

DIBS: Diversity Inducing Information Bottleneck in Model Ensembles

Samarth Sinha^{*1,2}, Homanga Bharadhwaj^{*1,2}, Anirudh Goyal³,
Hugo Larochelle³, Animesh Garg^{1,2}, Florian Shkurti^{1,2}

¹ University of Toronto

² Vector Institute

³ Mila

⁴ Google Brain

samarth.sinha@mail.utoronto.ca, homanga@cs.toronto.edu

Abstract

Although deep learning models have achieved state-of-the-art performance on a number of vision tasks, generalization over high dimensional multi-modal data, and reliable predictive uncertainty estimation are still active areas of research. Bayesian approaches including Bayesian Neural Nets (BNNs) do not scale well to modern computer vision tasks, as they are difficult to train, and have poor generalization under dataset-shift. This motivates the need for effective ensembles which can generalize and give reliable uncertainty estimates. In this paper, we target the problem of generating effective ensembles of neural networks by encouraging diversity in prediction. We explicitly optimize a diversity inducing adversarial loss for learning the stochastic latent variables and thereby obtain diversity in the output predictions necessary for modeling multi-modal data. We evaluate our method on benchmark datasets: MNIST, CIFAR100, TinyImageNet and MIT Places 2, and compared to the most competitive baselines show significant improvements in classification accuracy, under a shift in the data distribution and in out-of-distribution detection. : over 10% relative improvement in classification accuracy, over 5% relative improvement in generalizing under dataset shift, and over 5% better predictive uncertainty estimation as inferred by efficient out-of-distribution (OOD) detection.

Introduction

Deep Neural Networks (DNNs) have achieved state-of-the-art performance in a wide variety of vision tasks, where the goal is to perform a single task efficiently (He et al. 2016; Zagoruyko and Komodakis 2016; Yu, Koltun, and Funkhouser 2017; He et al. 2017). However, most state-of-the-art approaches in computer vision, train a single network for solving a particular task, which may not generalize when there is a change in the input distribution during evaluation. Related to the issue of generalization, the notion of predictive uncertainty quantification remains an open problem. To achieve this, it is important for the learned model to be uncertainty-aware, or to *know what it does not know*. One of the ways of estimating this is to show the network out-of-distribution (OOD) examples, and evaluate it on the effectiveness of the model to determine OOD samples (Liang, Li, and Srikant 2017).

Bayesian Neural Networks (BNNs) (Neal 2012) and MC-dropout (Gal and Ghahramani 2016) are theoretically motivated Bayesian methods, and have seen many applications in modeling predictive uncertainty. However, BNNs are: difficult to train, do not scale well to high-dimensional data, and do not perform well under dataset-shift (Lakshminarayanan, Pritzel, and Blundell 2017; Anonymous 2020; Sinha, Ebrahimi, and Darrell 2019). In addition, the choice of priors over the model weights is a crucial factor in their effectiveness. (Hafner et al. 2018; Lakshminarayanan, Pritzel, and Blundell 2017). MC-dropout is a fast and easy to train alternative to BNNs, and can be interpreted as an ensemble model followed by model averaging. However, recent works highlight its limitations in deep learning for uncertainty prediction (Sinha, Ebrahimi, and Darrell 2019; Sener and Savarese 2017), generalization, and predictive accuracy (Kendall and Gal 2017; Lakshminarayanan, Pritzel, and Blundell 2017).

Our work is motivated to provide better generalization and provide reliable uncertainty estimates, which we obtain from inferring multiple plausible hypotheses that are *sufficiently* diverse from each other. This is even more important in cases of high dimensional inputs, like images, because the data distribution is inherently multimodal. Ensemble learning is a natural candidate for learning multiple hypotheses from data. We address the problem of introducing diversity among the different ensemble components (Melville and Mooney 2004) and at the same time ensuring that the predictions balance data likelihood and diversity. To achieve this, we propose an adversarial diversity inducing objective with an information bottleneck (IB) constraint (Tishby, Pereira, and Bialek 2000; Alemi et al. 2016) to enforce forgetting the input as much as possible, while being predictive about the output. IB (Tishby, Pereira, and Bialek 2000) formalizes this in terms of minimizing the mutual information (MI) between the bottleneck representation layer with the input, while maximizing its MI with the correct output, which has been shown to improve generalization in neural networks (Alemi et al. 2016; Achille and Soatto 2018; Goyal et al. 2019).

Recent methods in ensemble learning (Lakshminarayanan, Pritzel, and Blundell 2017; Anonymous 2020) illustrate the drawbacks of applying classical ensembling

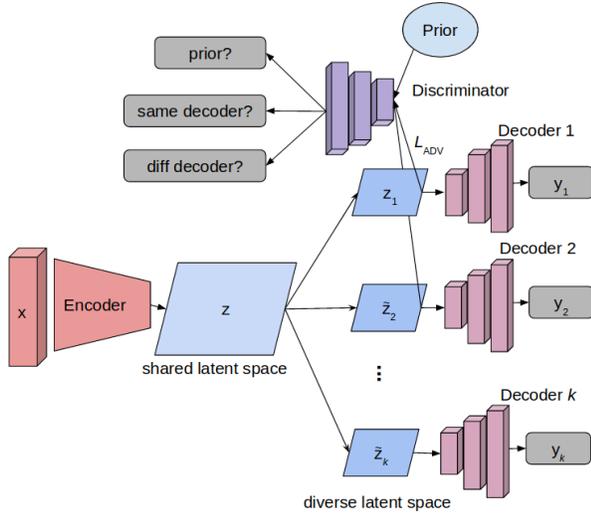


Figure 1: The basic structure of our proposed diverse ensembles approach. The input X is mapped to a shared latent variable Z through a deterministic encoder. The shared Z is mapped to K different stochastic variables \tilde{Z}_i which finally map to the K different outputs $Y_i, \forall i \in [1, \dots, K]$

