THE CONDITIONAL ENTROPY BOTTLENECK

Anonymous authors
Paper under double-blind review

ABSTRACT

We present a new family of objective functions, which we term the *Conditional Entropy Bottleneck* (CEB). We demonstrate the application of CEB to classification tasks. In our experiments, CEB gives: well-calibrated predictions; essentially perfect detection of challenging out-of-distribution examples and powerful whitebox adversarial examples; and robustness to the same. Finally, we report that CEB fails to learn a dataset with fixed random labels, providing a possible resolution to the problem of generalization observed in Zhang et al. (2016).

1 INTRODUCTION

The field of Machine Learning has suffered from the following well-known problems in recent years[1]:

- **Vulnerability to adversarial examples.** Essentially all machine-learned systems are currently believed by default to be highly vulnerable to adversarial examples. Many defenses have been proposed, but very few have demonstrated robustness against a powerful, general-purpose adversary. Lacking a clear theoretical framework for adversarial attacks, most proposed defenses are ad-hoc and fail in the presence of a concerted attacker (Carlini & Wagner, 2017a; Athalye et al., 2018).
- **Poor out-of-distribution detection.** Classifiers do a poor job of signaling that they have received data that is substantially different from the data they were trained on. Ideally, a trained classifier would give less confident predictions for data that was far from the training distribution (as well as for adversarial examples). Barring that, there would be a clear, principled statistic that could be extracted from the model to tell whether the model *should* have made a low-confidence prediction. Many different approaches to providing such a statistic have been proposed (Guo et al., 2017; Lakshminarayanan et al., 2016; Hendrycks & Gimpel, 2016; Liang et al., 2017; Lee et al., 2017; DeVries & Taylor, 2018), but most seem to do poorly on what humans intuitively view as obviously different data.
- **Miscalibrated predictions.** Related to the issues above, classifiers tend to be very overconfident in their predictions (Guo et al., 2017). This may be a symptom, rather than a cause, but miscalibration does not give practitioners confidence in their models.
- **Overfitting to the training data.** Zhang et al. (2016) demonstrated that classifiers can memorize fixed random labelings of training data, which means that it is possible (or even easy) to learn a classifier with perfect *inability* to generalize. This critical observation makes it clear that a fundamental test of generalization is that the model should *fail* to learn when given fixed random labels.

This paper does not set out to solve any of these problems. Instead, our sole interest is the learning of optimal representations. In pursuit of that goal, we attempt to be as general as possible, considering only how to define optimal representations, what objective function might be capable of learning them, and what requirements such an objective function places on the form of the model.

Given an optimal objective function, however, it is natural to explore the problems listed above, to see if such an objective function can ameliorate some of the core issues in the field of machine learning. We make those explorations in this paper, and find that our objective function, the *Conditional Entropy Bottleneck* (CEB) appears to impact all of the issues listed above.

[1]These problems existed before recent years, but not all of them were known. In particular, adversarial examples were unknown prior to 2013, and the severity of the overfitting problem was not known until 2016.

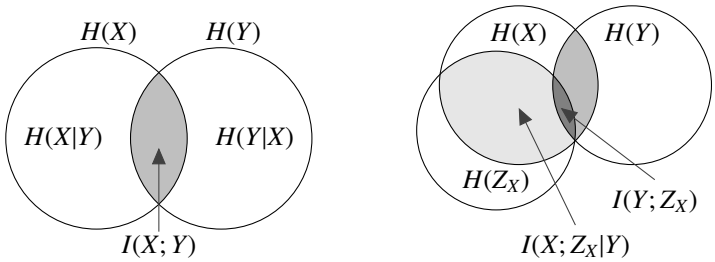

Figure 1: **(Left)**: Information Venn diagram showing the joint distribution over X, Y. **(Right)**: The joint distribution $Z_X \leftarrow X \leftrightarrow Y$. Z_X is carefully positioned to indicate its conditional independence from Y given X.

2 Optimal Representations

In the following discussion, we will take some liberty and conflate the scalar values of the information theoretic functionals, the entropy and the mutual information, with the underlying set-theoretic elements that those functionals measure. See Reza (1994) for a discussion of the set-theoretic analogies to information theory. When we talk about "covering" information quantities, we have these set-theoretic elements in mind.

Consider a joint distribution, $p(x, y)$, represented by the graphical model:

$$X \leftrightarrow Y$$

This joint distribution is our data, and may take any form. We don't presume to know how the data factors. It may factor as $p(x, y) = p(x)p(y|x)$, $p(x, y) = p(y)p(x|y)$, or even $p(x, y) = p(x)p(y)$.

The first two factorings are depicted in Figure 1 in a standard information diagram showing the various entropies and the mutual information. We can ask: given this generic setting, what is the optimal representation? It seems there are only two options worth considering: $H(X, Y)$ and $I(X; Y)$.

The field of lossless compression is concerned with representations that perfectly maintain $H(X, Y)$, as are the closely related studies of *Kolmogorov Complexity* (Kolmogorov, 1965) and *Minimum Description Length* (MDL) (Grünwald, 2007), all three of which are concerned with perfect *reconstruction* of inputs or messages.

In contrast, we think that the field of machine learning is primarily concerned with optimal *generalization* to unseen data. All of these fields recognize the importance of minimality, but the requirements of perfect reconstruction necessarily result in the retention of much more information in the model than may be needed for prediction or stochastic generation tasks. For most such machine learning tasks, this leaves only the representation corresponding to mutual information between X and Y.

The mutual information is defined in a variety of ways; we will use two (Cover & Thomas, 2006):

$$I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X) \tag{1}$$

$I(X; Y)$ is the unique, minimal representation *defined* by a dataset. The selection of a particular dataset entails this representation – all information not covered by $I(X; Y)$ is superfluous. For example, consider a labeled dataset, where X is high-dimensional and information-rich, and Y is a single integer. All of the information in X that is not needed to correctly predict the single value $Y = y$ is useless for the task defined by the dataset, and may actually be harmful to the performance of a machine learning system if retained in the learned representation, as we will show below.

In Appendix A, we define the *Minimal Necessary Information* (MNI) criterion for determining the optimality of a representation. In the two-variable case (X, Y), this corresponds exactly to the mutual information, $I(X; Y)$. Thus, we take the position that an optimal representation Z for two observed variables, X, Y, must "cover" $I(X; Y)$ in some sense. We show how to do this next.

3 The Conditional Entropy Bottleneck

Given the Minimum Necessary Information criterion, we would like to find a way to learn $I(X, Y)$ for an arbitrary dataset $P(X, Y)$. We can view much of machine learning as taking a data distribution outside of our control ($P(X, Y)$ in this case), and adding a new random variable, Z_X, that is under our control.[2] Doing this in the natural way gives the Markov chain $Z_X \leftarrow X \leftrightarrow Y$, shown as an information diagram in Figure 1 (Right). The placement of $H(Z_X)$ in that diagram carefully maintains the conditional independence between Y and Z_X given X, but is otherwise fully general. Some of the entropy of Z_X is unassociated with any other variable; some is only associated with X, and some is associated with X and Y together. Figure 1 (Right), then, shows diagrammatically the state of the learned representation early in training. We would like instead for $H(Z_X)$ to exactly cover $I(X; Y)$, as in the gray area in Figure 1. We can find an objective function that achieves this goal optimally by considering the equalities that must hold between the various entropies and mutual informations at the moment that $H(Z_X)$ covers $I(X; Y)$.

Broadly, then, our goal is to take the information diagram in Figure 1 (Right), and transform it into Figure 1 (Left), where the information that Z_X represents coincides exactly with the information shared between X and Y.

First, note the following equalities in that scenario:

$$I(X; Y; Z_X) = I(X; Y) = I(X; Z_X) = I(Y; Z_X)$$

These equalities can be read directly from the information diagram in Figure 1 (Left). Because of our Markov chain, we can say definitively that maximizing $I(X; Z_X)$ can never lead to a minimal representation. While doing so will cover $I(X; Y)$, there will be nothing in the objective function to prevent it from also covering all of $H(X)$. However, maximizing $I(Y; Z_X)$ is consistent with learning the Minimal Information and necessary for learning the Necessary Information. It is not *sufficient*, though, as $H(Z)$ can still cover $H(X)$ completely when $I(Y; Z_X)$ is maximal.

Now, consider the following equalities, also visible in Figure 1 (Left):

$$I(X; Y|Z_X) = I(X; Z_X|Y) = I(Y; Z_X|X) = 0$$

With our Markov chain, we have the following well-known equality (Cover & Thomas, 2006):

$$I(X; Z_X|Y) = I(X; Z_X) - I(Y; Z_X) \tag{2}$$

This is guaranteed to be non-negative, even for continuous random variables, as both terms are mutual informations, which are non-negative, and the Markov chain guarantees that $I(Y; Z_X)$ is no larger than $I(X; Z_X)$, by the data processing inequality. From an optimization perspective, this is ideal – we have a term that we can minimize, and we can directly know how far we are from the optimal value of 0 (measured in nats, so it is interpretable), when we are done (when it's close enough to 0 that we are satisfied), and when our model is insufficient for the task (i.e., when this term *isn't* close enough to 0).

