Dueling Decoders: Regularizing Variational Autoencoder Latent Spaces

Bryan Seybold, Emily Fertig, Alex Alemi, Ian Fischer
Google Research
{seybold,emilyaf,alemi,iansf}@google.com

Abstract

Variational autoencoders learn unsupervised data representations, but these models frequently converge to minima that fail to preserve meaningful semantic information. For example, variational autoencoders with autoregressive decoders often collapse into autodecoders, where they learn to ignore the encoder input. In this work, we demonstrate that adding an auxiliary decoder to regularize the latent space can prevent this collapse, but successful auxiliary decoding tasks are domain dependent. Auxiliary decoders can increase the amount of semantic information encoded in the latent space and visible in the reconstructions. The semantic information in the variational autoencoder's representation is only weakly correlated with its rate, distortion, or evidence lower bound. Compared to other popular strategies that modify the training objective, our regularization of the latent space generally increased the semantic information content.

1. Introduction

For many problem domains, large quantities of data are available without labels. One approach to leveraging such data is to convert the high-dimensional raw data into a low-dimensional *representation* using some unsupervised learning procedure. A desirable property in a learned representation is that nearby representations share semantically meaningful information about the data, such as task labels or world state [8]. Among the many uses of representation learning are simplifying reinforcement learning (e.g., [20]) and utilizing unlabeled data for semi-supervised classification tasks (e.g. [27, 3]).

One popular method for unsupervised representation learning is **variational autoencoders (VAEs)** [29]. VAEs are a probabilistic extension of autoencoders, which are themselves non-linear extensions of principal components analysis [29, 8]. VAEs learn a generative model of the data via approximate variational inference and reconstruct inputs encoded into a probabilistic latent space. The latent space is a learned representation that often has the property that
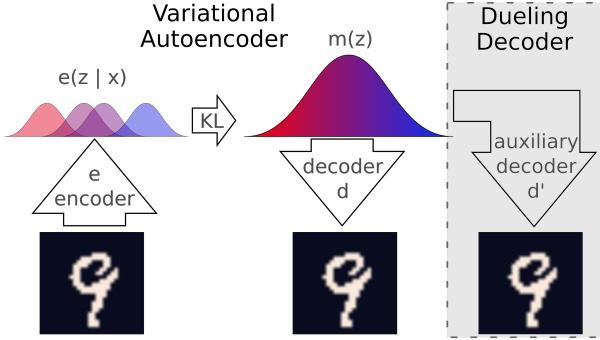
Figure 1. Overview of a variational autoencoder with auxiliary decoder for regularization. Images are encoded on the left into an encoder distribution, $e(z|x)$. The approximate marginal distribution, $m(z)$, learns to match the mixture of all of the encoder distributions across the training data, and samples from the encoder or the marginal can be decoded by sampling from the decoder $d(x|z)$. The additional, lightly shaded, auxiliary decoder $d'(x|z)$ is a central focus of our paper—a regularizer that helps to learn more semantically meaningful representations.

nearby points share the same task label. One criticism of early VAEs was that the reconstructions and samples they produce are often of lower visual quality compared to models such as generative adversarial networks [18]. The introduction of autoregressive decoders in VAEs led to models that produce outputs of better visual quality but often ignore the latent space [33, 40, 35, 19, 12, 28, 10]. Researchers have modified the VAE training objective in several ways, such as "free bits", KL annealing, Penalty-VAE, or adding an auxiliary decoder to encourage models to utilize the latent space [10, 28, 23, 2, 11, 48]. The theoretical basis of VAEs for representation learning is still being developed and the practical consequences of modifying the training objective for learned representations is not yet understood [12, 11, 2].

In this work, we explore how modifications of learning objective (but not the architecture or optimization procedure) influence the semantic information in the learned representation and reconstructions. We measure the ability of the network to learn a semantically meaningful representation by observing whether the encoded input can be decoded into the semantic label and whether reconstructed images

maintain the same label as the inputs. We demonstrate that common measures of VAE performance, including the **evidence lower bound (ELBO)**, rate, and distortion, may be poor proxies for measuring semantic information retained in the latent representation. We identify that adding an auxiliary decoder can greatly improve the semantic information in the learned representation, but that different auxiliary decoders are required for image labeling than for natural language processing.

To summarize, our main contributions are:

- a comparison across training objective variants showing that auxiliary tasks improve semantic information efficiently relative to other approaches;
- a demonstration that the effective auxiliary tasks for computer vision tasks can be distinct from effective auxiliary tasks for natural language processing;
- a simple procedure for measuring the semantic information maintained in VAE reconstructions that measures how well the decoder represents semantic information;
- that the semantic information retained by VAEs is not well predicted by the ELBO, distortion, or rate of the VAE;
- that VAEs with the same architecture and rate encode latent spaces with different amounts of semantic information;
- that autoregressive decoders suffer from local minima that lead to posterior collapse early in training, but after converging to a good ELBO, models are stable.

2. Related Work

2.1. Understanding Autoregressive Decoders Training Objectives

PixelRNN introduced the first tractable autoregressive artificial neural network generative model of images [33]. This was extended to feed-forward convolutional networks with PixelCNN and PixelCNN++ [33, 40, 35]. These generative models can achieve the same or better log probability as VAEs on many image tasks even without conditioning on an input image [29, 2, 33]. It is natural to consider combining these more powerful generative models with the VAE framework, but the naive approach results in the trained decoder ignoring the latent representation and becomes an autodecoder instead [12, 2, 41]. [12, 2] claim that any sufficiently powerful decoder will collapse into an autodecoder when optimizing the ELBO objective. Modifications to the ELBO have been proposed to mitigate this issue. The simplest such modification, the β-VAE [23], adds a Lagrange multiplier to the KL term of the ELBO. Other efforts have been made to better understand and control VAE optimization with autoregressive decoders, including KL annealing [10], allowing the model "free bits" [28, 12], Penalty-

VAE [2, 11], or adding an auxiliary decoding task [48]. More drastic changes modified the optimization procedure or modified the model architecture to try to circumvent these issues [47, 45, 26, 15, 21]. Changing the optimization procedure or architecture is in some senses orthogonal to changes in the loss, so we limit our study to understanding how the loss shapes the learned representation. In this work, we compare many modifications to the ELBO using classification-based measures of semantic information in the learned representations.

2.2. Quantifying Learned Representations

A common method of estimating the quality of the learned latent representation is through visual inspection [28] or reporting numbers that are based on the training objective, such as bits-per-pixel or rate. In contrast, we seek to measure the amount of information about image labels in the encoded space and reconstructions. While the rate is the upper-bound on the semantic information, the information regarding labels or other semantics can be much lower. One means of assessing learned representations is to assess "disentanglement" where semantically meaningful information falls along independent axes or manifold directions in the latent space [23, 11]. Both [23] and [11] increase disentanglement by increasing β, which can exacerbate posterior collapse with PixelCNN models [2]. It is unclear how discrete labels should map onto a small number of disentangled axes, so we do not evaluate disentanglement. An alternative method of quantifying learned representation quality is measuring the classification accuracy from latent space values to semantic labels. Linear classifiers are associated with interpretability, but work in low dimensional latent spaces shows intuitively that representations do not need to be linearly separable in order to be easily interpretable [8, 5, 4, 14, 6, 16]. None of these methods examine the end-to-end quality of the VAE including whether the reconstructed images respect semantic information in the input.

2.3. Regularization and Auxiliary Losses

Many approaches to ease training and improve generalization have been proposed for deep networks. Some of the most popular are BatchNorm and other forms of normalization [24, 7, 43], Dropout [37], adding skip-connections [22, 15], and L1 or L2 regularization. Auxiliary losses are commonly used when training deep networks in a variety of situations, including deterministic classification and object detection [38], reinforcement learning [25], and probabilistic models [42]. Recent work in probabilistic models has explored a variety of what may be considered auxiliary losses in the VAE family of objectives. [48] proposed a bag-of-words auxiliary decoder to a variational autoencoder for natural language processing. The unpublished [31] pro-

1

posed adding a similar pixel-independent auxiliary loss to a PixelVAE architecture. Neither [48] nor [31] measured the semantic information extracted by the autoencoder. [36] shares the same latent space across multiple encoders and decoders for cross-modal learning. [12] discusses limiting autoregressive models to only generated portions of an image. [47] adds an adversarial discriminator network to prevent the VAE from becoming an autodecoder. [27] and [3] added a classifier on the latent decoder to perform semi-supervised learning. We are unaware of any work that attempts to quantify the quality of the resulting learned representations, which is the core focus of this work.