techniques like bootstrapping (Breiman 1996) to deep neural nets. A recent paper (Anonymous 2020) analyzes the empirical success of ensembling using random initializations compared to Bayesian uncertainty estimation techniques such as BNNs (Neal 2012), and MC-dropout (Gal and Ghahramani 2016) and arrives at the conclusion that random ensembles successfully identify different modes in the data but they do not fit accurately to any mode while Bayesian methods fit accurately but to just one mode in the data. This motivates the need for an ensembling approach that both identifies different modes and fits accurately to each mode, thereby achieving high accuracy, high generalization, and precise uncertainty estimates.

We propose a principled scheme of ensemble learning, by jointly maximizing data likelihood, constraining information flow through a bottleneck to ensure the ensembles capture only *relevant* statistics of the input, and maximizing a diversity inducing objective to ensure that the multiple plausible hypotheses learned are diverse. Instead of K different neural nets, we have K different stochastic decoder heads, as shown in Fig. 1. We explicitly maximize diversity among the ensembles by an adversarial loss. Our ensemble learning scheme has several advantages as compared to randomly initialized ensembles and Bayesian NNs since the joint encoder helps us in learning shared ‘basic’ representations that will be useful for all the decoders. We are able to explicitly control the flow of information from the encoder to each of the decoders during training. Most importantly, we can explicitly enforce diversity among the decoder heads and do not have to rely on random initialization or a prior on the weights to yield diverse output. We show that this diversity enforcing objective helps capture the multiple modes in the

underlying data distribution.

In summary, we claim the following contributions:

1. We introduce diversity among the ensemble members through a novel adversarial loss that encourages samples from different stochastic latent variables to be separated and samples from the same stochastic latent variable to be close to each other.
2. We generalize the VIB (Alemi et al. 2016) formulation to multiple stochastic latent variables and balance diversity with high likelihood by enforcing an information bottleneck between the stochastic latent variables, \tilde{Z}_i , and the input X .

Through extensive experimentation, we demonstrate better generalization to dataset shift, better performance when training with few labels as compared to state-of-the-art baselines, and show better uncertainty estimation on OOD detection. Finally we demonstrate that we achieve consistently better performance with respect to baselines when varying the number of decoders (K) in the ensemble.

Preliminaries

Mutual Information, Information Bottleneck

Mutual Information (MI) is a measure of dependence between two random variables. The MI between random variables X and Y is defined as the KL divergence between the joint distribution and the product of the marginals:

$$\mathcal{I}(X, Y) = KL(\mathbb{P}_{XY} || \mathbb{P}_X \mathbb{P}_Y) \quad (1)$$

By the definition of KL divergence between two probability distributions \mathbb{P} and \mathbb{Q} , $KL(\mathbb{P} || \mathbb{Q}) = \mathbb{E}_{\mathbb{P}}[\log d\mathbb{P}/d\mathbb{Q}]$, we have:

$$\mathcal{I}(X, Y) = \int \int p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dx dy \quad (2)$$

In the variational information bottleneck (VIB) literature (Alemi et al. 2016), given input-output pairs X and Y , the goal is to learn a compressed latent space Z to maximize the mutual information between Z and Y , while minimizing the mutual information between Z and X to learn a representations Z that sufficiently forgets the spurious correlations that may exist in the input X , while still being predictive of Y . More formally:

$$\max_{\theta} \mathcal{I}(Z, Y; \theta) \quad \text{s.t.} \quad \mathcal{I}(X, Z; \theta) \leq I_c$$

Here I_c is some information constraint. This constrained optimization can be solved through Lagrange multipliers.

Notation

In this paper, we consider a network with a shared encoder $f(\cdot)$ and multiple stochastic task-specific decoders $g_i(\cdot)$ where both $f(\cdot)$ and $g_i(\cdot)$ are parameterized using neural networks. Each encoder encodes an image, X , to a shared latent space, Z , which is then used by each task specific decoder to obtain a prediction Y_i , where Y_i is the i -th prediction from the model. Fig 1 describes the architecture visually.

DIBS : Diverse Information Bottleneck in Ensembles

We propose a method for ensemble learning (Hansen and Salamon 1990) by promoting diversity among pairwise latent ensemble variables and by enforcing an information bottleneck (Alemi et al. 2016) between each latent \tilde{Z}_i and the input X . Formally, we consider a set of K decoders $\{\theta_1, \dots, \theta_K\} \sim p(\Theta)$ sampled from some given initial distribution $p(\Theta)$. Given input data X , we want to learn a shared encoding Z , and K decoders that map the latent state Z to \tilde{Z}_i 's and \tilde{Z}_i 's to the K output predictions $\{Y_1, \dots, Y_K\}$.

We posit that for effective learning through ensembles, there must be some diversity among the members of the ensemble, since each ensemble member is by assumption a weak learner, and individual performance is not as important as collective performance (Melville and Mooney 2004). However, promoting diversity randomly among the members is likely to result in uninformative/irrelevant aspects of data being captured by them. Hence, in addition to task-specific standard likelihood maximization, we introduce the need for a diversity enforcing constraint, and a bottleneck constraint. To accomplish the latter, we build upon the Variational Information Bottleneck (VIB) formulation (Alemi et al. 2016) by constraining the information flow from input X to the outputs Y_i , we introduce the information bottleneck term $-\mathcal{I}(\tilde{Z}_i, X; \theta)$. For diversity maximization between ensembles, we design an anti-clustering and diversity-inducing generative adversarial loss, described in the next section.