The above derivation gives the fully general *Conditional Entropy Bottleneck* objective:

$$\text{CEB} \equiv \min I(X; Z_X|Y) - I(Y; Z_X) \tag{3}$$

It is straightforward to turn this into a variational objective function, similar to the *Variational Information Bottleneck* (VIB) (Alemi et al., 2017). Taking the two terms in turn, we have:[3]

$$\begin{aligned}
\min I(X; Z_X|Y) &= I(X; Z_X) - I(Y; Z_X) && (4)\\
&= H(Z_X) - H(Z_X|X) - H(Z_X) + H(Z_X|Y) && (5)\\
&= -H(Z_X|X) + H(Z_X|Y) && (6)\\
&= \langle \log e(z_X|x) \rangle - \langle \log p(z_X|y) \rangle && (7)\\
&\leq \langle \log e(z_X|x) \rangle - \langle \log p(z_X|y) \rangle + \text{KL}[p(z_X|y) \| b(z_X|y)] && (8)\\
&= \langle \log e(z_X|x) \rangle - \langle \log b(z_X|y) \rangle && (9)
\end{aligned}$$

$e(z_X|x)$ is our *encoder*. It is not a variational approximation, even though it has learned parameters. $b(z_X|y)$ is the variational approximation to what we think of as the *backward encoder*.

[2] This is the perspective taken in the Information Bottleneck (Tishby et al., 2000).

[3] We write expectations $\langle \log e(z_X|x) \rangle$. They are always with respect to the joint distribution; e.g., $p(x, y, z_X)$.

The second term:

$$\max I(Y;Z_X) = H(Y) - H(Y|Z_X) \tag{10}$$
$$\Rightarrow \max\ -H(Y|Z_X) \tag{11}$$
$$= \langle \log p(y|z_X) \rangle \tag{12}$$
$$\geq \langle \log p(y|z_X) \rangle - \mathrm{KL}[p(y|z_X) \| c(y|z_X)] \tag{13}$$
$$= \langle \log p(y|z_X) \rangle - \langle \log p(y|z_X) \rangle + \langle \log c(y|z_X) \rangle \tag{14}$$
$$= \langle \log c(y|z_X) \rangle \tag{15}$$

$c(y|z_x)$ is the variational approximation to the *classifier* (although that name is arbitrary, given that Y may not be labels).

The variational bounds derived above give us a fully tractable objective function that works on large-scale problems, *Variational Conditional Entropy Bottleneck* (VCEB):

$$\mathrm{CEB} \equiv \min I(X;Z_X|Y) - I(Y;Z_X) \Rightarrow \min \langle \log e(z_X|x) \rangle - \langle \log b(z_X|y) \rangle - \langle \log c(y|z_X) \rangle \equiv \mathrm{VCEB} \tag{16}$$

The distributions with letters other than p are assumed to have learned parameters, which we otherwise omit in the notation. In other words, all three of $e(\cdot)$, $b(\cdot)$, and $c(\cdot)$ have learned parameters, just as in the encoder and decoder of a normal VAE (Kingma & Welling, 2014), or the encoder, classifier, and marginal in a VIB model (Alemi et al., 2017). Indeed, it is possible to switch between a CEB model and a VIB model simply by replacing the marginal $m(z_X)$ with the backward encoder, $b(z_X|y)$ and updating the loss function provided to the optimizer.

We will name the $I(X;Z_X|Y)$ term the *Residual Information* – this is the excess information in our representation beyond the information shared between X and Y:

$$Re_{X/Y} \equiv \langle \log e(z_X|x) \rangle - \langle \log b(z_X|y) \rangle \geq -H(Z_X|X) + H(Z_X|Y) = I(X;Z_X|Y) \tag{17}$$

There are a number of natural variations on this objective. We describe a few of them in Appendix C.

4 Variational Information Bottleneck

The Information Bottleneck (IB) (Tishby et al., 2000) attempts to learn a representation of X and Y subject to an information constraint:

$$\max I(Z;Y) \text{ subject to } I(Z;X) \leq R \tag{18}$$

where R is a constant *bottleneck*. This can be rewritten as an unconstrained Lagrangian optimization:

$$\max I(Z;Y) - \beta I(Z;X) \tag{19}$$

where β controls the size of the bottleneck.

A varational version of this objective is presented in Alemi et al. (2017). That objective is the *Variational Information Bottleneck* (VIB):

$$VIB \equiv \min \beta(\langle \log e(z_X|x) \rangle - \langle \log m(z_X) \rangle) - \langle \log c(y|z_X) \rangle \tag{20}$$

This is very similar to CEB, but instead of the backward encoder, VIB has a *marginal posterior*, $m(z_X)$, which is a variational approximation to $e(z_X) = \int dx\, p(x) e(z_X|x)$. Additionally, it has a hyperparameter, β. We show in Appendix B that the optimal value for $\beta = \frac{1}{2}$ when attempting to adhere to the MNI criterion.

Following Alemi et al. (2018), we define the *Rate* (R):

$$R \equiv \langle \log e(z_X|x) \rangle - \langle \log m(z_X) \rangle \geq I(X;Z_X) \tag{21}$$

5 Training

Because of the properties of $Re_{X/Y}$, we can consider different training algorithms. In particular, we can avoid needing to look at validation set performance in order to decide when to lower the learning

Table 1: Accuracy and rates (R) for each model. **Bold** indicates the best score in that column. Determ doesn't have a rate, since it doesn't have an explicit encoder distribution. The final rate for the other four models is reported, as well as the peak rate achieved during training. The true mutual information for Fashion MNIST is $I(X;Y) = 2.3$ nats, so achieving $R = 2.3$ is optimal according to MNI.

Model	Accuracy	Train R final (peak)
Determ	**92.7**	n/a
$\text{VIB}_{0.01}$	**93.0**	2.6 (11.6)
$\text{VIB}_{0.1}$	**92.7**	**2.3** (3.2)
$\text{VIB}_{0.5}$	90.0	**2.3** (2.4)
CEB	**92.9**	**2.3** (**2.3**)

rate. The closer we can get $Re_{X/Y}$ to 0 on the training set, the better we will generalize to data drawn from the same distribution. Consequently, one simple approach to training is to set a high initial learning rate (possibly with reverse annealing of the learning rate (Goyal et al., 2017)), and then lower the learning rate after any epoch of training that doesn't result in a new lowest mean residual information on the *training* data. This is equivalent to the *dev-decay* training algorithm of Wilson et al. (2017), but does not require the use of a validation set. Additionally, since the training set is typically much larger than a validation set would be, the average loss over the epoch is much more stable, so the learning rate is less likely to be lowered spuriously.

Of course, it would be inefficient to compute the mean residuals for the entire training set at the end of each epoch of training, but the same general argument holds for maintaining a running mean during the epoch. When training is making progress, this estimate is an upper bound on the training set residuals at the end of the epoch. However, it is possible for training to stall or even diverge during an epoch, and this estimate may not pick up immediately on those cases, delaying lowering the learning rate until more than an epoch of stalled or divergent training has happened. In our experience, this is not practically a problem – CEB models are very stable throughout training and can tolerate the occasional epoch where the learning rate is still too high. Indeed, we find it easiest to prevent any lowering of the learning rate for a large number of epochs (e.g., 40), and thereafter following this algorithm. We do not attempt to prove that this algorithm is optimal.

Remark. $Re_{X/Y}$ directly measures how far from optimal our learned representation is. If our optimization procedure is sufficently effective, the residual indicates that we could improve performance by increasing the capacity of our architecture or considering ways in which our model may be misspecified. Thus, CEB directly informs us of the possibility to improve our model.

6 Classification Experiments

Our primary experiments are focused on comparing the performance of otherwise identical models when we merely change the objective function. Consequently, we aren't interested in demonstrating state-of-the-art results in all things. Instead, we are interested in relative differences in performance that can be directly attributed to the difference in objective.

With that in mind, we present results for classification of Fashion MNIST (Xiao et al., 2017) for five different models. The five models are: a deterministic model (*Determ*); three VIB models, with $\beta \in \{\frac{1}{2}, 10^{-1}, 10^{-2}\}$ ($\text{VIB}_{0.5}$, $\text{VIB}_{0.1}$, $\text{VIB}_{0.01}$); and a CEB model.

All five models share the same core architecture mapping X to Y: a 7×2 Wide Resnet (Zagoruyko & Komodakis, 2016) for the encoder, with a final layer of $D = 4$ dimensions for the latent representation, followed by a two layer MLP classifier using ELU (Clevert et al., 2015) activations with a final categorical distribution over the 10 classes. The stochastic models parameterize the mean and variance of a $D = 4$ fully covariate multivariate Normal distribution with the output of the encoder. Samples from that distribution are passed into the classifier MLP. Apart from that difference, the stochastic models don't differ from Determ during evaluation. None of the five models uses any form of regularization (e.g., L_1, L_2, DropOut (Srivastava et al., 2014), BatchNorm (Ioffe & Szegedy, 2015)).

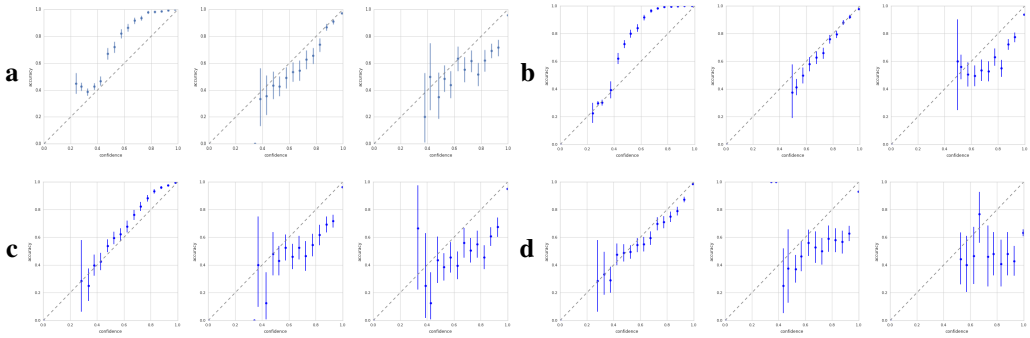

Figure 2: Calibration plots with 90% confidence intervals for four of the models after 2,000 steps, 20,000 steps, and 40,000 steps (left, center, and right of each trio, respectively): **a** is CEB, **b** is $\text{VIB}_{0.5}$, **c** is $\text{VIB}_{0.1}$, **d** is Determ. *Perfect calibration* corresponds to the dashed diagonal lines. *Underconfidence* occurs when the points are above the diagonal. *Overconfidence* occurs when the points are below the diagonal.

