
Inverting Supervised Representations with Autoregressive Neural Density Models

Charlie Nash
School of Informatics
University of Edinburgh, UK
charlie.nash@ed.ac.uk

Nate Kushman
Microsoft Research
Cambridge, UK
nkushman@microsoft.com

Christopher K.I. Williams
School of Informatics
University of Edinburgh, UK
Alan Turing Institute, London, UK
ckiw@inf.ed.ac.uk

Abstract

We present a method for feature interpretation that makes use of recent advances in autoregressive density estimation models to invert model representations. We train generative *inversion models* to express a distribution over input features conditioned on intermediate model representations. Insights into the invariances learned by supervised models can be gained by viewing samples from these inversion models. In addition, we can use these inversion models to estimate the mutual information between a model’s inputs and its intermediate representations, thus quantifying the amount of information preserved by the network at different stages. Using this method we examine the types of information preserved at different layers of convolutional neural networks, and explore the invariances induced by different architectural choices. Finally we show that the mutual information between inputs and network layers initially increases and then decreases over the course of training, supporting recent work by Shwartz-Ziv and Tishby (2017) on the information bottleneck theory of deep learning.

1 Introduction

The representations learned in supervised models are task specific; they discard irrelevant input information and preserve features that are useful for characterizing

their targets. This is the conventional wisdom taken for granted by many in the machine learning community. However the precise nature of what information is preserved across different layers of a neural network is generally unknown. A better understanding of this is desirable both for the interpretability of a particular network, and for the insights that can be gained for neural architecture design.

We expect supervised models to be invariant to certain transformations of the input data. For instance an effective image classifier should be invariant to translations of objects in the image, and such behaviour is encouraged through architectural choices like convolutions and pooling. As such we anticipate that the mapping from inputs to intermediate representations discards information, and that perfect recovery of the inputs is not possible. Recent work by Schwartz-Ziv and Tishby (2017) argues that compression of input data in network representations is a central reason for the success of deep models, particularly with respect to generalization performance. Despite this, attempts such as (Dosovitskiy and Brox, 2016b) have been made to invert representations using reconstructions that are optimized to minimize pixel losses such as mean squared error. This leads to blurry reconstructions from higher-level representations. Perceptual losses using features from a pretrained convolutional network or adversarial discriminator networks significantly improve the visual quality of results (Dosovitskiy and Brox, 2016a; Johnson et al., 2016), but fail to characterize the variability present in the inverse mapping. We propose instead to express a distribution over the inputs conditioned on a network representation. By sampling from this conditional distribution we can visualize the types of inputs that map to a given representation.

In recent years there have been significant advances in neural generative models of high-dimensional data (Goodfellow et al., 2014; Kingma and Welling, 2014;

Rezende et al., 2014). *Autoregressive* density models decompose a joint distribution into products of conditionals. By leveraging domain-specific structure, impressive results have been achieved with neural autoregressive models of images (van den Oord et al., 2016c,b) and audio (van den Oord et al., 2016a). Unlike alternative generative models such as variational autoencoders (Kingma and Welling, 2014; Rezende et al., 2014) or generative adversarial networks (Goodfellow et al., 2014), autoregressive models yield an exact density. We show in Section 4.1 that this density can be used to estimate a lower bound on the mutual information between inputs and model representations, which is a useful metric for the analysis of neural networks. In contrast to other methods for mutual information estimation this estimate is scalable to high-dimensionality inputs and network features with complex dependencies. In addition autoregressive models are straightforward to train in comparison to other generative models, with a single optimization objective and none of the instability associated with adversarial training. Autoregressive density models are therefore a strong choice for our desired goal of representation inversion.

In this work we present a method for the inversion of supervised representations that uses flexible autoregressive neural density models to express a distribution over inputs given an intermediate representation. We show how such models can be used to help understand how much and what kind of information is preserved at different hidden layers on a range of image datasets (Sec. 5.1). We use inversion models to visualise the invariances learned by classifiers with different architectures, and demonstrate advantages in interpretability compared to point-estimate approaches (Sec. 5.2). Finally we demonstrate that the mutual information between inputs and intermediate representations initially increases before decreasing over the course of training, reproducing the results of Schwartz-Ziv and Tishby (2017) in the context of ReLU-convolutional networks (Sec. 5.3).