3. Regularizing Latent Spaces with Dueling Decoders

3.1. Preliminaries

Variational inference is an approximate Bayesian method for fitting the parameters of generative models [9, 32]. Following the notation of [2], given data X, we may learn a representation $Z \leftarrow X$ by selecting a parameterized *encoder* $e(z|x)$, a parameterized *marginal* $m(z)$, and a parameterized *decoder* $d(x|z)$, and maximizing the ELBO [29]:

$$\mathbb{E}_{e(z|x)}[\log d(x|z)] - D_{KL}[e(z|x)||m(z)] \qquad (1)$$

where D_{KL} is the Kullbach-Leibler Divergence. This objective may be augmented with a hyperparameter, β [23]:

$$\mathbb{E}_{e(z|x)}[\log d(x|z)] - \beta D_{KL}[e(z|x)||m(z)] \qquad (2)$$

This hyperparameter may be viewed as a regularization that encourages disentangled representations Z [23, 11], or from an information-theoretic perspective where β controls the *rate* (R) of information Z maintains about X [2]:

$$R \equiv \int dx\, p(x) \int dz\, e(z|x) \log \frac{e(z|x)}{m(z)} \geq I(X; Z) \quad (3)$$

where $I(X; Z)$ is the mutual information between X and Z. R is the D_{KL} term from the ELBO and bounds the information extracted from the input. Similarly, we may identify the *distortion* (D) as the first term of the ELBO:

$$D \equiv \int dx\, p(x) \int dz\, e(z|x) \log d(x|z)$$
$$\leq I(X; Z) - H(X) \qquad (4)$$

Thus, the ELBO may be understood as simultaneously minimizing an upper bound and maximizing a lower bound on $I(X; Z)$, and β as a Lagrange multiplier that allows us to trade off between those two bounds.

3.2. Semantic Representation Accuracy

While the rate is an upper bound on information extracted from the VAE inputs, it does not require that semantically relevant information is extracted. We may directly measure the semantic information in the latent representation by separately training a classifier $P(Y|Z)$ on a labeled training set X, Y and corresponding Z sampled N times from $e(z|x)$ for each $x \in X$ and measuring the accuracy of the classifier on a held-out labeled test set:

$$\frac{1}{N} \sum_N \mathbf{1}_{Y=\mathrm{argmax}\, P(Y|Z)}, \qquad (5)$$

where $\mathbf{1}_p$ is the indicator function (1 when p is true, 0 otherwise). The label distortion is an information theoretic measurement of the semantic information in the latent encoding derived as the classifier's conditional entropy:

$$H(Y|Z) = H(Y) - I(Y; Z). \qquad (6)$$

3.3. Semantic Reconstruction Accuracy

To quantify semantic information from end-to-end of the VAE, we propose a straightforward measurement of the accuracy with which reconstructions maintain the original class label of a dataset. For a dataset, X, with labels, Y, and N reconstructions of the data from a VAE, X', we train a classifier to maximize the likelihood of $P(Y|X)$ on the original training set, X, and determine the **semantic reconstruction accuracy** of the model as

$$\frac{1}{N} \sum_N \mathbf{1}_{Y=\mathrm{argmax}\, P(Y|X')} \qquad (7)$$

on the reconstructions from the test set, X'. In other words, we train a classifier and then measure the classification accuracy on reconstructions $x' \sim p(x)e(z|x)d(x'|z)$.

3.4. Dueling Decoders

The typical VAE training objective has no term to explicitly maximize semantic information. The semantic labels needed for such a term are not available during unsupervised learning. The only option in unsupervised learning is modifying the optimization to encourage semantic information. One strategy is to find auxiliary tasks that correlate with semantic information. Both [48] and [31] added auxiliary tasks to VAEs, but only to prevent rate collapse. Neither studied which auxiliary tasks are appropriate for computer vision or measured how the auxiliary task changed semantic information extracted by the autoencoder.

We formalize auxiliary tasks as adding an auxiliary decoder, $d'(x|z)$. This decoder shares the same marginal distribution and latent space as the original decoder. Because two different decoders are competing to shape the latent

space, we call this method **Dueling Decoders**. The auxiliary decoder adds another reconstruction term to the training objective.

$$\mathbb{E}_{e(z|x)}[\log d(x|z) + \lambda \log d'(x|z)] - \beta D_{KL}[e(z|x)||m(z)] \tag{8}$$

It is worth noting that this regularization does not change the probabilistic interpretation of the distortion, rate, or ELBO of the original model. During evaluation, the auxiliary decoder is removed and the model is evaluated as a standard VAE. The auxiliary decoder only influences what is learned in the latent space of the model. This makes it distinct from methods that change the optimization or architecture [47, 45, 26, 15, 21].

This formalization holds across a wide range of auxiliary decoding tasks, but does not reveal the optimal auxiliary task to increase semantic information. The search for optimal auxiliary tasks is a primary pursuit in unsupervised learning. For example, if $d(x|z) \propto d'(x|z)$, Eq. 8 is just reweighting the the terms in the original ELBO. In contrast, if the primary decoder is a PixelCNN and the auxiliary decoder reconstructs the color histogram, the latent space feeding into the PixelCNN will be encouraged to encode additional color information. If this color information is related to label information, it will ensure that semantic information related to the labels is available. Alternatively, if the primary decoder is a PixelCNN and the auxiliary decoder is a pixel-independent CNN decoder, the model will be encouraged to learn global structure. The ideal auxiliary decoder will depend on the exact downstream task and model architecture.

Because the ELBO is bounded by the information in the data set, adding an auxiliary decoder can at best change where the model settles in the rate-distortion plane. Adding regularization breaks the symmetry of points within the ELBO objective similarly to how manipulating β changes the rate-distortion performance but can have little effect on the final ELBO [2]. However, an outstanding question is whether regularizing the latent space with an auxiliary decoder captures different amounts of semantic information than modifying the ELBO objective terms.

4. Experiments

4.1. Implementation Details

We evaluate VAEs with different decoder architectures and hyperparameters on 4 data sets. Two data sets are binary versions of MNIST [30]: Stochastic MNIST [34, 1] and Thesholded MNIST [17]. In Stochastic MNIST, pixel values were assigned to 0 or 1 with a probability dependent on their intensity. In Thresholded MNIST, all pixel values are assigned to 1 if greater than 127 and to 0 otherwise. Stochastic MNIST is commonly used in the VAE literature

[29, 2, 28, 27, 39], but the process of sampling pixel values independently decreases the degree of local correlation in the images. Thresholded MNIST preserves local correlations in each image. We also use two 8-bit data sets, MNIST [30] and Fashion MNIST [44]. Each data set has 10 labels and label entropy, $H(Y)$, of 2.3 nats.

For each data set, we train a range of VAEs. (The full architectures are listed in the supplemental). All VAEs use the same convolutional neural network architecture for the encoder and a 280 component VAMPPrior marginal [39]. The encoder outputs are means and lower triangular representations of covariance matrices of multivariate normal distributions. We use three different types of decoders: pixel-independent convolution decoders (CNN), autoregressive PixelCNN++ decoders [35] (PixelCNN), PixelCNN++ decoders with auxiliary decoders (Dueling Decoders) as described in Section 3.4. Unless otherwise noted, the auxiliary decoder is a pixel-independent convolutional decoder. Note that, at test time, the PixelCNN and Dueling Decoder architectures are identical. The architectures were chosen to achieve ELBOs between 80 and 90 nats on Stochastic MNIST (approximately matching results in the literature [2, 28]), and we fixed the architecture to see if the model could generalize across data sets. For binary data sets, the decoder output distribution is a Bernoulli, and for 8-bit data sets, the output distribution is a mixture of quantized logistics [35]. We perform a grid search over the following parameters for each model: β (0.1, 1.0), training batch size (32, 64), number of latent dimensions (2, 16, 64), and dueling decoder regularization λ (0.1, 1.0). We train each model for 20000 steps (12 or 25 epochs depending on batch size) with the Adam optimizer and time varying learning rate, $1^{-3} * 0.66^{t/1^{-4}} * (1 - 0.66^{t/1^{-3}}) + 1^{-10}$, where t is the step number. Reconstructions from different models with rate<10 are available in Figure 2.

