the prior ($p(\hat{y})$) provides more gain compared to the label difficulty ($p(y|\hat{y})$), demonstrating that it is critical to take data imbalance into account. We show an example of the two terms in action in Fig. 4. Our method balances the training data even when the dataset is imbalanced.

4.4.2 Robustness for various budgets and λ values

To analyse the sensitivity of the proposed framework against the user-defined parameters, we report the result of our method with various sampling budgets and λ values in Eq. (9): see Fig. 5(a). Regardless of the choice of the budget size, our method outperforms all compared methods. Smaller budget size tends to allow methods to react faster to the training outcomes and increase the effectiveness of active learning.




Figure 5. Results of ablation tests on *dominant* CIFAR-10.

In Fig. 5(b) we report the performance of our method with varying λ . We test with various λ , which do affect the performance of the proposed method to some degree. However, as before in the case of the budget, our method outperforms all compared methods regardless of the choice of λ . This further hints that the two-loss terms that we train for in Eq. (9) do not compete. Note that $\lambda = 0$ completely disables $p(\hat{y})$ and $p(y|\hat{y})$, as \hat{y} estimation becomes meaningless, which causes the method to perform significantly worse as shown earlier in Table 1, and we have therefore left it out in this figure to reduce clutter; see supplementary appendix for full results.

5. Conclusion

We have proposed a novel active learning method that incorporates class imbalance and label difficulty. Our method was derived through the Bayes’ rule, which results in three types of probabilities – data likelihood, label prior, label difficulty – being considered together. We implement our method via a VAE, that is regularised to behave as a conditional VAE. We have shown that this creates a significant difference for a real-world dataset that exhibits data imbalance, as well as in cases when data imbalance is introduced to CIFAR-10 and CIFAR-100 datasets.

While we limit our experiments to classification in this work, our method is application agnostic. In the future, we plan to extend our work to other discriminative tasks, for example, object detection and segmentation.

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, Compute Canada, and the Chung-Ang University Research Grants in 2020.

References

- [1] S. Al-Stouhi and C. K. Reddy. Transfer learning for class imbalance problems with inadequate data. *Knowledge and information systems*, 2016.
- [2] O. M. Aodha, N. Campbell, J. Kautz, and G.J. Brostow. Hierarchical subquery evaluation for active learning on a graph. In *Conference on Computer Vision and Pattern Recognition*, 2014.
- [3] A. Atzeni, M. Jansen, S. Ourselin, and J.E. Iglesias. A probabilistic model combining deep learning and multi-atlas segmentation for semi-automated labelling of histology. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2018.
- [4] W. Beluch, T. Genewein, A. Nurnberger, and J. Kohler. The power of ensembles for active learning in image classification. In *Conference on Computer Vision and Pattern Recognition*, 2018.
- [5] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner. Understanding disentangling in β -vae. *stat*, 2018.
- [6] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011.
- [7] J. Donahue, J. Hoffman, E. Rodner, K. Saenko, and T. Darrell. Semi-Supervised Domain Adaptation with Instance Constraints. In *Conference on Computer Vision and Pattern Recognition*, pages 668–675, 2013.
- [8] E. Elhamifar, G. Sapiro, A. Yang, and S.S. Sastry. A convex optimization framework for active learning. In *International Conference on Computer Vision*, 2013.
- [9] Y. Gal and Z. Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *International Conference on Machine Learning*, pages 1050–1059, 2016.
- [10] Y. Gal, R. Islam, and Z. Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, 2017.
- [11] Y. Guo. Active instance sampling via matrix partition. In *Advances in Neural Information Processing Systems*, 2010.
- [12] M. Hasan and A.K. Roy-Chowdhury. Context aware active learning of activity recognition models. In *International Conference on Computer Vision*, 2015.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [14] Y. He, K. Song, Q. Meng, and Y. Yan. An end-to-end steel surface defect detection approach via fusing multiple hierarchical features. *IEEE Transactions on Instrumentation and Measurement*, 2019.
- [15] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- [16] Y. Huo, Z. Xu, K. Aboud, P. Parvathaneni, S. Bao, C. Bermudez, S. Resnick, L.E. Cutting, and B.A. Landman. Spatially localized atlas network tiles enables 3d whole brain segmentation from limited data. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2018.
- [17] A. Joshi. Multi-class active learning for image classification. In *Conference on Computer Vision and Pattern Recognition*, 2009.
- [18] D.P. Kingma and J. Ba. Adam: A Method for Stochastic Optimisation. In *International Conference on Learning Representations*, 2015.
- [19] D.P. Kingma, S. Mohamed, D.J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *NIPS*, 2014.
- [20] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- [21] T.N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [23] 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*, pages 7167–7177, 2018.
- [24] D.D. Lewis and J. Catlett. Heterogeneous uncertainty sampling for supervised learning. *Machine Learning Proceedings*, 1994.
- [25] X. Li and Y. Guo. Adaptive active learning for image classification. In *Conference on Computer Vision and Pattern Recognition*, 2013.
- [26] X. Li and Y. Guo. Multi-level adaptive active learning for scene classification. In *European Conference on Computer Vision*, 2014.
- [27] H. Lin, B. Li, X. Wang, Y. Shu, and S. Niu. Automated defect inspection of led chip using deep convolutional neural network. *Journal of Intelligent Manufacturing*, 2019.
- [28] B. Liu and V. Ferrari. Active learning for human pose estimation. In *International Conference on Computer Vision*, 2017.
- [29] Z. Liu, J. Wang, S. Gong, H. Lu, and D. Tao. Deep reinforcement active learning for human-in-the-loop person re-identification. In *International Conference on Computer Vision*, 2019.
- [30] W. Luo, A. Schwing, and R. Urtasun. Latent structured active learning. In *Advances in Neural Information Processing Systems*, 2013.
- [31] A.K. McCallumzy and K. Nigamy. Employing em and pool-based active learning for text classification. In *International Conference on Machine Learning*, 1998.
- [32] H.T. Nguyen and A. Smeulders. Active learning using preclustering. In *International Conference on Machine Learning*, 2004.
- [33] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *Advances in Neural Information Processing Systems - Workshop*, 2017.