2 Related work

Our approach is related to many previous works on inverting neural networks. Although our approach has similar goals to optimization-based approaches to network inversion (Linden and Kindermann, 1989; Lee and Kil, 1994; Lu et al., 1999; Mahendran and Vedaldi, 2015) it is most closely related to methods which make use of another neural network that is trained to invert the hidden states (Dosovitskiy and Brox, 2016b,a; Johnson et al., 2016; Huang et al., 2017). In another related work Zeiler and Fergus (2014) invert individual features by explicitly reversing the filtering, pool-

ing and rectification operations in convolutional networks. Our work is distinct in its use of an autoregressive model to express a distribution over a network’s inputs.

Dosovitskiy and Brox (2016b) train an up-sampling convolutional network to map from a representation layer \mathbf{h} to inputs $\hat{\mathbf{x}} = \mathbf{f}(\mathbf{h})$, and optimize the mean squared error with respect to the true inputs $\mathbb{E}_{\mathbf{x}} \|\mathbf{x} - \mathbf{f}(\mathbf{h})\|^2$. With this method reconstructions become increasingly blurry as the amount of information preserved by the network about the inputs decreases with successive layers. The level of blurriness is quite useful as an indication of the amount of information preserved by the network, however it is a coarse measure that doesn’t demonstrate the variability of inputs consistent with a given representation. In addition, this approach is only appropriate for data in continuous spaces where mean squared error is a meaningful metric. Our method can be applied in any setting in which a distribution over inputs can be parameterized, such as language processing.

Johnson et al. (2016) augment a pixel loss with perceptual losses that make use of the feature space of a pre-trained classifier to invert VGG features. These additional constraints result in outputs that are visually appealing in comparison to simple pixel losses. Dosovitskiy and Brox (2016a) extend this approach with an adversarial loss that encourages reconstructed inputs to additionally “fool” a GAN discriminator. Although these approaches produce high quality outputs they are limited in that they produce a single image reconstruction for a given representation, rather than providing a distribution over plausible inputs consistent with the representation.

Our method is also similar to stacked generative adversarial networks (Huang et al., 2017), in which a series of GANs are used to map from higher to lower level representations of a pre-trained classifier. This model was presented primarily as a way to make use of supervised representations in order to improve sample quality. In order to avoid degenerate samples, an entropy loss was incorporated that encourages the auxiliary noise to be recoverable from samples. This loss provides a lower bound on the entropy of the conditional distributions, and results in diverse samples. However entropy maximization is not equivalent to maximizing the likelihood, and may result in poorly calibrated distributions.

Van den Oord et al. (2016b) use a conditional PixelCNN to generate images conditioned on portrait embeddings obtained from the top layer of a face-detection CNN trained with triplet loss on Flickr images. This is equivalent to our method, although in

that case the emphasis was on portrait generation rather than analysis of the learned representations.

3 Background

3.1 Autoregressive neural density models

Neural density models use neural networks to describe parametric distributions p_θ over random variables \mathbf{x} . Autoregressive models decompose the joint distribution into a series of D conditionals $p_\theta(\mathbf{x}) = \prod_i p_\theta(x_i | \mathbf{x}_{1:i-1})$ where the parameters for the i th conditional distribution are the outputs of a network $\theta_i = \mathbf{f}(\mathbf{x}_{1:i-1})$ that takes the preceding variables as input. Density models are typically trained to maximize the likelihood with respect to samples from the true data distribution. Various neural density models have been proposed; from general purpose models (Uria et al., 2013; Germain et al., 2015; Papamakarios et al., 2017) to domain specific models for images (van den Oord et al., 2016c,b), text (Sundermeyer et al., 2012), and audio (van den Oord et al., 2016a). Many neural density models make use of architectures that parallelize the computation of the D conditional distributions through a single pass of a network. This enables efficient computation as well as parameter sharing across conditional distributions. In order to ensure that each conditional only depends on the preceding variables, architectural tools such as causal convolutions and masking are used.

Conditional density models aim to model a conditional distribution $p(\mathbf{x}|\mathbf{h})$ of data variables \mathbf{x} given context \mathbf{h} . Typical examples include models of images conditioned on object classes, or speech models conditioned on speaker identity. Usually each conditional is allowed to depend on the context such that $p_\theta(\mathbf{x}|\mathbf{h}) = \prod_i p_\theta(x_i | \mathbf{x}_{1:i-1}, \mathbf{h})$. We make use of conditional density models in order to model the distribution over input data conditioned on supervised representations.