Adversarial Model for Diversity Maximization

We adopt an adversarial learning approach based on the intuition of diversity maximization among the K models. Our method is inspired by Adversarial Autoencoders (Makhzani et al. 2015), which proposes a natural scheme for combining adversarial training with variational inference. Here, our aim is to maximize separation in distribution between ensemble pairs $q(\tilde{z}_i|x), q(\tilde{z}_j|x)$, such that samples (\hat{z}_1, \hat{z}_2) are indistinguishable to a discriminator if $\hat{z}_1 \sim q(\tilde{z}_i|x), \hat{z}_2 \sim q(\tilde{z}_j|x)$ with $i = j$ and they are distinguishable if $i \neq j$. To this end, we frame the adversarial loss, such that the K generators $q(\tilde{z}_i|x) \forall i \in [1, \dots, K]$ trick the discriminator into thinking that samples from $q(\tilde{z}_i|x)$ and $q(\tilde{z}_j|x)$ are samples from different distributions.

We start with $r(\tilde{z})$, a prior distribution on \tilde{z} . In our case, this is normal $\mathcal{N}(0, \mathbb{I})$, but more complex priors are also supported, in the form of implicit models. We want to make each encoder $q(\tilde{z}_i|x)$ to be close in distribution to this prior, but sufficiently far from other encoders, so that overlap is minimized. Unlike typical GANs (Goodfellow et al. 2014), the discriminator of our diversity inducing loss takes in a pair of samples (\hat{z}_1, \hat{z}_2) instead of just one sample. Hence, we have the following possibilities for the different sources of a pair of latents: 1) $\hat{z}_1 \sim r(\tilde{z})$ and $\hat{z}_2 \sim r(\tilde{z})$, 2) $\hat{z}_1 \sim q(\tilde{z}_i|x)$ and $\hat{z}_2 \sim r(\tilde{z})$, 3) $\hat{z}_1 \sim q(\tilde{z}_i|x)$ and $\hat{z}_2 \sim q(\tilde{z}_i|x)$, and 4) $\hat{z}_1 \sim q(\tilde{z}_i|x)$ and $\hat{z}_2 \sim q(\tilde{z}_j|x)$, with $i \neq j$.

Let $D(\cdot)$ denote the discriminator, which is a feed-forward neural network that takes in a pair (\hat{z}_1, \hat{z}_2) as input and outputs a 0 (fake) or a 1 (real). There are K generators

corresponding to each $q(\tilde{z}_j|z)$, and the deterministic encoder $z = f(x)$. We denote the parameters of all these generators, as well as the deterministic encoder as G , to simplify notation. These generators are trained by minimizing the following loss over G :

$$L_G = \mathbb{E}_{\hat{z}_1 \sim q(\tilde{z}_i|x), \hat{z}_2 \sim q(\tilde{z}_j|x)} [\log D(\hat{z}_1, \hat{z}_2)] \\ + \mathbb{E}_{\hat{z}_1 \sim r(\tilde{z}), \hat{z}_2 \sim q(\tilde{z}_i|x)} [\log(1 - D(\hat{z}_1, \hat{z}_2))] \\ + \mathbb{E}_{\hat{z}_1 \sim q(\tilde{z}_i|x), \hat{z}_2 \sim q(\tilde{z}_i|x)} [\log(1 - D(\hat{z}_1, \hat{z}_2))] \quad (3)$$

Given a fixed discriminator D , the first term encourages pairs of different encoder heads to be distinguishable. The second term encourages each encoder to overlap with the prior. The third term encourages samples from the same encoder to be indistinguishable.

On the other hand, given a fixed generator G , the discriminator is trained by maximizing the following objective function with respect to D :

$$L_D = \mathbb{E}_{\hat{z}_1 \sim r(\tilde{z}), \hat{z}_2 \sim r(\tilde{z})} [\log D(\hat{z}_1, \hat{z}_2)] \\ + \mathbb{E}_{\hat{z}_1 \sim q(\tilde{z}_i|x), \hat{z}_2 \sim q(\tilde{z}_j|x)} [\log D(\hat{z}_1, \hat{z}_2)] \\ + \mathbb{E}_{\hat{z}_1 \sim r(\tilde{z}), \hat{z}_2 \sim q(\tilde{z}_i|x)} [\log(1 - D(\hat{z}_1, \hat{z}_2))] \quad (4)$$

The first term encourages the discriminator to not distinguish between samples from the prior. The second term aims to maximize overlap between different encoders, as an adversarial objective to what the generator is aiming to do in Eqn 3. The third term minimizes overlap between the prior and each encoder.

It is important to note that the generators do not explicitly appear in the loss function because they are implicitly represented through the samples $\hat{z}_1 \sim q(\tilde{z}_i|x), \hat{z}_2 \sim q(\tilde{z}_j|x)$. In each SGD step we backpropagate only through the generator corresponding to the respective (\hat{z}_1, \hat{z}_2) sample. We also note that we consider the pairs (\hat{z}_1, \hat{z}_2) to be unordered in the losses above, because we provide both orderings to the discriminator, to ensure symmetry.