The VIB models have an additional learned marginal, $m(z_X)$, which is a mixture of 240 $D = 4$ fully covariate multivariate Normal distributions. The CEB model instead has the backward encoder, $b(z_X|y)$ which is a $D = 4$ fully covariate multivariate Normal distribution parameterized by a 1 layer MLP mapping the label, $Y = y$, to the mean and variance. In order to simplify comparisons, for CEB we additionally train a marginal posterior identical in form to that used by the VIB models. However, for CEB, $m(z_X)$ is trained using a separate optimizer so that it doesn't impact training of the CEB objective in any way. Having $m(z_X)$ for both CEB and VIB allows us to compare the rate, R, of each model except Determ.

Since Fashion MNIST doesn't have a prespecified validation set, it offers an opportunity to test training algorithms that only look at training results, rather than relying on cross validation. To that end, the five models presented here are the first models with these hyperparameters that we trained on Fashion MNIST.[4] The learning rate for the CEB model was lowered according to the training algorithm described in Section 5. The other four models followed the same algorithm, but instead of tracking $Re_{X/Y}$, they simply tracked their training loss. All five models were required to retain the initial learning rate of 0.001 for 40 epochs before they could begin lowering the learning rate. At no point during training did any of the models exhibit non-monotonic test accuracy, so we do not believe that this approach harmed any performance – all five models converged essentially smoothly to their final, reported performance.

In the case of a simple classification problem with a uniform distribution over classes in the training set, we can directly compute $I(X;Y)$ as $\log C$, where C is the number of classes.[5] See Table 1 for a comparison of the rates between the four variational models, as well as their accuracies. All but $\text{VIB}_{0.5}$ achieve the same accuracy. All four stochastic models get close to the ideal rate of 2.3 nats, but they get there by different paths. For the VIB models, the lower β is, the higher the rate goes early in training, before converging down to (close to) 2.3 nats. CEB never goes above 2.3 nats.

7 Calibration

In Figure 2, we show calibration plots at various points during training for the four models. Calibration curves help analyze whether models are underconfident or overconfident. Each point in the plots corresponds to a 5% confidence range. Accuracy is measured in each bin. A *well-calibrated* model is correct half of the time it gives a confidence of 50% for its prediction.

All of the networks move from under- to overconfidence during training. However, CEB and $\text{VIB}_{0.5}$ are only barely overconfident, while reducing β to 0.1 is sufficent to make it nearly as overconfident

[4] Development focused on MNIST (LeCun et al., 1998) and 2 dimensional latent vectors for ease of visualization. See Figure 6 for an example.

[5] We are relying on the mild assumption that X (the high-dimensional data) has higher entropy than Y (the labels). $I(X;Y) \leq \min(H(X), H(Y))$.

Table 2: Results for out-of-distribution detection (*OoD*). *Thrsh.* is the threshold score used: H is the entropy of the classifier; R and $Re_{X/\hat{Y}}$ are defined in Section 8. Arrows denote whether higher or lower scores are better. **Bold** indicates the best score in that column for a particular OoD dataset.

OoD	Method	Thrsh.	FPR @ 95% TPR ↓	AUROC ↑	AUPR In ↑
U(0,1)	Determ	H	35.8	93.5	97.1
	$\text{VIB}_{0.01}$	H	41.1	92.5	96.0
		R	**0.0**	**100.0**	**100.0**
	$\text{VIB}_{0.1}$	H	43.5	94.5	96.2
		R	**0.0**	**100.0**	**100.0**
	$\text{VIB}_{0.5}$	H	73.2	87.0	90.5
		R	80.6	57.1	51.4
	CEB	H	63.4	92.8	95.1
		R	**0.0**	**100.0**	**100.0**
		$Re_{X/\hat{Y}}$	**0.0**	**100.0**	**100.0**
MNIST	Determ	H	59.0	88.4	90.0
	$\text{VIB}_{0.01}$	H	42.3	91.6	95.9
		R	**0.0**	**100.0**	**100.0**
	$\text{VIB}_{0.1}$	H	60.3	84.7	89.7
		R	**0.5**	86.8	**99.8**
	$\text{VIB}_{0.5}$	H	70.2	79.6	86.8
		R	12.3	66.7	91.1
	CEB	H	70.6	77.8	73.0
		R	**0.1**	94.4	**99.9**
		$Re_{X/\hat{Y}}$	**0.2**	92.0	**99.9**
Vertical Flip	Determ	H	66.8	88.6	90.2
	$\text{VIB}_{0.01}$	H	57.6	82.6	80.3
		R	**0.0**	**100.0**	**100.0**
	$\text{VIB}_{0.1}$	H	65.3	84.5	85.2
		R	**0.0**	**99.2**	**100.0**
	$\text{VIB}_{0.5}$	H	79.7	79.8	81.4
		R	17.3	52.7	91.3
	CEB	H	68.0	84.9	85.5
		R	**0.0**	90.7	**100.0**
		$Re_{X/\hat{Y}}$	**0.0**	92.6	**100.0**

as the deterministic model. This overconfidence is one of the issues that is correlated with exceeding the MNI during training (see Table 1). See Appendix B for a more in-depth explanation for how this can occur.

8 Out-of-Distribution Detection

We test the ability of the five models to detect three different out-of-distribution (OoD) detection settings. $U(0,1)$ is uniform noise in the image domain. MNIST uses the MNIST test set. Vertical Flip is the most challenging, using vertically flipped Fashion MNIST test images, as originally proposed in Alemi et al. (2018).

We use three different metrics for thresholding. The first two, H and R, were proposed in Alemi et al. (2018). H is the classifier entropy. R is the rate, defined in Section 4. The third metric is specific to CEB: $Re_{X/\hat{Y}}$. This is the predicted residual information – since we don't have access to the true value of Y at test time, we use $\hat{y} \sim c(y|z_X)$ to calculate $H(Z_X|\hat{Y})$. This is no longer a valid bound on $Re_{X/Y}$, as $\hat{y}$ may not be from the true distribution $p(x, y, z_X)$. However, the better the classifier, the closer the estimate should be.

These three threshold scores are used with the standard suite of proper scoring rules: *False Positive Rate at 95% True Positive Rate* (FPR 95% TPR), *Area Under the ROC Curve* (AUROC), and *Area Under the Precision-Recall Curve* (AUPR). See Lee et al. (2018) for definitions.

The core result is that $\text{VIB}_{0.5}$ performs much less well at the OoD tasks than the other two VIB models and CEB. We believe that this is another result of $\text{VIB}_{0.5}$ learning the right amount of information, but

not learning all of the *right* information, thereby demonstrating that it is not a valid MNI objective, as explored in Appendix B. On the other hand, the other two VIB objectives seem to perform extremely well, which is the benefit they get from capturing a bit more information about the training set. We will see below that there is a price for that information, however.

9 Adversarial Example Robustness and Detection

Table 3: Results for adversarial example detection (*Attack*). All attacks are targeting the "trousers" class in Fashion MNIST. *CW* is Carlini & Wagner (2017b). *CW,* $(C = 1)$ is CW with an additional confidence penalty set to 1. *CW,* $(C = 1)$ *Det.* is a custom CW attack targeting CEB's detection mechanism, $Re_{X/\hat{Y}}$. L_0, L_1, L_2, L_∞ report the corresponding norm (mean ±1 std.) of successful adversarial perturbations. Higher norms on CW indicate that the attack had a harder time finding adversarial perturbations, since it starts by looking for the smallest possible perturbation. The remaining columns are as in Table 2. Arrows denote whether higher or lower scores are better. **Bold** indicates the best score in that column for a particular adversarial attack. See Section 6 for details of the models and Section 9 for details of the attacks.