- [34] S. Paul, J. H. Bappy, and A. K. Roy-Chowdhury. The power of ensembles for active learning in image classification. In *Conference on Computer Vision and Pattern Recognition*, 2017.
- [35] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-Supervised Learning with Ladder Networks. In *Advances in Neural Information Processing Systems*, pages 3546–3554, 2015.
- [36] D. Roth and K. Small. Margin-based active learning for structured output spaces. In *European Conference on Machine Learning*, 2006.
- [37] K. Saito, Y. Ushiku, and T. Harada. Asymmetric Tri-Training for Unsupervised Domain Adaptation. In *International Conference on Machine Learning*, 2017.
- [38] O. Sener and S. Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018.
- [39] B. Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114, 2012.
- [40] B. Settles and M. Craven. An analysis of active learning strategies for sequence labeling tasks. In *Conference on Empirical Methods in Natural Language Processing*, 2008.
- [41] S. Sinha, S. Ebrahimi, and T. Darrell. Variational adversarial active learning. In *International Conference on Computer Vision*, 2019.
- [42] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 2015.
- [43] S. Tong and D. Koller. Support vector machine active learning with applications to text classification. *Journal of Machine Learning Research*, 2001.
- [44] S. Vijayanarasimhan and K. Grauman. Large-scale live active learning: Training object detectors with crawled data and crowds. *International Journal of Computer Vision*, 2014.
- [45] D. Wang and Y. Shang. A new active labeling method for deep learning. In *International Joint Conference on Neural Networks*, 2014.
- [46] H. Yan, Y. Ding, P. Li, Q. Wang, Y. Xu, and W. Zuo. Mind the Class Weight Bias: Weighted Maximum Mean Discrepancy for Unsupervised Domain Adaptation. In *Conference on Computer Vision and Pattern Recognition*, pages 2272–2281, 2017.
- [47] Y. Yang, Z. Ma, F. Nie, X. Chang, and A.G. Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 2015.
- [48] D. Yoo and I. Kweon. Learning loss for active learning. In *Conference on Computer Vision and Pattern Recognition*, 2019.
- [49] Y. J. Zhao, Y. H. Yan, and K. C. Song. Vision-based automatic detection of steel surface defects in the cold rolling process: considering the influence of industrial liquids and surface textures. *The International Journal of Advanced Manufacturing Technology*, 2017.
- [50] Z. Zhou, J. Shin, L. Zhang, S. Gurudu, M. Gotway, and J. Liang. The power of ensembles for active learning in image classification. In *Conference on Computer Vision and Pattern Recognition*, 2017.
- [51] T. Zhu, Y. Lin, and Y. Liu. Synthetic minority oversampling technique for multiclass imbalance problems. *Pattern Recognition*, 2017.