Overall Optimization

The previous sub-section described the diversity inducing adversarial loss. In addition to this, we have the likelihood, and information bottleneck loss terms, denoted together by $\mathbb{L}(\theta)$ below. Here, $\theta = (\theta_D, \theta_G, \Theta)$ denotes the parameters of the discriminator, the generators, and the decoders.

$$\mathbb{L}(\theta) = \sum_{i=1}^m \alpha_i \mathcal{I}(\tilde{Z}_i, Y_i; \theta) - \sum_{i=1}^m \beta_i \mathcal{I}(\tilde{Z}_i, X; \theta)$$

For notational convenience, we omit θ in subsequent discussions. The first term can be lower bounded, as in (Alemi et al. 2016):

$$\mathcal{I}(\tilde{Z}_i, Y_i) \geq \int p(y_i, \tilde{z}_i) \log \frac{q(y_i|\tilde{z}_i)}{p(y_i)} dy_i d\tilde{z}_i \quad (5) \\ = \int p(x) p(y_i|x) p(\tilde{z}_i|x) \log q(y_i|\tilde{z}_i) dx dy_i d\tilde{z}_i + H(Y)$$

The inequality here is a result of $KL(p(y_i|\tilde{z}_i) || q(y_i|\tilde{z}_i)) \geq 0$, where $q(y_i|\tilde{z}_i)$ is a variational approximation to the true

distribution $p(y_i|\tilde{z}_i)$ and denotes our i^{th} decoder. Since the entropy of output labels $H(Y)$ is independent of θ , it can be ignored in the subsequent discussions. Formally, the second term can be formulated as

$$\mathcal{I}(\tilde{Z}_i, X) \leq \int p(\tilde{z}_i|x)p(x) \log \frac{p(\tilde{z}_i|x)}{\psi(\tilde{z}_i)} d\tilde{z}_i dx \quad (6)$$

The inequality here also results from the non-negativity of the KL divergence. The marginal $p(\tilde{z}_j)$ has been approximated by a variational approximation $\psi(\tilde{z}_j)$. Following the approach in VIB (Alemi et al. 2016), to approximate $p(x, y_i)$ in practice we can use the empirical data-distribution $p(x, y_i) = \frac{1}{N} \sum_{n=1}^N \delta_{x^n}(x) \delta_{y_i^n}(y_i)$. We also note that $z^n = f(x^n)$ is the shared encoder latents, where n denotes the n^{th} datapoint among a total of N datapoints. Now, using the re-parameterization trick, we write $\tilde{z}_i = g_i(z, \epsilon)$, where ϵ is a zero mean unit variance Gaussian noise, such that $p(\tilde{z}_i|z) = \mathcal{N}(\tilde{z}_i|g_i^\mu(z), g_i^\Sigma(z))$. We finally obtain the following lower-bound approximation of the the loss function. The detailed derivation is in the Appendix.

$$\begin{aligned} \mathbb{L} \approx & \frac{1}{N} \sum_{n=1}^N \left[\mathbb{E}_{\epsilon \sim p(\epsilon)} \left[\sum_{i=1}^m \alpha_i \log q(y_i^n | g_i(f(x^n), \epsilon)) \right. \right. \\ & \left. \left. - \sum_{i=1}^m \beta_i \text{KL} \left(p(\tilde{z}_i|x^n) \parallel \psi(\tilde{z}_i) \right) \right] \right] \quad (7) \end{aligned}$$

In our experiments we set $\psi(\tilde{z}_j) = \mathcal{N}(\tilde{z}_j|0, \mathbb{I})$. To make predictions in classification tasks, we output the modal class of the set of class predictions by each ensemble member.

Similar to GANs (Goodfellow et al. 2014), the model is optimized using alternating optimization where we alternate among objectives $\max_{\theta} \mathbb{L}(\theta)$, $\min_{\theta_G} L_G$, and $\max_{\theta_D} L_D$. It is important to note that we do not explicitly optimize the KL-divergence term above, but implicitly do it during the process of adversarial learning using \mathcal{L}_{adv} . In Section 3.1, the case $\hat{z}_1 \sim q(\tilde{z}_1|x)$ and $\hat{z}_2 \sim r(\tilde{z})$ corresponds to minimizing this KL-divergence term. This is done similarly to (Makhzani et al. 2015).

Predictive Uncertainty Estimation

Our proposed method is able to meaningfully capture both epistemic and aleatoric uncertainty. Aleatoric uncertainty is typically modeled as the variance of the output distribution, which can be obtained by outputting a distribution, say a normal $p(y|x, \theta) \sim \mathcal{N}(\mu_\theta(x), \sigma_\theta(x))$ (Hafner et al. 2018).

Epistemic uncertainty in traditional Bayesian Neural Networks (BNNs) is captured by defining a prior (often an uninformative prior) over model weights $p(\theta)$, updating it based on the data likelihood $p(\mathcal{D}|\theta)$, where \mathcal{D} is the dataset and θ is the parameters, in order to learn a posterior over model weights $p(\theta|\mathcal{D})$. In practice, for DNNs since the true posterior cannot be computed exactly, we need to resort to samples from some approximate posterior distribution $q(\theta|\mathcal{D}) \approx p(\theta|\mathcal{D})$ (Gustafsson, Danelljan, and Schön 2019).

In our approach, for epistemic uncertainty, we note that although ensembles do not consider priors over weights

(unlike BNNs), they correspond to learning multiple models $\{\theta_k\}_{k=1}^K$ which can be considered to be samples from some approximate posterior $q(\theta|\mathcal{D})$ (Gustafsson, Danelljan, and Schön 2019), where \mathcal{D} is the training dataset. We note that $p(y|x, \mathcal{D}) = \int_{\tilde{z}, \theta} p(y|\tilde{z}, \theta) p(\tilde{z}|x, \theta) p(\theta|\mathcal{D})$ and a typical Bayesian NN would directly approximate $p(\theta|\mathcal{D})$, which would require a prior over weights $p(\theta)$, whose selection is problematic. DIBS avoids this issue by turning sampling into optimization of a set of θ_k such that $p(\tilde{z}|x, \theta_k)$ are diverse, but still predictive of $p(\tilde{z}|x)$, without explicitly approximating $p(\theta|\mathcal{D})$. As a result there is also no notion of a *true* posterior over weights $p(\theta|\mathcal{D})$ (unlike in BNNs).