Attack	Model	Attack Success ↓	L_0 ↑	L_1 ↑	L_2 ↑	L_∞ ↑	Thrsh.	FPR @ 95% TPR ↓	AUROC ↑	AUPR In ↑
CW	Determ	100.0%	377.1 ±100.3	16.2 ±10.2	1.4 ±1.7	0.2 ±0.1	H	15.4	90.7	86.0
	$\text{VIB}_{0.01}$	55.2%	389.6 ±100.9	17.1 ±10.3	1.5 ±1.8	0.2 ±0.1	H R	11.2 **0.0**	59.9 **100.0**	90.0 **100.0**
	$\text{VIB}_{0.1}$	68.8%	392.1 ±101.6	29.2 ±18.1	5.1 ±7.5	0.4 ±0.2	H R	16.5 **0.0**	77.4 **100.0**	80.0 **100.0**
	$\text{VIB}_{0.5}$	**35.8%**	**432.0** ±99.6	**40.1** ±32.1	**9.4** ±14.4	**0.5** ±0.3	H R	64.2 **0.0**	62.5 98.7	55.3 **100.0**
	CEB	**35.8%**	416.4 ±97.7	33.6 ±30.3	7.4 ±15.0	0.3 ±0.2	H R $Re_{X/\hat{Y}}$	62.2 **0.0** **0.0**	65.2 **99.7** **99.5**	57.1 **100.0** **100.0**
CW $(C = 1)$	Determ	100.0%	378.7 ±100.3	16.6 ±10.4	1.4 ±1.9	0.2 ±0.1	H	17.9	90.9	85.7
	$\text{VIB}_{0.01}$	96.7%	381.3 ±101.5	17.4 ±10.5	1.6 ±1.9	0.2 ±0.1	H R	19.6 **0.0**	72.1 **100.0**	89.6 **100.0**
	$\text{VIB}_{0.1}$	97.3%	382.8 ±100.4	28.2 ±17.2	4.8 ±7.4	**0.4** ±0.2	H R	28.7 **0.0**	86.0 **100.0**	79.1 **100.0**
	$\text{VIB}_{0.5}$	50.4%	**422.0** ±101.3	**36.4** ±28.6	**7.8** ±12.3	**0.4** ±0.2	H R	86.5 **0.1**	59.8 96.2	54.1 **100.0**
	CEB	**48.0%**	**417.6** ±95.5	33.3 ±29.8	**7.3** ±15.4	**0.4** ±0.2	H R $Re_{X/\hat{Y}}$	77.4 **0.0** **0.0**	63.5 **99.3** 98.7	56.4 **100.0** **100.0**
CW $(C = 1)$ Det.	CEB	**25.1%**	416.4 ±92.2	**84.1** ±44.0	**34.4** ±22.8	**0.9** ±0.1	H R $Re_{X/\hat{Y}}$	95.1 66.5 72.9	56.4 69.3 69.9	45.0 88.5 87.6

Adversarial examples were first noted in Szegedy et al. (2013). The first practical attack, *Fast Gradient Method* (FGM) was introduced shortly after (Goodfellow et al., 2015). Since then, many new attacks have been proposed. Most relevant to us is the Carlini-Wagner (CW) attack (Carlini & Wagner, 2017b), which was the first practical attack to directly use a blackbox optimizer to find minimal perturbations.[6] Many defenses have also been proposed, but almost all of them are broken (Carlini & Wagner, 2017a; Athalye et al., 2018). This work may be seen as a natural continuation of the adversarial analysis of Alemi et al. (2017), which showed that VIB naturally had robustness to whitebox adversaries, including CW. In that work, the authors did not train any VIB models with a learned $m(z_X)$, which results in much weaker models, as shown in Alemi et al. (2018). We believe this is the first work exploring learning a marginal and using that marginal in an adversarial setting.

[6]Szegedy et al. (2013) initially used L-BFGS to find the adversaries. Carlini & Wagner (2017b) showed that it was possible to use Adam (Kingma & Ba, 2015), which is much faster.

We consider CW in the whitebox setting to be the current gold standard attack, even though it is much more expensive than FGM or the various iterative attacks like DeepFool (Moosavi-Dezfooli et al., 2016) or iterative variants of FGM (Kurakin et al., 2016). Running an optimizer directly on the model to find the perturbation that can fool the model tells us much more about the robustness of the model than approaches that only take a small, fixed number of gradient steps. Additionally searching over the space of perturbation magnitudes makes the attack very hard to defend against, and consequently the current best option for testing robustness.

Here, we explore three variants of the CW L_2 targeted attack. The implementation the first two CW attacks are from Papernot et al. (2018). CW and CW ($C = 1$) are the baseline CW attack, and CW with a confidence adjustment of 1. Note that in order for these attacks to succeed at all on CEB, we had to increase the default CW learning rate to 5×10^{-1}. Without that increase, CW found almost no adversaries in our early experiments. All other parameters are left at their defaults for CW, apart from setting the clip ranges to $[0, 1]$. The final attack, CW ($C = 1$) Det. is a modified version of CW ($C = 1$) that additionally incorporates a detection tensor into the loss. For CEB, we had it target minimizing $Re_{X/\hat{Y}}$ in order to try to break the network's ability to detect the attack.

All of the attacks are targeting the *trouser* class of Fashion MNIST, as that is the most distinctive class. Targeting a less distinctive class, such as one of the shirt classes, would confuse the difficulty of classifying the different shirts and the difficulty of the adversary. We run each of the first three attacks on the *entire* Fashion MNIST test set (all 10,000 images). For the stochastic networks, we permit 32 encoder samples and take the mean classification result (these samples are also used for gradient generation in the attacks to be fair to the attacker). CW is expensive, but we are able to run these on a single GPU in about 30 minutes. However, CW ($C = 1$) Det. ends up being about 200 times more expensive – we were only able to run 1000 images, permitting only 8 samples from the encoder, and it took $2\frac{1}{2}$ hours. Consequently, we only run CW ($C = 1$) Det. on the CEB model, and the results are less significant.

Our metric for robustness is the following: we count the number of adversarial examples that change a correct prediction to an incorrect prediction of the target class, and divide by the number of correct predictions the model makes on the non-adversarial inputs. We additionally measure the size of the resulting perturbations. For CW, a larger perturbation generally indicates that the attack had to work harder to find an adversarial example, making this a secondary indication of robustness. Finally, we measure detection using the same thresholding techniques from Table 2.

The results of these experiments are in Table 3. We show all 20,000 images for four of the models in Figure 7. The most striking pattern in the models is how well $\text{VIB}_{0.01}$ and $\text{VIB}_{0.1}$ do at detection, while $\text{VIB}_{0.5}$ is dramatically more robust. We think that this is the most compelling indication of the importance of not overshooting $I(X; Y)$ – even minor amounts of overshooting appear to destroy the robustness of the model. On the other hand, $\text{VIB}_{0.5}$ has a hard time with detection, which indicates that, while it has learned a highly compressed representation, it has not learned the optimal set of bits. Thus, as we discuss in Appendix B, VIB trades off between learning the *necessary* information, which allows it to detect attacks perfectly, and learning the *minimum* information, which allows it to be robust to attacks.

In the end, however, CEB permits both – it maintains the necessary information for detecting powerful whitebox attacks, but also retains the minimum information, providing robustness. This is again visible in the CW ($C = 1$) Det. attack, which directly targets CEB's detection mechanism. Even though it no longer does well detecting the attack, the model becomes *more* robust to the attack, as indicated both by the much lower attack success rate and the much larger perturbation magnitudes.

10 Generalization Experiments

We replicate the basic experiment from Zhang et al. (2016): we use the images from Fashion MNIST, but replace the training labels with fixed random labels (i.e., the same random value every epoch of training). We use that dataset to train multiple deterministic models, CEB models, and a range of VIB models. We find that the CEB model *never* learns (even after 100 epochs of training), the deterministic model always learns (after about 40 epochs of training it begins to memorize the random labels), and the VIB models only learn with $\beta \leq 0.001$.

The fact that CEB and VIB with β near $\frac{1}{2}$ manage to resist memorizing random labels is our final empirical demonstration that MNI is a powerful criterion for objective functions.

11 Conclusion

We have presented the basic form of the Conditional Entropy Bottleneck (CEB), motivated by the Minimum Necessary Information (MNI) criterion for optimal representations. We have shown through careful experimentation that simply by switching to CEB, you can expect substantial improvements in OoD detection, adversarial example detection and robustness, calibration, and generalization. Additionally, we have shown that it is possible to get all of these advantages without using any additional form of regularization, and without any new hyperparameters. We have argued empirically that objective hyperparameters can lead to hard-to-predict suboptimal behavior, such as memorizing random labels, or reducing robustness to adversarial examples. In Appendix C and in future work, we will show how to generalize CEB beyond the simple case of two observed variables.

It is our perspective that all of the issues explored here – miscalibration, failure at OoD tasks, vulnerability to adversarial examples, and dataset memorization – stem from the same underlying issue, which is retaining too much information in the representation learned, either implicitly or explicitly, for the training inputs. We believe that the MNI criterion and CEB show a path forward for many tasks in machine learning, permitting fast, amortized inference while ameliorating major problems.

Acknowledgments

REDACTED

References

A. A. Alemi, B. Poole, I. Fischer, J. V. Dillon, R. A. Saurous, and K. Murphy. Fixing a Broken ELBO. *ICML2018*, 2018. URL http://arxiv.org/abs/1711.00464.

Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep Variational Information Bottleneck. In *International Conference on Learning Representations*, 2017. URL http://arxiv.org/abs/1612.00410.

Alexander A Alemi, Ian Fischer, and Joshua V Dillon. Uncertainty in the variational information bottleneck. *arXiv preprint arXiv:1807.00906*, 2018.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

William Bialek and Naftali Tishby. Predictive information. *arXiv preprint cond-mat/9902341*, 1999.

Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 3–14. ACM, 2017a.

Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 39–57. IEEE, 2017b.

Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.

Thomas M Cover and Joy A Thomas. *Elements of information theory 2nd edition*. John Wiley & Sons, 2006.

T. DeVries and G. W. Taylor. Learning Confidence for Out-of-Distribution Detection in Neural Networks. *arXiv: 1802.04865*, 2018. URL https://arxiv.org/abs/1802.04865.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *CoRR*, 2015.

Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

Peter D Grünwald. *The Minimum Description Length Principle*. MIT press, 2007.

C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On Calibration of Modern Neural Networks. *arXiv: 1706.04599*, 2017. URL https://arxiv.org/abs/1706.04599.

D. Hendrycks and K. Gimpel. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *arXiv: 1610.02136*, 2016. URL https://arxiv.org/abs/1610.02136.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.

Andrei N Kolmogorov. Three approaches to the quantitative definition ofinformation'. *Problems of information transmission*, 1(1):1–7, 1965.

Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.