To measure the semantic representation accuracy, we trained a linear classifier and a multilayer perceptron (MLP) from the encoded latent values from the test set and report the accuracy of those classifiers on the test set as described in Section 3.2.

To measure the semantic reconstruction accuracy, we train a convolutional neural network classifier on each data set. The network is trained on the original training set and validated on the test set to have >98% accuracy on each MNIST variant and 94% accuracy on the more difficult Fashion MNIST data set. The classifier is only trained on the original training set, not on reconstructions. The classifier was applied to 300 evenly sampled reconstructions from the test set to test how well the reconstructions match the semantics of the original data set as described in Section 3.3.
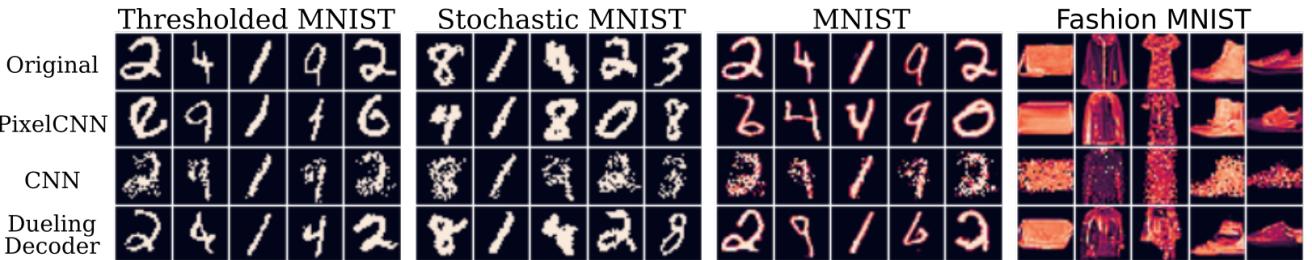
Figure 2. Reconstructions from each decoder class with rate<10 and the best semantic reconstruction accuracy. The top row is the original image from the data set. The rows beneath are reconstructions from the PixelCNN, CNN, and Dueling Decoder decoders respectively.

Decoder	β	ELBO	Distortion	Rate	Latent Accuracy MLP	Label Distortion MLP	Latent Accuracy Linear	Reconstruction Accuracy
STOCHASTIC MNIST								
PixelCNN	0.1	83.2	74.6	8.57 ± 0.16	0.48 ± 0.03	1.34 ± 0.08	0.36 ± 0.05	0.36 ± 0.03
CNN	1.0	136.7	129.8	6.91 ± 0.09	0.81 ± 0.03	0.65 ± 0.04	$\mathbf{0.55 \pm 0.12}$	$\mathbf{0.73 \pm 0.02}$
Dueling Decoder	1.0	81.3	72.1	9.25 ± 0.72	$\mathbf{0.91 \pm 0.02}$	$\mathbf{0.31 \pm 0.06}$	$\mathbf{0.72 \pm 0.06}$	$\mathbf{0.79 \pm 0.03}$
THRESHOLDED MNIST								
PixelCNN	0.1	60.6	52.6	8.01 ± 0.18	0.58 ± 0.06	1.07 ± 0.14	0.39 ± 0.02	0.47 ± 0.07
CNN	1.0	130.4	123.4	7.04 ± 0.11	0.82 ± 0.01	0.59 ± 0.02	$\mathbf{0.62 \pm 0.07}$	0.81 ± 0.02
Dueling Decoder	1.0	56.4	48.4	7.93 ± 0.18	$\mathbf{0.92 \pm 0.01}$	$\mathbf{0.27 \pm 0.02}$	$\mathbf{0.70 \pm 0.02}$	$\mathbf{0.88 \pm 0.01}$
MNIST								
PixelCNN	0.1	560.1	552.8	7.29 ± 0.08	0.42 ± 0.04	1.52 ± 0.06	0.27 ± 0.03	0.37 ± 0.02
CNN	1.0	891.6	883.6	7.94 ± 0.28	$\mathbf{0.79 \pm 0.01}$	$\mathbf{0.58 \pm 0.03}$	0.51 ± 0.06	$\mathbf{0.78 \pm 0.01}$
Dueling Decoder	1.0	564.6	554.8	9.78 ± 0.96	$\mathbf{0.88 \pm 0.04}$	$\mathbf{0.41 \pm 0.11}$	$\mathbf{0.58 \pm 0.05}$	$\mathbf{0.79 \pm 0.04}$
FASHION MNIST								
PixelCNN	1.0	1492.2	1484.0	8.24 ± 0.17	$\mathbf{0.78 \pm 0.01}$	$\mathbf{0.63 \pm 0.03}$	$\mathbf{0.52 \pm 0.02}$	$\mathbf{0.76 \pm 0.02}$
CNN	1.0	1995.3	1986.0	9.31 ± 0.04	0.71 ± 0.01	0.82 ± 0.02	$\mathbf{0.54 \pm 0.03}$	0.56 ± 0.01
Dueling Decoder	1.0	1495.1	1486.4	8.76 ± 1.34	$\mathbf{0.72 \pm 0.05}$	$\mathbf{0.79 \pm 0.17}$	$\mathbf{0.47 \pm 0.08}$	$\mathbf{0.69 \pm 0.03}$
PixelCNN+	1.0	1473.3	1469.2	4.08 ± 0.81	$\mathbf{0.49 \pm 0.14}$	$\mathbf{1.23 \pm 0.28}$	$\mathbf{0.30 \pm 0.05}$	0.43 ± 0.12
Dueling Decoder+	1.0	1480.2	1471.2	9.03 ± 0.11	$\mathbf{0.68 \pm 0.02}$	$\mathbf{0.94 \pm 0.05}$	$\mathbf{0.36 \pm 0.06}$	$\mathbf{0.67 \pm 0.01}$

Table 1. For each data set and each decoder, we report results for the model with rate<10 and the highest semantic representation accuracy (MLP). Values are the means of three randomly initialized training runs $\pm$ one standard deviation. Bold values are not significantly different from the best value (t-test, equal variance, α=0.05, Bonferroni corrected n=2). Decoders marked with + are enlarged networks.

4.2. Dueling Decoders exceed the best of independent and autoregressive decoders.

Although decoder architectures are not the focus of this paper, it is illustrative to show how the choice of decoder changes the semantic information the VAE learns to extract (Figure 2 and Table 1). In agreement with [2], models with the same decoder demonstrate a trade off in rate-distortion space due to the limited amount of information in the data set (see supplement). We focus further analysis on these low-rate models because the rate versus accuracy curve in Figure 3 has an inflection point just above 10 indicating that accuracy increases most rapidly at low rates, the entropy of all four training sets is upper-bounded by $\log 60000 = 11$ nats (the number of data values), and compact representations are a desirable property. Results across