For aleatoric noise, we note that we have stochastic latent variables $\tilde{z}_k \sim p(\tilde{z}_k|x)$ and obtain respective outputs $p(y_k|\tilde{z}_k)$. By sampling multiple times (say M times) from $p(\tilde{z}_k|x)$, we obtain an empirical distribution $\{y_{k,i}\}_{i=1}^M$. The empirical variance of the output distributions of all the ensembles $\{y_{k,i}\}_{i=1, k=1}^{M, K}$ gives us a measure of aleatoric uncertainty

The posterior predictive distribution gives a measure of the combined predictive uncertainty (epistemic+aleatoric), which for our approach can be calculated as follows:

$$\hat{p}(y^*|x^*) = \frac{1}{MK} \sum_{k=1}^K \sum_{i=1}^M p(y_{k,i}^*|x, \theta_k)$$

Since we enforce diversity among the K ensemble members through the adversarial loss described in Section , we expect to obtain more reliable aleatoric uncertainty estimate and hence better predictive uncertainty overall. We perform experimental evaluation of predictive uncertainty estimation through OOD detection experiments in the next section.

Experiments

In the section, we show how our method is able to achieve:

- 1. Better accuracy:** How do the proposed approach and baselines perform on the task of image classification in the face of limited data?
- 2. Better generalization:** How well do the models generalize when the evaluation distribution is different from the training distribution?
- 3. Better uncertainty estimation:** Are we are able to obtain better uncertainty estimates compared to the baselines as evidenced by OOD detection?

We compare our approach to four external baselines: **ABE** (Kim et al. 2018), **NCP** (Hafner et al. 2018), **MC-Dropout** (Gal and Ghahramani 2016), and the state-of-the-art deep ensemble scheme of (Lakshminarayanan, Pritzel, and Blundell 2017), that considers ensembles to be randomly initialized and trained neural networks. We henceforth call this method **Random**. For NCP, we impose the NCP priors on the input and output for each NN architecture that we evaluate. For images, the input prior amounts to an additive Gaussian noise on each pixel. ABE (Kim et al. 2018) considers a diversity inducing loss based on pairwise squared difference among the ensemble outputs, and is a recently proposed strong baseline. MC-Dropout (Gal and Ghahramani 2016) is a Bayesian method that samples

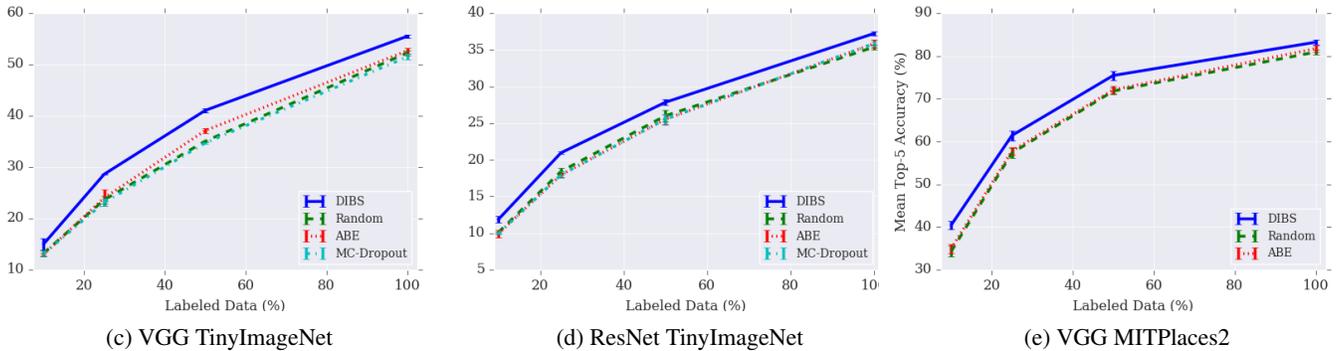
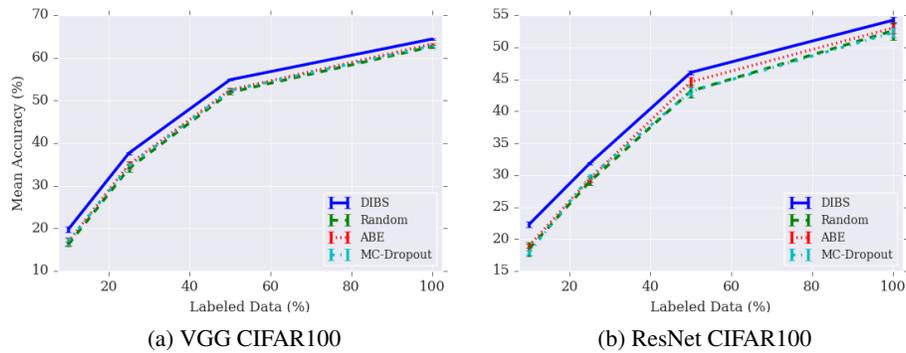


Figure 2: Performance of baselines, Random (Lakshminarayanan, Pritzel, and Blundell 2017), MC-Dropout (Gal and Ghahramani 2016) and ABE (Kim et al. 2018) against our proposed approach DIBS on four datasets with two backbone architectures. All results show % accuracy on the test dataset. The y-axis label on (a) propagates to figures (b), (c), (d), and (e). We show results for different % of labels of the dataset used during training. It is evident that when less data is used for training, DIBS relatively performs much better than the baseline. The specific details of the architectural variants are in the Appendix.

dropout masks repeatedly to produce different predictions from the model. We evaluate the performance of these baselines along our model DIBS in five benchmark datasets: MNIST, CIFAR-10, CIFAR-100 (Krizhevsky, Hinton et al. 2009), TinyImageNet (Russakovsky et al. 2015), and MIT Places 2 (Zhou et al. 2017). Performing experiments on the Places 2 dataset confirms that our method is able to scale well to large scale settings as it is a scene recognition dataset with over 1.8 M images and 365 unique classes.