B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *arXiv: 1612.01474*, 2016. URL https://arxiv.org/abs/1612.01474.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

K. Lee, H. Lee, K. Lee, and J. Shin. Training Confidence-calibrated Classifiers for Detecting Out-of-Distribution Samples. *arXiv: 1711.09325*, 2017. URL https://arxiv.org/abs/1711.09325.

Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *arXiv preprint arXiv:1807.03888*, 2018.

S. Liang, Y. Li, and R. Srikant. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. *arXiv: 1706.02690*, 2017. URL https://arxiv.org/abs/1706.02690.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2574–2582, 2016.

Nicolas Papernot, Fartash Faghri, Nicholas Carlini, Ian Goodfellow, Reuben Feinman, Alexey Kurakin, Cihang Xie, Yash Sharma, Tom Brown, Aurko Roy, Alexander Matyasko, Vahid Behzadan, Karen Hambardzumyan, Zhishuai Zhang, Yi-Lin Juang, Zhi Li, Ryan Sheatsley, Abhibhav Garg, Jonathan Uesato, Willi Gierke, Yinpeng Dong, David Berthelot, Paul Hendricks, Jonas Rauber, and Rujun Long. Technical report on the cleverhans v2.1.0 adversarial examples library. *arXiv preprint arXiv:1610.00768*, 2018.

Fazlollah M Reza. *An introduction to information theory*. Courier Corporation, 1994.

Claude Elwood Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27:379–423, 1948.

Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv: 1312.6199*, 2013. URL https://arxiv.org/abs/1312.6199.

N. Tishby and N. Zaslavsky. Deep Learning and the Information Bottleneck Principle. *arXiv: 1503.02406*, 2015. URL https://arxiv.org/abs/1503.02406.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *SSW*, pp. 125, 2016.

Ramakrishna Vedantam, Ian Fischer, Jonathan Huang, and Kevin Murphy. Generative models of visually grounded imagination. *International Conference on Learning Representations*, 2018. URL https://arxiv.org/abs/1705.10762.

Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103. ACM, 2008.

A. C. Wilson, R. Roelofs, M. Stern, N. Srebro, and B. Recht. The Marginal Value of Adaptive Gradient Methods in Machine Learning. *arXiv: 1705.08292*, 2017. URL https://arxiv.org/abs/1705.08292.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

S. Zagoruyko and N. Komodakis. Wide Residual Networks. *arXiv: 1605.07146*, 2016. URL https://arxiv.org/abs/1605.07146.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Here we collect a number of results that are not critical to the core of the paper, but may be of interest to particular audiences.

A Minimum Necessary Information

We propose the *Minimum Necessary Information* (MNI) criterion for a learned representation. We think it is necessary to explicitly name this, because it does not seem to be specified in other proposed criteria. As stated in Section 2, Kolmogorov Complexity and MDL are only interested in lossless compression. The Information Bottleneck Principle of Tishby et al. (2000) comes close, but it fails to take a strong stance on how to determine the correct amount of information in the representation, instead leaving that critical decision to the practitioner. Of course, it is well-known (Cover & Thomas, 2006) that the mutual information is a minimal sufficient statistic. Thus, everything that we are saying here should apply to any other minimal sufficient statistic of X, Y. The reason to consider the mutual information to be a special minimal sufficient statistic, and to define MNI separate from minimal sufficient statistics, is because the mutual information is part of the broader information theory which provides a very general set of mechanisms for manipulating this particular minimal sufficient statistic. As we have demonstrated in Section 3 and appendix C, use of information theory makes derivation of optimal MNI objective functions easy.

Thus, Minimum Necessary Information is a much stronger and more restrictive criterion for representations than the others, justifying its separate identification. However, we do not want to overclaim novelty, as the Minimum Necessary Information in the bivariate case exactly maps to the well-known, thoroughly studied concept of mutual information: the mutual information is the only quantity that we know of that satisfies both minimality and necessity, and it is certainly the only information-theoretic metric to do so.

We can define MNI in three rather obvious pieces. First is *Information*: we would like a representation that captures semantically meaningful information. In order to measure how successfully we capture meaningful information, we must first know how to measure information. Thus, the criterion prefers information-theoretic approaches, given the uniqueness of entropy as a measure of information (Shannon, 1948). The process that determines which information is meaningful is, by lack of any reasonable alternative, a task, which is specified by the set of variables in the dataset. I.e., the dataset X, Y defines two potential tasks: predict Y given X, or predict X given Y. This brings us to *Necessity*: the information we capture in our representations must be necessary to solve the task. This aspect of the criterion is the source of any semantic value of the information in the representation. Finally, *Minimality*: this simply refers to the amount of information – given that we learn a representation that can solve the task, we require that the representation we learn retain the smallest amount of information about the task out of the set of all representations that solve the task. This part of the criterion restricts us from incorporating "non-semantic" information into our representation, such as noise or spurious correlation.

B Analysis of CEB and IB

By setting $\beta = \frac{1}{2}$ IB can achieve a representation that exactly matches the quantity of information $I(X; Y)$. It is worth noting that this is a somewhat surprising theoretical result. The analysis performed in Alemi et al. (2018), naively applied to VIB, would suggest that modifying the architecture should entail having to adjust β to maintain the same constraint on $I(X; Z_X)$. This result shows that this is not the case for $\beta = \frac{1}{2}$ – regardless of architecture and modeling choices, if the model is capable of hitting the Minimum Necessary Information target for a dataset, it will do so. The reason the analysis from Alemi et al. (2018) does not apply, however, is perhaps more clear now: the graphical models $Z_X \leftarrow X$ (β VAE) and $Z_X \leftarrow X \leftrightarrow Y$ (VIB,CEB) give different information constraints, and VIB is able to leverage (with proper β) the same conditional independence constraint for the task that CEB uses directly. In contrast, the unsupervised learning graphical model has no such conditional independence constraint, so modeling and architectural choices play a clear role when selecting β.

However, setting $\beta = \frac{1}{2}$ still is not sufficient to achieve the Minimum Necessary Information – it simply will perfectly *target* that quantity, but the particular bits selected by the learned representation

may not be the optimal bits. One way of viewing the VIB objective is the following:

$$CEB \equiv \min \langle \log e(z_X|x) \rangle - \langle \log p(z_X|y) \rangle - \langle \log p(y|z_X) \rangle \tag{22}$$

$$\leq \min \langle \log e(z_X|x) \rangle - \langle \log p(z_X|y) \rangle - \langle \log p(y|z_X) \rangle + \mathrm{KL}\left[p(z_X|y) \| \int dy\, p(y) p(z_X|y) \right] \tag{23}$$

$$= \min \langle \log e(z_X|x) \rangle - \left\langle \log \int dy\, p(y) p(z_X|y) \right\rangle - \langle \log p(y|z_X) \rangle \tag{24}$$

$$= \min \langle \log e(z_X|x) \rangle - \langle \log p(z_X) \rangle - \langle \log p(y|z_X) \rangle \tag{25}$$

$$= \min -H(Z_X|X) + H(Z_X) + H(Y|Z_X) \equiv IB \tag{26}$$

In the fourth line, we are simply identifying $p(z_X)$ as the marginal posterior of $p(z_X|y)$, rather than as the marginal posterior of $p(z_X|x)$ in the normal VIB derivation. Both of these marginals refer to the same random variable Z_X, thus giving the same entropy, $H(Z_X)$. This derivation of IB shows explicitly that we can view it as a variational approximation to CEB, where we approximate $p(z_X|y)$ with $p(z_X) = \int dy\, p(y) p(z_X|y)$, the true marginal. This is a valid variational approximation with the rather unusual property of not having any of its own parameters – to get an objective we can optimize, we would still need to make a second parameterized variational approximation using $\mathrm{KL}[p(z_X) \| m(z_X)]$, as in VIB.

We can ask the question: how tight is the unparameterized variational approximation?

$$\min -H(Z_X|X) + H(Z_X|Y) + H(Y|Z_X) = \min -H(Z_X|X) + H(Z_X) + H(Y|Z_X) - K \tag{27}$$

We can trivially solve for K:

$$K = I(Y; Z_X)$$

This gives us a "corrected" IB objective:

$$IB_{corrected} = \min -H(Z_X|X) + H(Z_X) + H(Y|Z_X) - I(Y; Z_X) \tag{28}$$

$$= \min -H(Z_X|X) + H(Z_X) + 2H(Y|Z_X) - H(Y) \tag{29}$$

$$\Leftrightarrow \min -H(Z_X|X) + H(Z_X) + 2H(Y|Z_X) \tag{30}$$

Note that the choices we make for the correction above are the only choices that lead to valid variational approximations. It is conceivable that non-variational approximations to the mutual information would permit different choices in this analysis. This analysis again shows the factor of $\frac{1}{2}$ (dividing the final result by 2 gives the normal formulation with $\beta = \frac{1}{2}$) we identified above that allows us to get to exactly the correct *amount* of information in the representation, regardless of model.