3.2 PixelCNN

The PixelCNN (van den Oord et al., 2016c) is an autoregressive neural density model for images that uses a convolutional neural network to parameterize conditional distributions for each sub-pixel in an image. Pixel values are sampled one at a time: from left to right and from top to bottom. Causality in the conditional distributions is maintained using masked convolutions that only allow connections from previously observed pixels. The PixelCNN and its variants (van den Oord et al., 2016c,b; Salimans et al., 2017) are powerful models of images, and currently are the state of the art with respect to log-likelihood scores on natural

images. In our experiments we make use of the PixelCNN++ (Salimans et al., 2017), which incorporates a number of changes to the original model including the use of an alternative mixture-based pixel likelihood, downsampling to increase receptive field sizes and short-cut connections. Conditioning information is incorporated by regressing a context vector to biases which are added to intermediate feature maps. For full details see Salimans et al., (2017) and the implementation at <https://github.com/openai/pixel-cnn>.

3.3 Mutual information estimation in neural networks

A quantity of particular interest in the analysis of network representations is the mutual information $I(\mathbf{x}; \mathbf{h})$ between inputs \mathbf{x} and a model representation \mathbf{h} :

$$I(\mathbf{x}; \mathbf{h}) = H(\mathbf{x}) - H(\mathbf{x}|\mathbf{h}) = H(\mathbf{h}) - H(\mathbf{h}|\mathbf{x}), \quad (1)$$

where $H(\mathbf{x})$ denotes the entropy of \mathbf{x} . The mutual information represents the reduction in uncertainty about \mathbf{x} that we obtain if we know \mathbf{h} , and can be thought of in this context as the amount of information preserved in the transformation $\mathbf{x} \rightarrow \mathbf{h}$. In general we are unable to obtain this quantity as we don't have access to the true distributions $p(\mathbf{x})$, $p(\mathbf{h})$, $p(\mathbf{x}|\mathbf{h})$ or $p(\mathbf{h}|\mathbf{x})$. However, the mutual information can be approximated in various ways. In this section we focus on mutual information estimation via density estimation. For more details on alternative approaches see Appendix B. It is important to first consider the implications of working with discrete vs continuous probability distributions.

Discrete vs continuous entropy. For deterministic functions of continuous inputs the conditional distribution $p(\mathbf{h}|\mathbf{x})$ is degenerate and so the differential entropy $H(\mathbf{h}|\mathbf{x})$ is negative infinity. For finite $H(\mathbf{h})$ this implies that the mutual information is infinite. This poses a problem for the analysis of neural networks with continuous inputs, as network representations are typically a deterministic function of the inputs. In this work we deal exclusively with discrete image inputs, and can avoid this issue by using discrete entropy for which $H(\mathbf{h}|\mathbf{x})$ is zero rather than infinite. However, for models with continuous inputs care must be taken to either add noise or to discretize the continuous space. For a more detailed discussion of related issues see Saxe et al. (2018).

Density estimation. For networks operating with continuous input spaces one method for mutual information estimation is to add noise to network activations $\mathbf{h}^\epsilon = \mathbf{h} + \epsilon$ and use a parametric or non-parametric model p^* to estimate $p(\mathbf{h}^\epsilon)$. As \mathbf{h} is a deterministic function of \mathbf{x} the conditional entropy $H(\mathbf{h}|\mathbf{x})$

is simply equal to the entropy of the Gaussian noise $H(\epsilon)$. The approximate model can then be used to estimate the cross entropy $H(p(\mathbf{h}^\epsilon), p^*(\mathbf{h}^\epsilon))$. An upper bound on the mutual information can then be established as follows:

$$I(\mathbf{x}; \mathbf{h}^\epsilon) = H(\mathbf{h}^\epsilon) - H(\epsilon) \leq H(p(\mathbf{h}^\epsilon), p^*(\mathbf{h}^\epsilon)) - c, \quad (2)$$

where $c = H(\epsilon)$, and we use the fact that $H(p(\mathbf{h}^\epsilon), p^*(\mathbf{h}^\epsilon)) \geq H(\mathbf{h}^\epsilon)$. Kolchinsky & Tracey (2017); Kolchinsky et al. (2017) use a kernel density estimate (KDE) of $p(\mathbf{h}^\epsilon)$ to obtain a bound on the mutual information as above. We note that the performance of kernel density estimates deteriorate significantly in higher dimensions. Theis et al. (2016) show that even for a very large number of samples, kernel density methods greatly underestimate the true log-likelihood of simple models trained on 6×6 image patches. As such these methods are not appropriate for large networks. A parametric estimate $p_\theta(\mathbf{h}^\epsilon)$ using e.g. neural autoregressive density models would potentially scale better to higher dimensions, although to our knowledge this hasn't been explored in previous work.