Experimental Setup

We run experiments with two standard vision architectures as the backbone, namely VGG19 (Simonyan and Zisserman 2014) and ResNet18 (He et al. 2016). For optimization, we use Stochastic Gradient Descent (SGD) (Bottou 2010) with a learning rate of 0.05 and momentum of 0.9 (Sutskever et al. 2013). We decay the learning rate by a factor of 10 every 30 epochs of training.

Performance

Experiments on CIFAR10, CIFAR100, TinyImageNet, and MIT Places 2 show that DIBS outperforms all baselines on the task of image classification. We evaluate the performance of the proposed approach DIBS and all the baselines on four datasets: MNIST, CIFAR100, Tiny-

ImageNet, and MIT Places 2 and evaluating the classification accuracy. To demonstrate good performance “at-scale,” we consider three base architectures: a simple 4-layer CNN, VGG Networks (Simonyan and Zisserman 2014), and ResNets (He et al. 2016). Specifics of these architectures are mentioned in the Appendix. Fig. 2 show results in terms of % accuracy on the CIFAR-10, CIFAR-100, TinyImageNet, and MIT Places 2 datasets when there are respectively 100%, 50%, 25%, and 10% of the labeled dataset used during training. For the MIT Places 2 dataset, we considered the top-5 classification accuracy in order to be consistent with the evaluation metric in the original challenge (Zhou et al. 2017). For all the other datasets, we consider the top-1 classification accuracy. We randomly sampled examples from the entire training dataset to create these smaller training sets.

It is interesting to note that when less data is used during training, DIBS performs relatively much better than the baselines indicating better generalization. As evident from Fig. 2, DIBS consistently performs better than all the baseline schemes with all the base architectures. The results on the Places 2 dataset demonstrates that our approach can effectively scale to a significantly larger dataset. From Fig. 2, we can also see that the relevant magnitude of performance improvement of DIBS over baselines increase as the dataset

size increases (MIT Places 2, TinyImageNet, CIFAR100). This suggests the efficacy of our approach in the image classification task.

Generalization and Transfer Experiments

In this section we consider experiments of generalization to changes in the data distribution (without finetuning) and transfer under dataset shift to a different test distribution (with finetuning). For all the experiments here we use a simple 4 layer feedforward CNN with maxpool layer and ReLU non-linearity after every layer. Details are mentioned in the Appendix. We use this instead of a VGGNet or ResNet due to the small scale of the datasets involved in the experiments.

DIBS effectively generalizes to dataset change under translation, and rotation of digits. In this section, we consider the problem of generalization, through image translation, and image rotation experiments on MNIST. The generalization experiments on MNIST are described below:

1. **Translate (Trans):** Training on normal MNIST images, and testing by randomly translating the images by 0-5 pixels, 0-8 pixels, and 0-10 pixels respectively.
2. **Rotate:** Training on normal MNIST digits, and testing by randomly rotating the images by 0-30 degrees and 0-45 degrees respectively.
3. **Interpolation-Extrapolation-Translate (IETrans):** We train on images translated randomly in the range [-5,5] pixels and test on images translated randomly in the range [-10,10] pixels.
4. **Interpolation-Extrapolation-Rotate (IERotate):** We train on images rotated randomly in the range [-22,22] degrees and test on images rotated randomly in the range [-45,45] degrees.
5. **Color:** We train on Normal MNIST images and test on colored MNIST images (Kim et al. 2019) by randomly changing foreground color to red, green, or blue.

The interpolation-extrapolation experiments help us understand the generalization of the models on the data distribution that it was trained on (interpolation), as well as on a data distribution objectively different from training (extrapolation). Hence, we consider the testing distributions to be a superset of the training distribution in these two experiments. Table 1 summarizes the results of these experiments. We observe that DIBS achieves over 2% higher accuracy compared to the baseline of Random (Lakshminarayanan, Pritzel, and Blundell 2017) in the experiment of generalizing under translation shift and over 1% higher accuracy in the rotational shift experiment.

Here we show that DIBS effectively transfers when trained on a source dataset and finetuned and evaluated on a target dataset. We train our model DIBS on one dataset which we call the source and finetune it on another dataset which we call the target, and finally evaluate it on the test set of the target dataset. We consider the following experiment:

1. **Source:** MNIST; **Target:** SVHN (Netzer et al. 2011)

For finetuning in the target dataset, we fix the encoder(s) of DIBS and the baselines Random (Lakshminarayanan, Pritzel, and Blundell 2017) and ABE (Kim et al. 2018), and