We now consider whether this choice of double variational approximation should have any additional impact on final performance. By this, we mean that we started with a term $H(Z_X|Y)$, made a variational approximation to it with the true $H(Z_X)$, but corrected for that approximation by doubling the weight of the classifier, $H(Y|Z_X)$. Now, we will make the second variational approximation, where $H(Z_X)$ is replaced with the learned marginal, $m(z_X)$. We would like to know how that compares to the direct variational approximation of $H(Z_X|Y)$ with $b(z_X|y)$. To do this, we can just take the difference between the two variational objectives:

$$\delta \equiv CEB - VIB_{corrected} \tag{31}$$

$$= \langle \log e(z_X|x) \rangle - \langle \log b(z_X|y) \rangle - \langle \log c(y|z_X) \rangle - \langle \log e(z_X|x) \rangle + \langle \log m(z_X) \rangle + 2 \langle \log c(y|z_X) \rangle \tag{32}$$

$$= - \langle \log b(z_X|y) \rangle + \langle \log m(z_X) \rangle + \langle \log c(y|z_X) \rangle \tag{33}$$

If δ is positive, CEB has a looser bound. Otherwise, $VIB_{corrected}$ has a looser bound. If we make the mild assumptions that all three variational approximations are sufficiently expressive to have negligible error at convergence, that we are measuring at convergence, and that Y is discrete, we can rewrite δ using entropies:

$$\delta \approxeq H(Z_X|Y) - H(Z_X) - H(Y|Z_X) \tag{34}$$

$$= -I(Y; Z_X) - H(Y|Z_X) \leq 0 \tag{35}$$

This shows clearly that IB does not satisfy the Minimum Necessary Information criterion, as when it reaches the minimum information, that information does not contain all of the nats that are *necessary*

to solve the task. Equivalently, when IB solves the task, it does so with more than the minimum necessary number of nats.

There is one further question we can ask about this analysis: how many nats is δ at the end of learning a representation with $VIB_{corrected}$? Because we are doing representation learning, at every step t of training Z_X^t is a valid representation of X. This makes the analysis very simple: at each step of training, we throw away $\lambda_t \delta_t$ nats of information when we update our parameters using VIB with $\beta = \frac{1}{2}$ as our objective function, where λ_t is the learning rate at time t, and δ_t is the correction term. Thus, at step T of training, we have lost this much necessary information (or overshot the minimum information by this amount):

$$\Delta_T \equiv \sum_t^T \lambda_t(-I(Y; Z_X^t) - H(Y|Z_X^t)) \tag{36}$$

In principle, Δ_T can be approximated by separately training a $b(z_X|y)$ distribution to convergence at every step of IB training. In practice, the other two variational approximations would also need to be trained to convergence at every update of $e(z_X|x)$ as well, otherwise the measurement error early in training when $c(y|z_X)$ and $m(z_X)$ are not tight will dominate the estimate.

Remark. Recent research has explored whether Deep Learning might work because it is doing performing Information Bottleneck implicitly (Tishby & Zaslavsky, 2015; Shwartz-Ziv & Tishby, 2017). We won't venture an opinion on the question, apart from pointing out that it is clear to us that maximum likelihood on a DNN is evidently not doing $IB_{corrected}$, given its well-known propensity to memorize training data, which is incompatible with MNI objectives.

C Additional CEB Objectives

Here we describe a few of the more obvious variants of the CEB objective.

C.1 Conditional Generation

In the above presentation of CEB, we derived the objective for what may be termed "classification" tasks (although there is nothing in the derivation that restricts the form of either X or Y). However, CEB is fully symmetric, so it is natural to consider the second task defined by our choice of dataset, conditional generation of X given $Y = y$.

In this case, we can augment our graphical model with a new variable, Z_Y, and derive the same CEB objective for that variable:

$$\min I(Y; Z_Y|X) = \min I(Y; Z_Y) - I(X; Z_Y) \tag{37}$$
$$\Rightarrow \min -H(Z_Y|Y) + H(Z_Y|X) \tag{38}$$
$$\tag{39}$$
$$\max I(X; Z_Y) = \max H(X) - H(X|Z_Y) \tag{40}$$
$$\Rightarrow \max -H(X|Z_Y) \tag{41}$$

In the same manner as above, we can derive variational bounds on $H(Z_Y|X)$ and $H(X|Z_Y)$. In particular, we can variationally bound $p(z_Y|x)$ with $e(z_Y|x)$. Additionally, we can bound $p(x|z_Y)$ with a decoder distribution of our choice, $d(x|z_Y)$.

Because the decoder is maximizing a lower bound on the mutual information between Z_Y and X, it can never memorize X. It is directly limited during training to use exactly $H(Y)$ nats of information from Z_Y to decode X. For a mean field decoder, this means that the decoder will only output a canonical member of each class. For a powerful decoder, such as an autoregressive decoder, it will learn to select a random member of the class.

For discrete Y, this model can trivially be turned into an unconditional generative model by first sampling Y from the training data or using any other appropriate procedure, such as sampling Y uniformly at random.

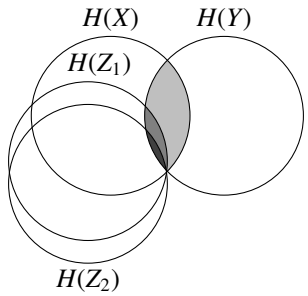

Figure 3: Information diagram for the basic hierarchical CEB model, $Z_2 \leftarrow Z_1 \leftarrow X \leftrightarrow Y$.

C.2 Bidirectional Generation

Given the presentation of conditional generation above, it is natural to consider that both $c(y|z)$ and $d(x|z)$ are conditional generative models of Y and X, respectively, and to learn a Z that can handle both tasks. This can be done easily with the following bidirectional CEB model: $Z_X \leftarrow X \leftrightarrow Y \rightarrow Z_Y$. This corresponds to the following factorization: $p(x, y, z_X, z_Y) \equiv p(x, y)e(z_X|x)b(z_Y|y)$. The two objectives from above then become the following single objective:

$$\min\ -H(Z_X|X) + H(Z_X|Y) + H(Y|Z_X) \tag{42}$$

$$-\ H(Z_Y|Y) + H(Z_Y|X) + H(X|Z_Y) \tag{43}$$

A natural question is how to ensure that Z_X and Z_Y are consistent with each other. Fortunately, that consistency is trivial to achieve by making the natural variational approximations: $p(z_Y|x) \rightarrow e(z_Y|x)$ and $p(z_X|y) \rightarrow b(z_X|y)$. The full bidirection variational CEB objective then becomes:

$$\begin{aligned}\min\ &\langle \log e(z_X|x)\rangle - \langle \log b(z_X|y)\rangle - \langle \log c(y|z_X)\rangle \\ &+ \langle \log b(z_Y|y)\rangle - \langle \log e(z_Y|x)\rangle - \langle \log d(x|z_Y)\rangle\end{aligned} \tag{44}$$

Here, we learn a unified Z that is consistent with both Z_X and Z_Y, permitting generation of either output given either input in the trained model, in the same spirit as Vedantam et al. (2018), but without any hyperparameter tuning.

C.3 Hierarchical CEB

Thus far, we have focused on learning a single latent representation (possibly composed of multiple latent variables at the same level). Here, we consider how to learn a hierarchical model with CEB.

Consider the graphical model $Z_2 \leftarrow Z_1 \leftarrow X \leftrightarrow Y$. This is the simplest hierarchical supervised representation learning model. The general form of its information diagram is given in Figure 3.

The key observation for generalizing CEB to hierarchical models is that the target mutual information doesn't change. By this, we mean that all of the Z_i in the hierarchy should cover $I(X; Y)$ at convergence, which means maximizing $I(Y; Z_i)$. It is reasonable to ask why we would want to train such a model, given that the final set of representations are presumably all effectively identical in terms of information content. The answer is simple: doing so allows us to train deep models in a principled manner such that all layers of the network are consistent with each other and with the data. We need to be more careful when considering the residual information terms, though – it is not the case that we want to minimize $I(X; Z_i|Y)$, which is not consistent with the graphical model. Instead, we want to minimize $I(Z_{i-1}; Z_i|Y)$, defining $Z_0 = X$.

This gives the following simple *Hierarchical CEB* objective:

$$CEB_{\text{hier}} \equiv \min \sum_i I(Z_{i-1}; Z_i|Y) - I(Y; Z_i) \tag{45}$$

$$\Leftrightarrow \min \sum_i -H(Z_i|Z_{i-1}) + H(Z_i|Y) + H(Y|Z_i) \tag{46}$$

Because all of the Z_i are targetting Y, this objective is as stable as regular CEB. Note that if all of the Z_i have the same dimensionality, in principle they may all use the same networks for $b(z_i|Y)$ and/or $c(y|z_i)$, which may substantially reduce the number of parameters in the model. All of the individual loss terms in the objective must still appear, of course. There is no requirement, however, that the Z_i have the same latent dimensionality, although doing so may give a unified hiearchical representation.

C.4 Sequence Learning

Many of the richest problems in machine learning vary over time. In Bialek & Tishby (1999), the authors define the *Predictive Information*:

$$I(X_{past}, X_{future}) = \left\langle \log \frac{p(x_{past}, x_{future})}{p(x_{past})p(x_{future})} \right\rangle$$

This is of course just the mutual information between the past and the future. However, under an assumption of temporal invariance (any time of fixed length is expected to have the same entropy), they are able to characterize the predictive information, and show that it is a subextensive quantity: $\lim_{T\to\infty} I(T)/T \to 0$, where $I(T)$ is the predictive information over a time window of length $2T$ (T steps of the past predicting T steps into the future). This concise statement tells us that past observations contain vanishingly small information about the future as the time window increases.

The application of CEB to extracting the predictive information is straightforward. Given the Markov chain $X_{<t} \to X_{\geq t}$, we learn a representation Z_t that optimally covers $I(X_{<t}, X_{\geq t})$ in *Predictive CEB*:

$$CEB_{\text{pred}} \equiv \min I(X_{<t}; Z_t|X_{\geq t}) - I(X_{\geq t}, Z_t) \tag{47}$$

$$\Rightarrow \min -H(Z_t|X_{<t}) + H(Z_t|X_{\geq t}) + H(X_{\geq t}|Z_t) \tag{48}$$

Note that the model entailed by this objective function *does not* rely on $Z_{<t}$ when predicting $X_{\geq t}$. A single Z_t captures all of the information in $X_{<t}$ and is to be used to predict as far forward as is desired. "Rolling out" Z_t to make predictions is a modeling error according to the predictive information.