the range of low rate models (<20 nats) show similar trends (Fig. 3). Across Thresholded, Stochastic, and the original MNIST, VAEs with autogregressive PixelCNN decoders achieved low distortions, but encoded little semantic information (all semantic accuracies < 0.6). In contrast, VAEs with independent CNN decoders encoded more semantic information (reconstruction and MLP semantic accuracies > 0.7), but suffered from high distortion. (At higher rates, CNN decoders can also achieve good distortions on MNIST, but not at rates less than 10 nats.) In contrast, using a Pixel-CNN architecture with an auxiliary CNN decoder in a Dueling Decoder setup achieves both low distortions and high semantic information (reconstruction and MLP accuracies > 0.7). Notably, the Dueling Decoder strategy achieves better distortions, rates, and trends towards better seman-
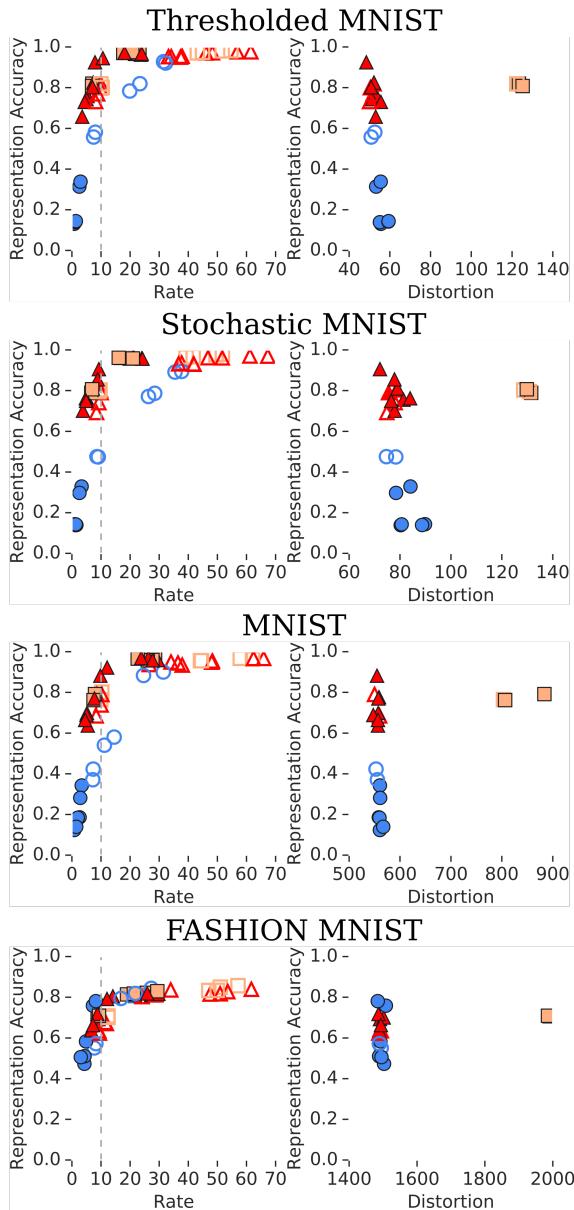
Figure 3. Correlations between semantic representation accuracy and the rate or distortion. **left:** Rate versus accuracy for models marked by decoder class: PixelCNN (●), CNN (■), and Dueling Decoder (▲). Filled or empty markers denote $\beta = 1.0$ or $\beta = 0.1$ respectively. The vertical dashed line indicates the 10 nat threshold for low-rate. **right:** Distortion versus accuracy for the low-rate models. Best viewed in color.

tic accuracies on MNIST than either the PixelCNN or CNN decoders alone.

It worth noting that the PixelCNN models with the highest representation accuracy on MNIST variants were trained with $\beta = 0.1$ because the rate term collapses to near zero for $\beta = 1.0$ (Figure 3). Because the rate is an upper bound on any information, models with zero rate cannot contain semantic information. Decreasing β or another manipula-

Modification	Rate	Latent Accuracy MLP	Reconstr. Accuracy
PixelCNN $\beta = 1$	0.97	0.14 ± 0.01	0.11 ± 0.01
PixelCNN $\beta = 0.1$	31.67	0.93 ± 0.01	0.91 ± 0.01
Dueling Decoder	8.13	$\mathbf{0.92 \pm 0.01}$	$\mathbf{0.88 \pm 0.01}$
KL Annealing	9.16	$\mathbf{0.92 \pm 0.01}$	$\mathbf{0.89 \pm 0.02}$
Free Bits	10.16	0.81 ± 0.01	0.74 ± 0.02
Penalty-VAE	10.04	0.53 ± 0.05	0.41 ± 0.05

Table 2. Comparison of training modifications to avoid collapsing into an autodecoder on Thresholded MNIST. Bold font indicates values not significantly different from the best value in the bottom four rows (t-test, equal variance, α=0.05, Bonferroni corrected n=3).

tion is necessary to potentially increase semantic information. In contrast, both CNN decoders and Dueling Decoders maintained a rate larger than the entropy in the labels without manipulating β. The semantic information retained by CNN and Dueling Decoders is higher than PixelCNN with $\beta = 0.1$ even at similar rates. This indicates that the increased rate achieved by decreasing β in these models encodes style information or noise rather than semantic information. Decreasing β does not preferentially guide the model to encode semantic information.

On Fashion MNIST, PixelCNN at $\beta = 1.0$ does not collapse with all hyperparameter combinations (particularly for larger batch sizes). When the model does not collapse, it encodes relevant semantic information well. However, adding a Dueling Decoder consistently prevented collapse and increased semantic information. Adding a Dueling Decoder can therefore simplify hyperparameter searches without hurting performance if unnecessary. Furthermore, requiring a larger batch size to avoid collapse limits the memory available for the network architecture. Increasing the number of layers and number of kernels in the PixelCNN decoder can yield better ELBOs and distortions (1469.2 $\pm$ 3.2 versus 1484.0 $\pm$ 1.7), but limits the batch size to 32. The model collapses again. In contrast, the enlarged Dueling Decoder model trends towards maintaining more information in the latents and generates significantly more semantically accurate reconstructions.

4.3. Comparing different modified training objectives.

Because the Dueling Decoder and the PixelCNN have the same architecture at test time, the difference between the two models is only in the objective that they optimize. We compared our modification of the training objective to other modifications of the training objective from the literature in Table. 2. All models used the same architecture and procedure at inference time. Manipulating β resulted in either a near-zero rate and low-accuracy reconstructions or a very-high rate. Dueling Decoders and KL Annealing [10]

both outperformed Free Bits [28] and Penalty-VAE [11, 2] approaches (0.92 and 0.92 semantic representation accuracy vs 0.81 and 0.53). Dueling Decoders achieved similar accuracy at a marginally lower rate and was easier to train because the objective does not change over time as in KL Annealing. As KL annealing and Dueling Decoders modify different portions of the loss, these two could be combined in the future.

4.4. Auxiliary decoders requiring spatial information increase semantic information.

All Dueling Decoder results up to this point in the paper have used a pixel-independent CNN auxiliary decoder, but it is initially unclear what makes this auxiliary task effective at increasing semantic information. Previously published work has used a bag-of-words decoder in a natural language processing model [48], which corresponds to a color or intensity histogram decoder. To determine how the choice of auxiliary tasks determines the semantic information learned, we tested four decoders (Table 3): 1) an intensity histogram decoder that reconstructing the distribution of intensity values across the entire image, 2) a marginal histogram decoder that reconstructs the distribution of intensity values in each row and each column independently, 3) the CNN decoder that reconstructs the distribution of intensity values in each pixel independently, and 4) an image gradient decoder that reconstructs the distribution of local intensity changes in the vertical and horizontal directions. Later decoders require encoding additional spatial information. Convolutional and image gradient decoders, which require more spatial information, increase the amount of semantic information learned by the VAE over marginal and histogram decoders (latent accuracy - MLP: 0.94 and 0.91 vs 0.82 and 0.50). Notably, an auxiliary histogram decoder does not increase the semantic information relative to the PixelCNN baseline (latent accuracy - MLP: 0.58 vs 0.50). Training supervised classifiers on the reconstruction targets (e.g. intensity histogram) yields the same order of performance. Therefore, it is important to identify the correct auxiliary decoder for the semantics of interest, supervised performance on reconstruction targets is a useful proxy for effectiveness, and natural language processing and image classification require different auxiliary tasks.