Generalization experiments				
Exp	Test details	Random	ABE	DIBS
Trans	[-5,5] p	76.38±0.70	77.53±0.68	79.46±0.50
	[-10,10] p	37.84±1.05	39.06±0.98	41.38±1.00
Rotate	[-30,30] d	95.79±0.18	94.96±0.18	96.36±0.11
	[-45,45] d	87.09±0.31	87.82±0.16	88.90±0.17
IETrans	[-10,10] p	93.04±0.22	92.76±0.30	94.21±0.41
IERotate	[-45,45] d	97.79±0.10	97.85±0.15	98.35±0.37
Color	R,G,B	97.63±0.49	97.61±0.38	98.20±0.63
Transfer experiments				
Source	Target	Random	ABE	DIBS
MNIST	SVHN	43.11±2.10	47.05±1.62	50.09±0.97

Table 1: Generalization and Transfer experiments on MNIST. Details of the experiments are mentioned in Section . Results show that for all the experiments, DIBS outperforms the baselines. For the transfer experiment, we train for 50 epochs on MNIST, free the encoder(s) and fine-tune in the training dataset of SVHN for 20 epochs and report the % accuracy on the test dataset of SVHN. Here p denotes pixel and d degrees. The backbone is a simple 4-layer CNN described in the Appendix.

update the parameters of the decoders for a few fixed iterations. We train on MNIST for 50 epochs and fine-tune on the training dataset of SVHN for 20 epochs before evaluating on the test dataset of SVHN. Details of the exact procedure are in the Appendix. Results in Table 1 show that DIBS achieves higher accuracy under transfer to the target environment in this experiment as compared to the baselines.

Uncertainty Estimation Through OOD Examples

DIBS achieves accurate predictive uncertainty estimates necessary for reliable OOD detection. We follow the scheme of (Hafner et al. 2018; Liang, Li, and Srikant 2017) for evaluating on Out-of-Distribution (OOD) examples. We train our model on CIFAR (Krizhevsky, Hinton et al. 2009) and then at test time, consider images sampled from the dataset to be in-distribution and images sampled from a different dataset, say Tiny-Imagenet to be OOD. In our experiments, we use four OOD datasets, namely Imagenet-cropped (Russakovsky et al. 2015) (randomly cropping image patches of size 32x32), Imagenet-resized (Russakovsky et al. 2015) (downsampling images to size 32x32), synthetic uniform-noise, and Gaussian-Noise. The details are same as (Liang, Li, and Srikant 2017).

To elaborate on the specifics of OOD detection at test time, given input image x , we calculate the softmax score of the input with respect to each ensembles $\mathcal{S}_i(x)$ and compare the score to a threshold δ_i . For DIBS, each $\mathcal{S}_i(x)$ corresponds to a particular decoder head. The aggregated ensemble prediction is given by the mode of the individual predictions. The details of this procedure are mentioned in the Ap-

Trained on CIFAR-10/ Trained on CIFAR-100				
		ImageNet Crop	ImageNet Resize	Gaussian
ROC	Rand	84.61/68.87	80.53/73.14	77.09/68.24
	NCP	83.13/66.13	78.11/68.93	75.31/72.43
	ABE	85.24/69.94	79.09/73.45	82.16/76.45
	DIBS	87.57/70.06	82.47/75.71	86.56/80.61
		ImageNet Crop	ImageNet Resize	Gaussian
PR	Rand	82.08/67.72	77.76/73.69	88.52/84.66
	NCP	82.60/66.15	75.34/70.16	85.29/84.06
	ABE	84.33/68.11	77.71/74.86	89.10/89.21
	DIBS	88.67/68.28	81.10/74.98	93.36/91.00

Table 2: Comparison of performance on OOD data classification. DIBS and the baselines, Random, ABE, and NCP are trained on a particular dataset (CIFAR-10/CIFAR-100) and are then tasked with prediction of images to be in-distribution or out-of-distribution (OOD). For OOD examples, we use four datasets as described in Section . For evaluation, we use the metrics AUROC (Area Under the Receiver Operating Characteristic Curve) and AUPR (Area under the Precision-Recall curve). Results show that DIBS significantly outperforms all the baselines on both the metrics.

pendix. Table 2 compares the performance of DIBS against baselines, Random (Lakshminarayanan, Pritzel, and Blundell 2017) and NCP (Hafner et al. 2018). It is evident that DIBS consistently outperforms both the baselines on both the AUROC and AUPR metrics.

Related Work

Ensembles have been used in fields ranging from computer vision (Huang et al. 2017; Lee et al. 2016) to reinforcement learning and imitation learning for planning and control (Chua et al. 2018; Li, Song, and Ermon 2017). Traditionally, ensembles have been proposed to tackle the problem of effective generalization (Hansen and Salamon 1990), and algorithms like random forests (Breiman 2001), and broad-approaches like boosting (Freund, Schapire, and Abe 1999), and bagging (Breiman 1996) are common ensemble learning techniques. In ensemble learning, multiple models are trained to solve the same problem. Each individual learner model is a simple model, or a ‘weak learner’ while the aggregate model is a ‘strong learner.’

Diversity among ensemble learners, important for generalization (Lee et al. 2016; Zhou, Wu, and Tang 2002; Hansen and Salamon 1990), has traditionally been ensured by training each weak learner on a separate held-out portion of the training data (bagging) (Breiman 1996), adding random noise to the output predictions, randomly initializing model weights (Lakshminarayanan, Pritzel, and Blundell 2017; Anonymous 2020), having stochastic model weights (Neal 2012), or by manipulating the features and attributes (Kim et al. 2018; Lee et al. 2016) of the model. As demonstrated by (Lakshminarayanan, Pritzel, and Blundell 2017), bagging

is not a good diversity inducing mechanism, when the underlying base learner has multiple local optima, as is the case with neural net architectures, which are the focus of this paper. BNNs (Neal 2012) provide reasonable epistemic uncertainty estimates but do not necessarily capture the inherent aleatoric uncertainty, and so are not capable of successfully inducing diversity in the output predictions for effectively modeling multi-modal data (Anonymous 2020; Kendall and Gal 2017).