Also note that, given a dataset of sequences, CEB_{pred} may be extended to a bidirectional model, as in Appendix C.2. In this case, two representations are learned, $Z_{<t}$ and $Z_{\geq t}$. Both representations are for timestep t, the first representing the observations before t, and the second representing the observations from t onwards. As in the normal bidirectional model, using the same encoder and backwards encoder for both parts of the bidirectional CEB objective ties the two representations together.

Modeling and architectural choices. As with all of the variants of CEB, whatever entropy remains in the data after capturing the entropy of the mutual information in the representation must be modeled by the decoder. In this case, a natural modeling choice would be a probalistic RNN with powerful decoders per time-step to be predicted. However, it is worth noting that such a decoder would need to sample at each future step to decode the subsequent step. An alternative, if the prediction horizon is short or the predicted data are small, is to decode the entire sequence from Z_t in a single, feed-forward network (possibly as a single autoregression over all outputs in some natural sequence). Given the subextensivity of the predictive information, that may be a reasonable choice in stochastic environments, as the useful prediction window may be small.

Multi-scale sequence learning. As in WaveNet (Van Den Oord et al., 2016), it is natural to consider sequence learning at multiple different temporal scales. Combining an architecture like time-dilated WaveNet with CEB is as simple as combining CEB_{pred} with CEB_{hier} (Appendix C.3). In this case, each of the Z_i would represent a wider time dilation conditioned on the aggregate Z_{i-1}. The advantage of such an objective over that used in WaveNet is avoiding unnecessary memorization of earlier timesteps.

C.5 Unsupervised CEB

Pure unsupervised learning is fundamentally an ill-posed problem. Without knowing what the task is, it is impossible to define an optimal representation directly. We think that this core issue is what lead the authors of Bengio et al. (2013) to prefer barely compressed representations. But by that line of

reasoning, it seems that unsupervised learning devolves to lossless compression – perhaps the correct representation is the one that allows you to answer the question: "What is the color of the fourth pixel in the second row?"

On the other hand, it also seems challenging to put the decision about what information should be kept into objective function hyperparameters, as in the β VAE and penalty VAE (Alemi et al., 2018) objectives. That work showed that it is possible to constrain the amount of information in the learned representation, but it is unclear how those objective functions keep only the "correct" bits of information for the downstream tasks you might care about. This is in contrast to all of the preceeding discussion, where the task clearly defines the both the correct amount of information and which bits are likely to be important.

However, unsupervised representation learning is still an interesting problem, even if it is ill-posed. Our perspective on the importance of defining a task in order to constrain the information in the representation suggests that we can turn the problem into a data modeling problem in which the practitioner who selects the dataset also "models" the likely form of the useful bits in the dataset for the downstream task of interest.

In particular, given a dataset X, we propose selecting a function $f(X) \to X'$ that transforms X into a new random variable X'. This defines a paired dataset, $P(X, X')$, on which we can use CEB as normal. Note that choosing the identity function for f results in maximal mutual information between X and X' ($H(X)$ nats), which will result in a representation that is far from the MNI for normal downstream tasks. In other words, representations learned by true autoencoders are unlikely to be any better than simply using the raw X.

It may seem that we have not proposed anything useful, as the selection of $f(.)$ is unconstrained, and seems much more daunting than selecting β in a β VAE or σ in a penalty VAE. However, there is a very powerful class of functions that makes this problem much simpler, and that also make it clear using CEB will *only* select bits from X that are useful. That class of functions is the noise functions.

C.5.1 Denoising CEB Autoencoder

Given a dataset X without labels or other targets, and some set of tasks in mind to be solved by a learned representation, we may select a random noise variable U, and function $X' = f(X, U)$ that we believe will destroy the irrelevant information in X. We may then add representation variables $Z_X, Z_{X'}$ to the model, giving the joint distribution $p(x, x', u, z_X, z_{X'}) \equiv p(x)p(u)p(x'|f(x, u))e(z_X|x)b(z_{X'}|x')$. This joint distribution is represented in Figure 4.

Denoising Autoencoders were originally proposed in Vincent et al. (2008). In that work, the authors argue informally that reconstruction of corrupted inputs is a desirable property of learned representations. In this paper's notation, we could describe their proposed objective as $\min H(X|Z_{X'})$, or equivalently $\min \langle \log d(x|z_{X'} = f(x, \eta)) \rangle_{x,\eta \sim p(x)p(\theta)}$.

Here we make this idea somewhat more formal through the MNI criterion and the derivation of CEB as the optimal objective for that criterion. We also note that, practically speaking, we would like to learn a representation that is consistent with uncorrupted inputs as well. Consequently, we are going to use a bidirectional model.

$$CEB_{\text{denoise}} \equiv \min I(X; Z_X|X') - I(X'; Z_X) + I(X'; Z_{X'}|X) - I(X; Z_{X'}) \tag{49}$$

$$\Rightarrow \min -H(Z_X|X) + H(Z_X|X') + H(X'|Z_X) - H(Z_{X'}|X') + H(Z_{X'}|X) + H(X|Z_{X'}) \tag{50}$$

This requires two encoders and two decoders, which may seem expensive, but it permits a consistent learned representation that can be used cleanly for downstream tasks. Using a single encoder/decoder pair would result in either an encoder that does not work well with uncorrupted inputs, or a decoder that only generates noisy outputs.

If you are only interested in the learned representation and not in generating good reconstructions, the objective simplifies to the first three terms. In that case, the objective is properly called a *Noising CEB Autoencoder*, as the model predicts the noisy X' from X:

$$CEB_{\text{noise}} \equiv \min I(X; Z_X|X') - I(X'; Z_X) \tag{51}$$

$$\Rightarrow \min -H(Z_X|X) + H(Z_X|X') + H(X'|Z_X) \tag{52}$$

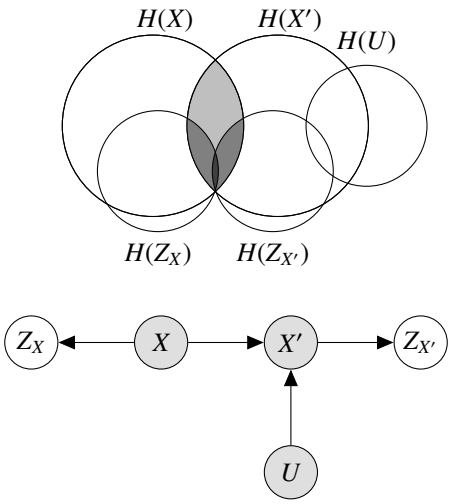

Figure 4: Information diagram and graphical model for the Denoising CEB Autoencoder.

In these models, the noise function, $X' = f(X, U)$ must encode the practitioner's assumptions about the structure of information in the data. This obviously will vary per type of data, and even per desired downstream task.

However, we don't need to work too hard to find the perfect noise function initially. A natural first choice for f is:[7]

$$f(x, \eta) = \text{clip}(x + \eta, \mathcal{D}) \tag{53}$$

$$\eta \sim \lambda U(-1, 1) * \mathcal{D} \tag{54}$$

$$\mathcal{D} = \text{domain}(X) \tag{55}$$

In other words, add uniform noise scaled to the domain of X and by a hyperparameter λ, and clip the result to the domain of X. When $\lambda = 1$, X' is indistinguishable from uniform noise. As $\lambda \to 0$, this maintains more and more of the original information from X in X'. For some value of $\lambda > 0$, most of the irrelevant information is destroyed and most of the relevant information is maintained, if we assume that higher frequency content in the domain of X is less likely to contain the desired information. That information is what will be retained in the learned representation.

Theoretical optimality of noise functions. Above we claimed that this learning procedure will only select bits that are useful for the downstream task, given that we select the proper noise function. Here we prove that claim constructively. Imagine an oracle that knows which bits of information should be destroyed, and which retained in order to solve the future task of interest. Further imagine, for simplicity, that the task of interest is classification. What noise function must that oracle implement in order to ensure that $CEB_{denoise}$ can only learn exactly the bits needed for classification? The answer is simple: for every $X = x_i$, select $X' = x'_i$ uniformly at random from among all of the $X = x_j$ that should have the same class label as $X = x_i$. Now, the only way for CEB to maximize $I(X; Z_{X'})$ and minimize $I(X'; Z_{X'})$ is by learning a representation that is isomorphic to classification, and that encodes exactly $I(X; Y)$ nats of information, even though it was only trained "unsupervisedly" on X, X' pairs. Thus, if we can choose the correct noise function that destroys only the bits we don't care about, $CEB_{denoise}$ will learn the desired representation and nothing else (caveated by model, architecture, and optimizer selection, as usual).

C.6 Semi-Supervised CEB

Given *any* amount of paired data X, Y immediately improves our ability to learn a semantic representation. Fortunately, it is easy to reincorporate paired data in combination with noising and

[7] White noise is probably a better choice for audio signals, and may be the right choice for most real-valued signals, including images and videos.