4.5. The ELBO defines a Valley of Stable Points

Dueling Decoder models demonstrate much better semantic reconstruction performance on MNIST tasks than purely autoregressive models even at the same rate, distortion, and model architecture. This indicates autoregressive decoders are capable of utilizing semantic information present in the latent space at $\beta = 1.0$. However, autoregressive models all converge to similar ELBO values (approximately 56 nats on Thresholded MNIST) even when the

Auxiliary	Rate	Latent Accuracy MLP	Reconst. Accuracy	Supervised Accuracy on Target
None	7.99	0.58 ± 0.06	0.41 ± 0.05	
Conv	9.75	$\mathbf{0.94 \pm 0.01}$	$\mathbf{0.90 \pm 0.03}$	0.99
Gradient	9.38	0.91 ± 0.01	$\mathbf{0.88 \pm 0.01}$	0.99
Marginals	9.03	0.82 ± 0.02	0.76 ± 0.02	0.90
Histogram	8.09	0.50 ± 0.01	0.42 ± 0.01	0.32

Table 3. Best semantic accuracies at rates less than 10 nats for Thresholded MNIST for each auxiliary decoder and the performance of supervised classifiers trained with each auxiliary target as input. Other data sets omitted for brevity but are similar.

training objective is modified to change the rate-distortion trade-off [2]. It is not immediately obvious whether all points in parameter space with the same ELBO are stable or whether the structure of PixelCNN decoders encourages collapsing into an autodecoder. Stated another way, do all points with the same ELBO form a valley in optimization space or are autoregressive models biased towards autodecoders? To test this, we took Dueling Decoder models at various points during training, disabled the regularization term, and observed the ELBO, rate, and distortion the models converged to. (The converse is less interesting because adding regularization will necessarily pull the model toward a new point in parameter space.) When removing regularization from models during training, we discovered that the models continued to leverage the latent spacce, but each model converged to a different rate-distortion point with similar ELBOs (Figure 4). This is analogous to KL-annealing to $\beta = 1.0$ performing well [10, 12]. (Under KL-annealing, β is slowly increased over time, but Dueling Decoders does not require modifying the objective over time.) These results suggests that models at many points with the same ELBO are stable regardless of where they lie in the rate-distortion plane.
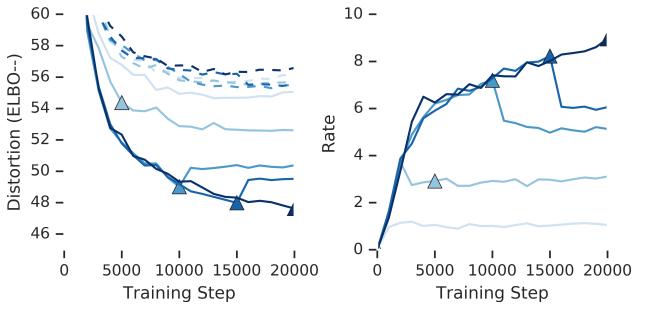
Figure 4. Removing regularization during training produces different convergence points. **Left:** Distortion (solid lines) and ELBO (dashed lines) when removing regularization at 0, 5000, 10000, or 20000 steps (marked by a ▲ in distortion curve). **Right:** Similar curve for the rate.
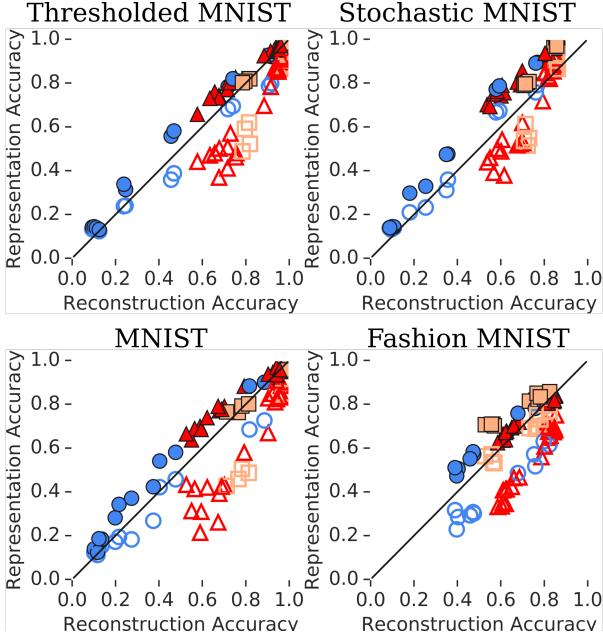
Figure 5. The relationship between the semantic reconstruction accuracy (classifier on X') and the semantic representation accuracy (MLP or linear classifier on Z). The MLP is filled and the linear classifier is empty. Symbols and colors mark the decoder types: PixelCNN (●), CNN (■), and Dueling Decoder (▲). Best viewed in color.

4.6. Latent Structure Explains Semantic Reconstruction Accuracy

The semantic representation accuracy and semantic reconstruction accuracy are different means of estimating the quality of the learned representations. The linear or MLP representation accuracy only incorporates the quality of the encoder, while the reconstruction accuracy depends on the decoder and separate classifier. Even though these are potentially very different, Figure 5 shows that all three measures are highly correlated ($r^2 > 0.8$, Table 4). As the horizontal axis is the semantic reconstruction accuracy and the vertical axis also measures accuracy, the line with unity slope is the equivalent performance to the semantic reconstruction accuracy. Most MLP representation classifiers fall above the line for reconstruction accuracy and most linear representation classifiers fall below the line for reconstruction accuracy. The VAE decoder and subsequent classifier can never increase the amount of information about the input (due to the Data Processing Inequality), so it is useful to estimate how much semantic information is lost. While there is some variability across data subsets, on average, the decoder and classifier maintain 88%±11% of the semantic information available to an MLP in the latent space and the linear classifier only maintains 77%±15% of the information. Therefore the semantic reconstruction accuracy is a good proxy for the amount of information the VAE has extracted and also measures the performance of the decoder,

Dataset	Linear	MLP
Thresholded MNIST	0.904	0.990
Stochastic MNIST	0.901	0.988
MNIST	0.875	0.981
Fashion MNIST	0.804	0.933

Table 4. The r^2 of the relationship between the semantic reconstruction accuracy and the linear or MLP latent decoder accuracy.
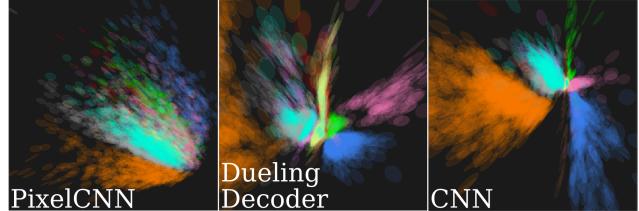
Figure 6. Visualization of the 2 dimensional latent space on Thresholded MNIST for the listed decoder models. Each ellipse is located at the mean of each encoder distribution with axes defined by the covariance matrix. The ellipses are color-coded by the label. Best viewed in color.

which the representation accuracy does not.

To understand why different models have different degrees of semantic reconstruction accuracy, we explored the latent spaces of these models. In Figure 6, we visualize the 2D encoder distributions for each test set image and color coded each with the class label. The PixelCNN models exhibit very little structure in the latent space. In contrast, the CNN and Dueling Decoder models exhibit more clustering in the latent space. Similar visualizations for other data sets and t-SNE embedding visulizations of the best low-rate model regardless of latent space dimensionality are available in the supplemental. Each visualization shows a consistent trend: when the model had high semantic reconstruction or representation accuracy, inputs with the same label are visibly clustered together in the latent space.

5. Conclusions

We find that autoregressive models are capable of using structure in their latent space when it is present. Adding a pixel-independent auxiliary decoder is an effective way to regularize the latent space for image labeling and helps autoregressive decoders escape the local minima that lead to autodecoders. Regularizing the latent space improves the extraction of semantic information over many other objective manipulations that prevent collapse. The rate, distortion, and ELBO are all poor proxies for semantic information which may explain why optimizing re-weighted ELBO objectives does not generally maximize semantic information. Further work will be needed to understand and explore the right auxiliary tasks for VAEs that maximize semantic information.

References

[1] A. Alemi. "Here is a one line MNIST Dataset Loader in Python in a tweet.", 2018. https://twitter.com/alemi/status/1042834067173957632.

[2] A. Alemi, B. Poole, I. Fischer, J. Dillon, R. A. Saurous, and K. Murphy. Fixing a broken elbo. In *International Conference on Machine Learning*, pages 159–168, 2018.

[3] A. A. Alemi and I. Fischer. Therml: Thermodynamics of machine learning, 2018.

[4] A. A. Alemi, I. Fischer, and J. V. Dillon. Uncertainty in the variational information bottleneck. *arXiv preprint arXiv:1807.00906*, 2018.

[5] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

[6] Anonymous. The conditional entropy bottleneck. In *Submitted to International Conference on Learning Representations*, 2019. under review.

[7] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

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

[9] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[10] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio. Generating sentences from a continuous space, 2015.

[11] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner. Understanding disentangling in $beta$-VAE. *arXiv preprint arXiv:1804.03599*, 2018.

[12] X. Chen, D. P. Kingma, T. Salimans, Y. Duan, P. Dhariwal, J. Schulman, I. Sutskever, and P. Abbeel. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.