(Lakshminarayanan, Pritzel, and Blundell 2017) proposes a mechanism of randomly initializing the weights of a neural net architecture, and hence obtaining an ensemble of neural network models, treated as an uniformly weighted mixture of Gaussians. This approach outperforms bagging and BNNs in terms of both predictive accuracy and uncertainty estimation, however, as pointed out in (Anonymous 2020) the number of ensembles needed to accurately identify different modes and model each mode sufficiently requires a large number of models, and is computationally expensive.

Motivated by this, instead of adopting a random initialization approach, we proposed a principled scheme of diversity maximization among latent ensemble variables, so that different modes in the data distribution are identified, and constrained the diversity of the latent variables through an information bottleneck. We adopted the approach of having a shared encoder and K -headed stochastic decoder, with each head of the decoder representing one model of the ensemble and utilize an adversarial loss to promote meaningful diversity. (Kim et al. 2018) proposes a similar architecture, but for enforcing diversity among the decoders, the authors explicitly maximize the Euclidean distance between every pair of feature embeddings (for each datapoint), and is not guaranteed to separate the multiple data modes “in-distribution” in the embedding space.

Another important component of our architecture is an information bottleneck constraint, that constrains the flow of information from the input layer X to each of the K stochastic latent decoder variables \tilde{Z}_i ’s, so that the predictions don’t become arbitrarily diverse due to the diversity inducing loss. This relates to the work in (Alemi et al. 2016), which we extend to K latent variables instead of just one.

Conclusion

In this paper we addressed the issue of enforcing diversity in a learned ensemble through a novel adversarial loss, while ensuring high likelihood of the predictions, through the notion of variational information bottleneck. We demonstrate through extensive experimentation that the proposed approach outperforms state-of-the-art baseline ensemble and Bayesian learning methods on four benchmark datasets in terms of accuracy under sparse training data, uncertainty estimation for OOD detection, and generalization to a test distribution significantly different from the training data distribution. Our technique is generic and applicable to any latent variable model.

References

- Achille, A.; and Soatto, S. 2018. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence* 40(12): 2897–2905.
- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Anonymous. 2020. Deep Ensembles: A Loss Landscape Perspective. In *Submitted to International Conference on Learning Representations*. URL <https://openreview.net/forum?id=r1xZAKrFPr>. Under review.
- Bottou, L. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, 177–186. Springer.
- Breiman, L. 1996. Bagging predictors. *Machine learning* 24(2): 123–140.
- Breiman, L. 2001. Random forests. *Machine learning* 45(1): 5–32.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, 4754–4765.
- Freund, Y.; Schapire, R.; and Abe, N. 1999. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence* 14(771-780): 1612.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, 1050–1059.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Goyal, A.; Islam, R.; Strouse, D.; Ahmed, Z.; Botvinick, M.; Larochelle, H.; Levine, S.; and Bengio, Y. 2019. Infobot: Transfer and exploration via the information bottleneck. *arXiv preprint arXiv:1901.10902*.
- Gustafsson, F. K.; Danelljan, M.; and Schön, T. B. 2019. Evaluating Scalable Bayesian Deep Learning Methods for Robust Computer Vision. *arXiv preprint arXiv:1906.01620*.
- Hafner, D.; Tran, D.; Irpan, A.; Lillicrap, T.; and Davidson, J. 2018. Reliable uncertainty estimates in deep neural networks using noise contrastive priors. *arXiv preprint arXiv:1807.09289*.
- Hansen, L. K.; and Salamon, P. 1990. Neural network ensembles. *IEEE Transactions on Pattern Analysis & Machine Intelligence* (10): 993–1001.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, G.; Li, Y.; Pleiss, G.; Liu, Z.; Hopcroft, J. E.; and Weinberger, K. Q. 2017. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, 5574–5584.
- Kim, B.; Kim, H.; Kim, K.; Kim, S.; and Kim, J. 2019. Learning Not to Learn: Training Deep Neural Networks with Biased Data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9012–9020.
- Kim, W.; Goyal, B.; Chawla, K.; Lee, J.; and Kwon, K. 2018. Attention-based ensemble for deep metric learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 736–751.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Technical report, Citeseer.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 6402–6413.
- Lee, S.; Prakash, S. P. S.; Cogswell, M.; Ranjan, V.; Crandall, D.; and Batra, D. 2016. Stochastic multiple choice learning for training diverse deep ensembles. In *Advances in Neural Information Processing Systems*, 2119–2127.
- Li, Y.; Song, J.; and Ermon, S. 2017. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, 3812–3822.
- Liang, S.; Li, Y.; and Srikant, R. 2017. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*.
- Makhzani, A.; Shlens, J.; Jaitly, N.; Goodfellow, I.; and Frey, B. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Melville, P.; and Mooney, R. J. 2004. Diverse ensembles for active learning. In *Proceedings of the twenty-first international conference on Machine learning*, 74.
- Neal, R. M. 2012. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115(3): 211–252.
- Sener, O.; and Savarese, S. 2017. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sinha, S.; Ebrahimi, S.; and Darrell, T. 2019. Variational Adversarial Active Learning. *arXiv preprint arXiv:1904.00370*.
- Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, 1139–1147.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Yu, F.; Koltun, V.; and Funkhouser, T. 2017. Dilated residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 472–480.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146*.

Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* 40(6): 1452–1464.

Zhou, Z.-H.; Wu, J.; and Tang, W. 2002. Ensembling neural networks: many could be better than all. *Artificial intelligence* 137(1-2): 239–263.