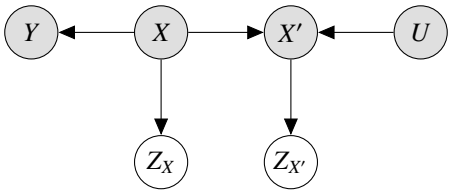

Figure 5: Graphical model for Semi-Supervised CEB.

denoising CEB, introduced above. We present the assumed graphical model in Figure 5. We give the corresponding *Semi-Supervised CEB* directly:

$$
\begin{aligned}
CEB_{semi} &\equiv \min I(X; Z_X|X') - I(X'; Z_X) + I(X'; Z_{X'}|X) - I(X; Z_{X'}) && (56)\\
&\quad + \mathbf{1}_{Y\in(X,Y)}\big[I(X'; Z_{X'}|Y) - I(Y; Z_{X'})\big] && (57)\\
&\Rightarrow \min -H(Z_X|X) + H(Z_X|X') + H(X'|Z_X) - H(Z_{X'}|X') + H(Z_{X'}|X) + H(X|Z_{X'}) && (58)\\
&\quad + \mathbf{1}_{Y\in(X,Y)}\big[-H(Z_{X'}|Y) + H(Z_{X'}|Y) + H(Y|Z_{X'})\big] && (59)
\end{aligned}
$$

$\mathbf{1}_{Y\in(X,Y)}$ is the indicator function, equal to 1 when a Y is part of the paired data, and equal to 0 otherwise. In other words, if we have $Y = y$ paired with a given $X = x$, we can include those terms in the objective. If we do not have that, we can simply leave them out.

Note that it is straightforward to generalize this to semisupervised learning with two or more observations that are both being learned unsupervisedly, but also have some amount of paired data. For example, images and natural language, assuming we have a reasonable noise model for unsupervisedly learning natural language.

D Visualizations

Here we provide some visualizations of the Fashion MNIST tasks.

In Figure 6, we show a trained 2D CEB latent representation of Fashion MNIST. The model learned to locate closely related concepts together, including the cluster of "shirt" classes near the center, and the cluster of "shoe" classes toward the lower right. In spite of the restriction to 2 dimensions, this model achieves $\sim 92\%$ on the test set.

In Figure 7, the 10,000 test images and their 10,000 adversaries are shown for four of the models. It is easy to see at a glance that the CEB model organizes all of the adversaries into the "trousers" class, with a crisp devision between the true examples and the adversaries. In contrast, the two VIB models have adversaries mixed throughout. However, all three models are clearly preferable to the deterministic model, which has all of the adversaries mixed into the "trousers" class with no ability to distinguish between adversaries and true examples.

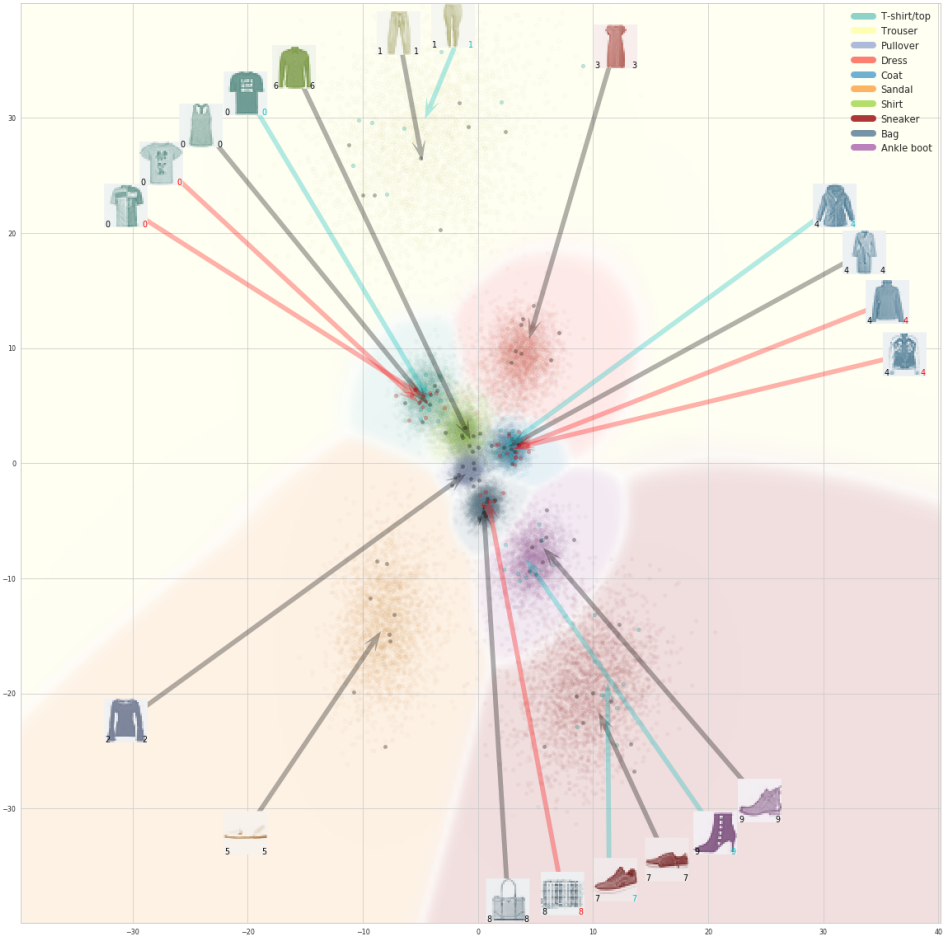

Figure 6: A trained 2D latent space for a Fashion MNIST CEB model.

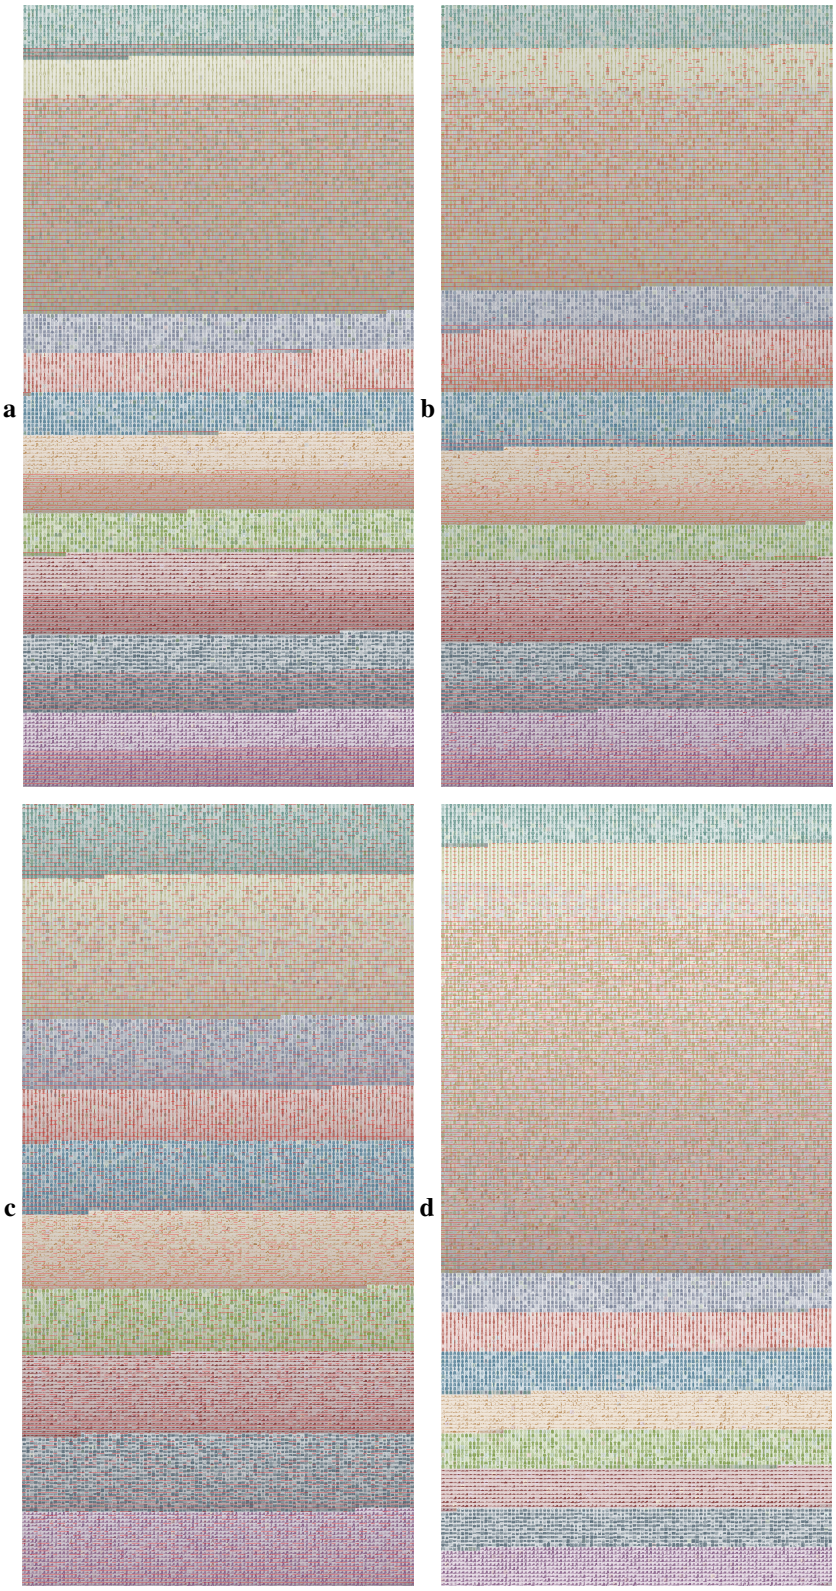

Figure 7: All adversarial images sorted by predicted class, showing the difference between robust and non-robust models. Each predicted class is sorted by the model's rate, R (H is used for **d**), from low to high. Images with a red bar along their top are adversarial. **a** is CEB, **b** is $\text{VIB}_{0.5}$, **c** is $\text{VIB}_{0.01}$, **d** is Determ.