[13] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus), 2015.

[14] T. R. Davidson, L. Falorsi, N. De Cao, T. Kipf, and J. M. Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018.

[15] A. B. Dieng, Y. Kim, A. M. Rush, and D. M. Blei. Avoiding latent variable collapse with generative skip models. *arXiv preprint arXiv:1807.04863*, 2018.

[16] E. Fertig, A. Arbabi, and A. A. Alemi. β-vaes can retain label information even at high compression. *Bayesian Deep Learning Workshop at NeurIPS 2018*, 2018.

[17] A. Fischer and C. Igel. Training restricted boltzmann machines: An introduction. *Pattern Recognition*, 47(1):25–39, 2014.

[18] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[19] I. Gulrajani, K. Kumar, F. Ahmed, A. A. Taiga, F. Visin, D. Vazquez, and A. Courville. Pixelvae: A latent variable model for natural images. *arXiv preprint arXiv:1611.05013*, 2016.

[20] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning Latent Dynamics for Planning from Pixels. *ArXiv e-prints*, 2018.

[21] J. He, D. Spokoyny, G. Neubig, and T. Berg-Kirkpatrick. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*, 2019.

[22] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[23] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.

[24] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

[25] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.

[26] Y. Kim, S. Wiseman, A. Miller, D. Sontag, and A. Rush. Semi-amortized variational autoencoders. In *International Conference on Machine Learning*, pages 2683–2692, 2018.

[27] D. P. Kingma, S. Mohamed, D. Jimenez Rezende, and M. Welling. Semi-supervised learning with deep generative models. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3581–3589. Curran Associates, Inc., 2014.

[28] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improving variational inference with inverse autoregressive flow, 2016.

[29] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[30] Y. LeCun, C. Cortes, and C. Burges. Mnist handwritten digit database. *AT&T Labs [Online]. Available: http://yann. lecun. com/exdb/mnist*, 2, 2010.

[31] T. Lucas and J. Verbeek. Auxiliary guided autoregressive variational autoencoders, 2018.

[32] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.

[33] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.

[34] R. Salakhutdinov and I. Murray. On the quantitative analysis of deep belief networks. In *Proceedings of the 25th international conference on Machine learning*, pages 872–879. ACM, 2008.

[35] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017.

[36] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal deep variational hand pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[37] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[38] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.

[39] J. M. Tomczak and M. Welling. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*, 2017.

[40] A. van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, et al. Conditional image generation with pixel-cnn decoders. In *Advances in Neural Information Processing Systems*, pages 4790–4798, 2016.

[41] A. van den Oord, O. Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pages 6306–6315, 2017.

[42] R. Vedantam, I. Fischer, J. Huang, and K. Murphy. Generative models of visually grounded imagination. *arXiv preprint arXiv:1705.10762*, 2017.

[43] Y. Wu and K. He. Group normalization. *arXiv preprint arXiv:1803.08494*, 2018.

[44] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

[45] Z. Yang, Z. Hu, R. Salakhutdinov, and T. Berg-Kirkpatrick. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3881–3890. JMLR. org, 2017.

[46] S. Zagoruyko and N. Komodakis. Wide residual networks, 2016.

[47] S. Zhao, J. Song, and S. Ermon. InfoVAE: Information Maximizing Variational Autoencoders. jun 2017.

[48] T. Zhao, R. Zhao, and M. Eskenazi. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, 2017.

Supplemental

5.1. Network Architecture Details

All models used in this paper are convolutional models unless specified as otherwise. The architecture for the encoder models is listed in Supplemental Table 5. The architecture for the pixel-independent convolutional decoder models is listed in Supplemental Table 6. The architecture for the PixelCNN decoder is based off of [35] but slightly modified as in Supplemental Table 7. The auxiliary decoders for intensity histograms or marginal histograms were two fully connected layers of 256 units and ReLU activations feeding into the distribution. The classifier architecture for MNIST based models is the same as the encoder architecture in Supplemental Table 5 but the distribution parameters are replaced with a 10-way softmax classification head. The classifier for the more challenging Fashion MNIST data set is the WideResnet model of [46]. The linear semantic representation decoder is a single layer, fully-connected network connected to a softmax classifier. The MLP semantic representation decoder has two layers of 200 fully-connected units with Exponential Linear Unit activations followed by a softmax classifier [13].

5.2. Hyperparameters for Best Low-Rate Models

The hyperparameters for training the best models with rates<10 for each data set are provided in Supplemental Table 8 in the same order as in the main text.

5.3. Distortion-Rate Trade-Off

As in [2], across the range of models tested models the ELBO is bounded by the limited information content of the data set, but models can trade off between rate and distortion. In Supplemental Figure 7, the models are plotted in the rate-distortion plane and color coded for semantic reconstruction accuracy. As shown in the main text, the accuracy increases rapidly at low rates and is often stable as the rate increases further. For the MNIST and Fashion MNIST data sets, we do not plot the CNN models in Supplemental Figure 7 because their distortions are hundreds of nats worse and visualizing them obscures the rest of the data.

5.4. Latent Space Visualizations

In Supplemental Figure 8, we visualize the 2 dimensional encoder distributions for each test set image and color coded each with the class label for each data set. We can only directly visualize the best models with 2 dimensional latent spaces, so we use t-SNE for the best low-rate models regardless of latent space dimensionality in Supplemental Figure 9. As reported in the main text, each visualization shows a consistent trend: when the model had high semantic reconstruction accuracy, inputs with the same label are

visibly clustered together in the latent space. The color code for each label is provided in Supplemental Table 9.
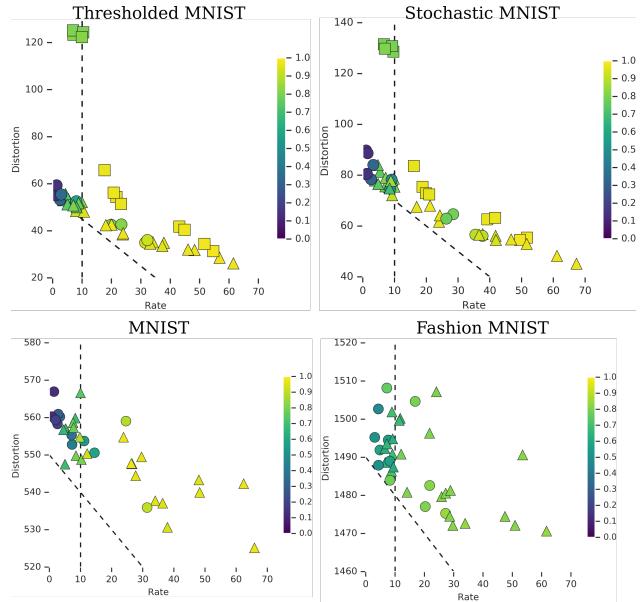
Figure 7. Rate, distortion, and accuracy for models on each dataset. Shape indicates the architecture: PixelCNN (●), CNN (■), and Dueling Decoder (▲). Color represents the semantic reconstruction accuracy. The dashed line at a rate of 10 nats marks the cutoff for the low-rate models. The diagonal line marks a constant ELBO contour. For the MNIST and Fashion MNIST data sets, the CNN decoder models have much higher distortions than are visible in this plot and are not represented here.

Encoder layer	Filters	Shape	Activation	Stride	Padding
1	32	[5, 5]	relu	1	same
2	64	[5, 5]	relu	2	same
3	64	[5, 5]	relu	1	same
4	64	[5, 5]	relu	2	same
5	128	[7, 7]	relu	1	valid
6	distribution parameters	FC	None		

Table 5. Encoder Architecture. The network for the convolutional encoder. The first 5 layers are all convolutional layers. The last layer is a fully-connected (FC) layer to output the mean and lower-triangular covariance of the Gaussian distribution for the latent size of the current model.

Decoder layer	Filters	Shape	Activation	Stride	Padding
1	128	[7, 7]	relu	1	valid
2	128	[5, 5]	relu	1	same
3	128	[5, 5]	relu	2	same
4	64	[5, 5]	relu	1	same
5	64	[5, 5]	relu	2	same
6	32	[5, 5]	relu	1	same
7	distribution parameters	[5, 5]	None	1	same

Table 6. Decoder Architecture. The network for the transposed convolution decoder. The first 5 layers are all transposed convolutional layers. The last layer outputs the probability for the Bernoulli distribution or 5 quantized logistic parameters (center, slope, confidence) for each pixel.

PixelCNN block	Filters	Shape	Activation	Stride	Type
1	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
2	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
3	64 / 128	[2, {3, 2}]	sigmoid * tanh	2	Down, Right
4	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
5	64 / 128	[2, {3, 2}]	sigmoid * tanh	2	Down, Right
6	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
7	64 / 128	[2, {3, 2}]	sigmoid * tanh	2	Deconv: Down, Right
8	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
9	64 / 128	[2, {3, 2}]	sigmoid * tanh	2	Deconv: Down, Right
10	64 / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
	- / 128	[2, {3, 2}]	sigmoid * tanh	1	Down, Right
11	distribution parameters	[1, 1]	None	1	Conv

Table 7. PixelCNN Architecture. The PixelCNN++ network architecture [35]. Each line is block of down and a down-right masked convolution (implemented via shifts). Where the type is listed as "Deconv", the transposed convolution with stride 2 increases the dimensions. Skip connections are applied as in the paper between the lower and upper halves of the network. The encoded latent vector is added as a bias to each layer. The output distribution parameters follow the conventions of Table 6. The un-numbered blocks and larger number of filters are the sizes for the enlarged PixelCNN used in a subset of Fashion MNIST experiments.

Decoder	ELBO	Distortion	Rate	SR Accuracy	β	batch size	latent size	regularization
THRESHOLDED MNIST								
PixelCNN	60.6	52.6 ± 0.2	8.01 ± 0.18	0.47 ± 0.07	0.1	64	2	0.0
CNN	130.4	123.4 ± 0.3	7.04 ± 0.11	0.81 ± 0.02	1.0	64	2	0.0
Dueling Decoder	56.4	48.4 ± 0.2	7.93 ± 0.18	0.88 ± 0.01	1.0	32	16	0.1
STOCHASTIC MNIST								
PixelCNN	83.2	74.6 ± 0.1	8.57 ± 0.16	0.36 ± 0.03	0.1	32	2	0.0
CNN	136.7	129.8 ± 0.3	6.91 ± 0.09	0.73 ± 0.02	1.0	64	2	0.0
Dueling Decoder	81.3	72.1 ± 0.8	9.25 ± 0.72	0.79 ± 0.03	1.0	32	16	0.1
MNIST								
PixelCNN	560.1	552.8 ± 4.5	7.29 ± 0.08	0.37 ± 0.02	0.1	64	2	0.0
CNN	891.6	883.6 ± 73.2	7.94 ± 0.28	0.78 ± 0.01	1.0	64	2	0.0
Dueling Decoder	564.6	554.8 ± 2.7	9.78 ± 0.96	0.79 ± 0.04	1.0	32	16	0.1
FASHION MNIST								
PixelCNN	1492.2	1484.0 ± 1.7	8.24 ± 0.17	0.76 ± 0.02	1.0	64	16	0.0
CNN	1995.3	1986.0 ± 1.3	9.31 ± 0.04	0.56 ± 0.00	1.0	64	2	0.0
Dueling Decoder	1495.1	1486.4 ± 2.0	8.76 ± 1.34	0.69 ± 0.03	1.0	64	64	0.1
PixelCNN+	1473.3	1469.2 ± 3.2	4.08 ± 0.81	0.43 ± 0.12	1.0	32	16	0.0
Dueling Decoder+	1480.2	1471.2 ± 4.0	9.03 ± 0.11	0.67 ± 0.01	1.0	32	2	1.0

Table 8. For each data set and each decoder, we report results for the model with rate<10 and the highest semantic representation accuracy. The reported ELBO is the sum of the distortion and rate. The distortion, rate, and reconstruction accuracy are the means of three randomly initialized training runs $\pm$ one standard deviation. The hyperparameters, β, batch size, number of latent dimensions, and degree of Dueling Decoder regularization, λ, used to train these models are on the right. The decoders with + are enlarged PixelCNN decoders described in Table 7
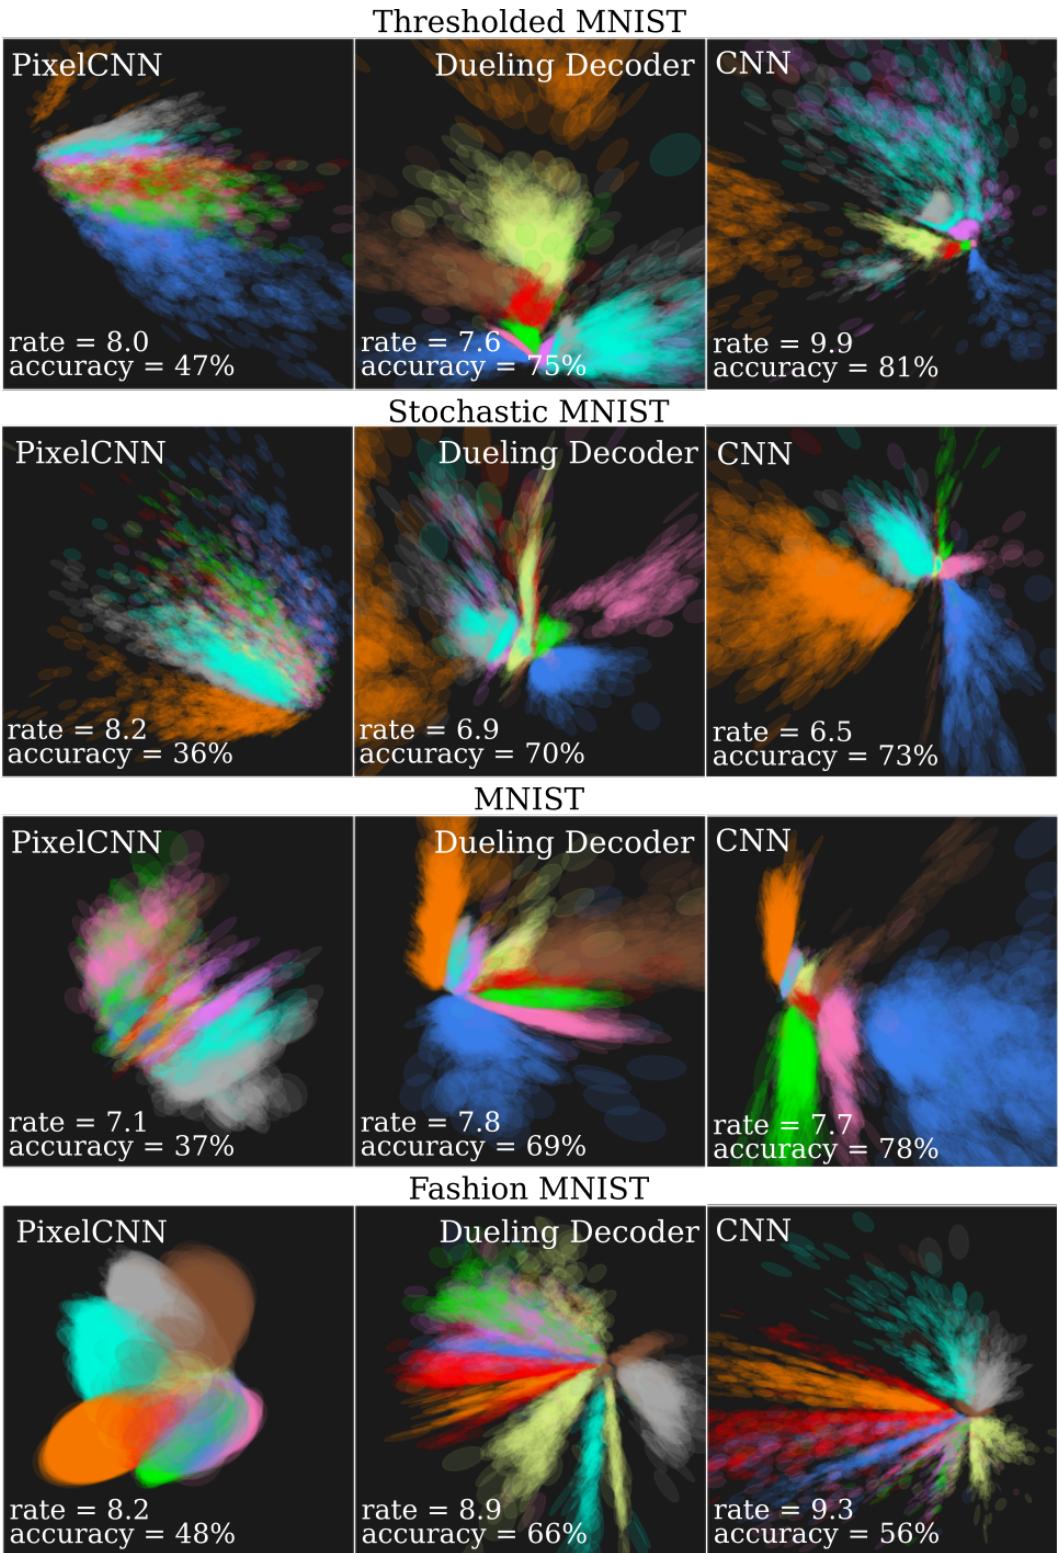
Figure 8. Visualization of the 2 dimensional latent space on each data set for most accurate low-rate decoder models with 2 latent dimensions. Each ellipse is located at the mean of each encoder distribution with axes defined by the covariance matrix. The ellipses are color-coded by the label. The rate and semantic reconstruction accuracy are listed for each model. Best viewed in color.

Thresholded MNIST

PixelCNN

rate = 8.0
accuracy = 47%

Dueling Decoder

rate = 8.3
accuracy = 88%

CNN

rate = 9.9
accuracy = 81%

Stochastic MNIST

PixelCNN

rate = 8.2
accuracy = 36%

Dueling Decoder

rate = 8.9
accuracy = 79%

CNN

rate = 6.5
accuracy = 73%

MNIST

PixelCNN

rate = 7.3
accuracy = 37%

Dueling Decoder

rate = 9.8
accuracy = 79%

CNN

rate = 7.7
accuracy = 78%

Fashion MNIST

PixelCNN

rate = 8.2
accuracy = 76%

Dueling Decoder

rate = 8.6
accuracy = 69%

CNN

rate = 9.3
accuracy = 56%
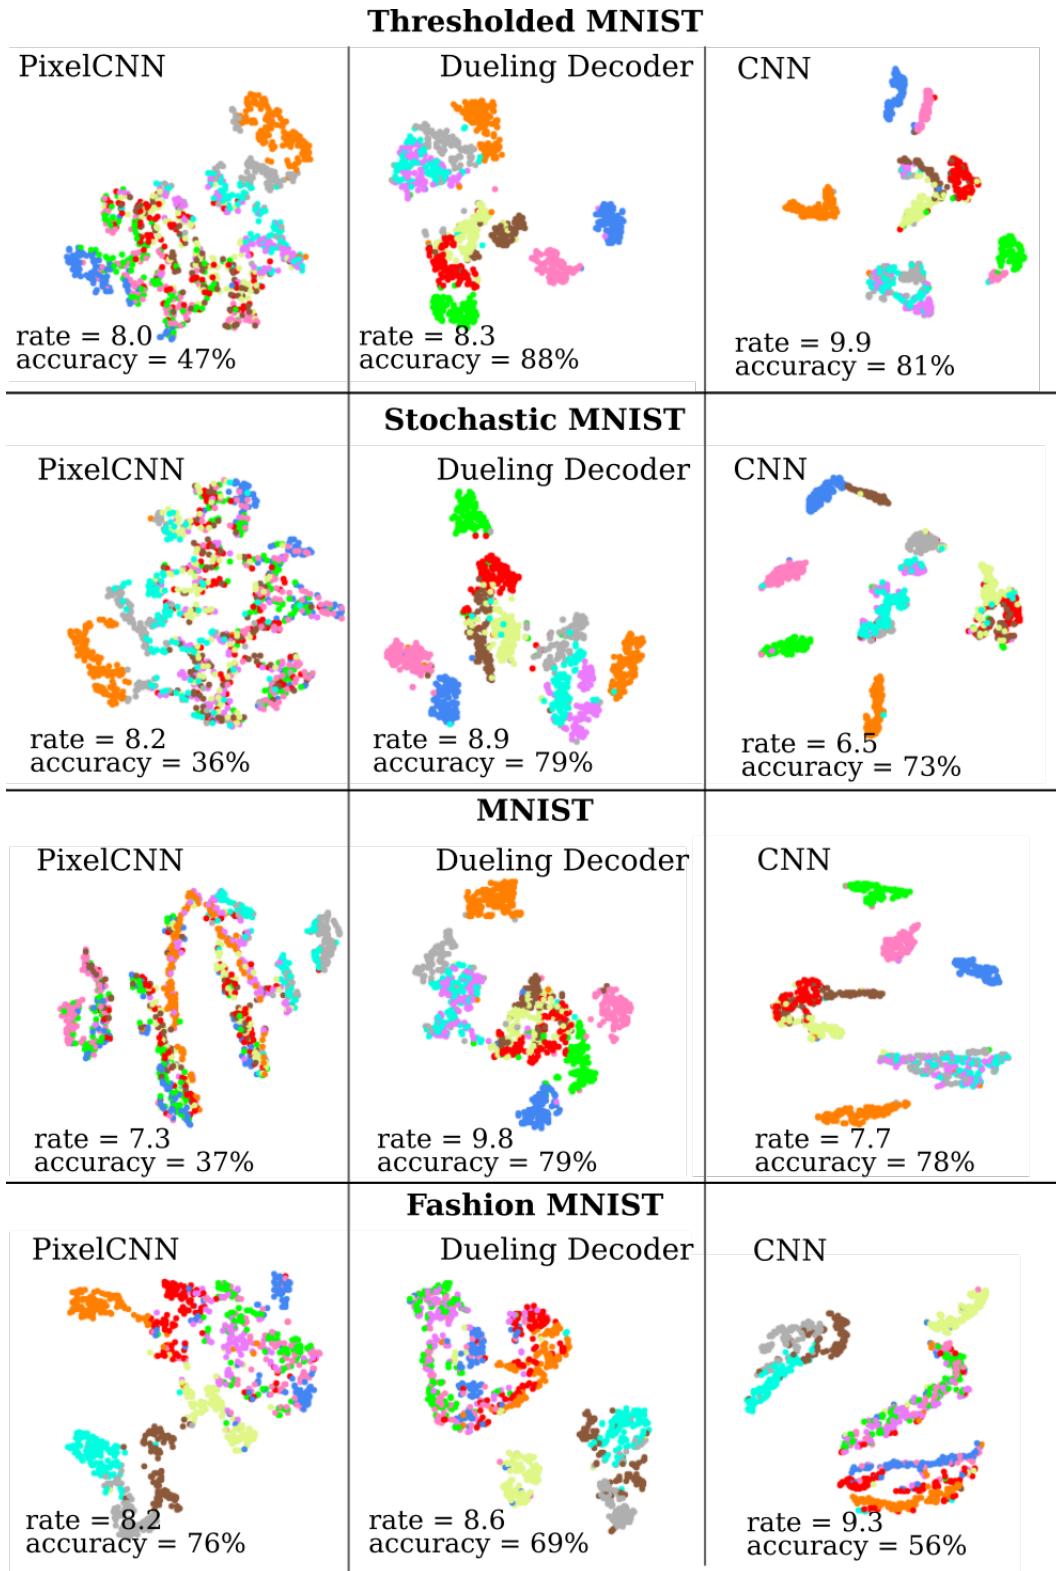
Figure 9. Visualization of the 2 dimensional t-SNE of the latent latent space on each data set for most accurate low-rate decoder models with any number of latent dimensions. Each point is the t-SNE value for the mean of the encoder distribution on the test set, color-coded by the label. The rate and semantic reconstruction accuracy are listed for each model. Best viewed in color.

Color	MNIST	Fashion MNIST
●	0	T-shirt/top
●	1	Trouser
●	2	Pullover
●	3	Dress
●	4	Coat
●	5	Sandal
●	6	Shirt
●	7	Sneaker
●	8	Bag
●	9	Ankle boot

Table 9. The mapping between color and data set labels used in the latent space visualizations. Best viewed in color.