4 Inverting Supervised Representations

Previous approaches to representation inversion (Dosovitskiy and Brox, 2016a) optimize a parameterized inversion function \mathbf{f}_θ with respect to the mean squared error of an image and its reconstruction:

$$\mathcal{L}_\theta^{\text{MSE}} = \mathbb{E}_{p(\mathbf{x}, \mathbf{h})} \|\mathbf{x} - \mathbf{f}_\theta(\mathbf{h})\|^2. \quad (3)$$

This use of a point estimate in order to invert an information-lossy transformation results in blurry reconstructions and does not provide information about the variability inherent in the inverse mapping. Our method instead minimizes the negative log-likelihood of a parameterized *inversion* model:

$$\mathcal{L}_\theta^{\text{NLL}} = -\mathbb{E}_{p(\mathbf{x}, \mathbf{h})} [\log p_\theta(\mathbf{x}|\mathbf{h})]. \quad (4)$$

This enables us to query whether a given input is a good match for a particular representation. In addition we can sample our trained model to get a sense of the degree of constraint present in \mathbf{h} . Optimization of Equation 3 is equivalent to our proposed maximum likelihood criterion if the conditional distribution p_θ is a Gaussian with spherical covariance. As such our method simply extends Equation 3 by using a more flexible class of conditional probability models. As we have chosen to focus on supervised models of images, we use a conditional variant of the PixelCNN++ as our inversion model. Although we do not

explore it here, we note that our method is not tied to any particular density model, and that equivalent domain-appropriate models could be used for e.g. text classification or speech recognition.

4.1 Bounding the mutual information

Inversion models can be used to compute an upper bound on $H(\mathbf{x}|\mathbf{h})$ using the conditional cross entropy between $p(\mathbf{x}|\mathbf{h})$ and the inversion model distribution $p_\theta(\mathbf{x}|\mathbf{h})$:

$$H(p(\mathbf{x}|\mathbf{h}), p_\theta(\mathbf{x}|\mathbf{h})) = -\mathbb{E}_{p(\mathbf{x}, \mathbf{h})} [\log p_\theta(\mathbf{x}|\mathbf{h})] \quad (5)$$

$$= H(\mathbf{x}|\mathbf{h}) + \mathbb{E}_{p(\mathbf{h})} [D_{\text{KL}}[p(\mathbf{x}|\mathbf{h})||p_\theta(\mathbf{x}|\mathbf{h})]] \quad (6)$$

$$\geq H(\mathbf{x}|\mathbf{h}). \quad (7)$$

This enables us to bound the mutual information as follows:

$$I(\mathbf{x}; \mathbf{h}) = H(\mathbf{x}) - H(\mathbf{x}|\mathbf{h}) \quad (8)$$

$$\geq H(\mathbf{x}) - H(p(\mathbf{x}|\mathbf{h}), p_\theta(\mathbf{x}|\mathbf{h})). \quad (9)$$

This is similar to the density estimation approach described in Section 3.3, however instead of estimating $p(\mathbf{h}^\epsilon)$ we are estimating $p(\mathbf{x}|\mathbf{h})$. The gap between the true conditional entropy and the conditional cross entropy is given by the KL divergence between the true conditional distribution $p(\mathbf{x}|\mathbf{h})$ and our approximating distribution $p_\theta(\mathbf{x}|\mathbf{h})$ averaged over \mathbf{h} . Therefore the stronger our density model, the better the approximation to the conditional entropy will be. The PixelCNN++ architecture we use incorporates conditioning information in a large number of convolutional layers, at a range of spatial scales, and is therefore expected to have the expressive capacity to capture dependencies if they exist. This architecture has previously been shown to be effective for learning distributions over images conditioned on object classes, and so we expect it to be similarly effective in this setting.

In practice we use an empirical estimate of the conditional cross entropy by averaging across T , a held-out test set of (\mathbf{x}, \mathbf{h}) pairs, :

$$H^*(\mathbf{x}|\mathbf{h}) = H^*(p(\mathbf{x}|\mathbf{h}), p_\theta(\mathbf{x}|\mathbf{h})) \quad (10)$$

$$= -\frac{1}{|T|} \sum_{(\mathbf{x}, \mathbf{h}) \in T} \log p_\theta(\mathbf{x}|\mathbf{h}) \quad (11)$$

We could take this a step further, and directly estimate the mutual information by using an unconditional model $p_\theta(\mathbf{x})$ to estimate the marginal data distribution $p(\mathbf{x})$ and hence the entropy $H(\mathbf{x})$. However, we don't typically need the absolute value of the mutual information for our analyses, but just the trends in how the mutual information changes between different network layers and settings. Thus, we directly report the estimated negative cross entropy (NCE), $-H^*(\mathbf{x}|\mathbf{h})$, which is equivalent to the mutual information bound, up to the constant $H(\mathbf{x})$.

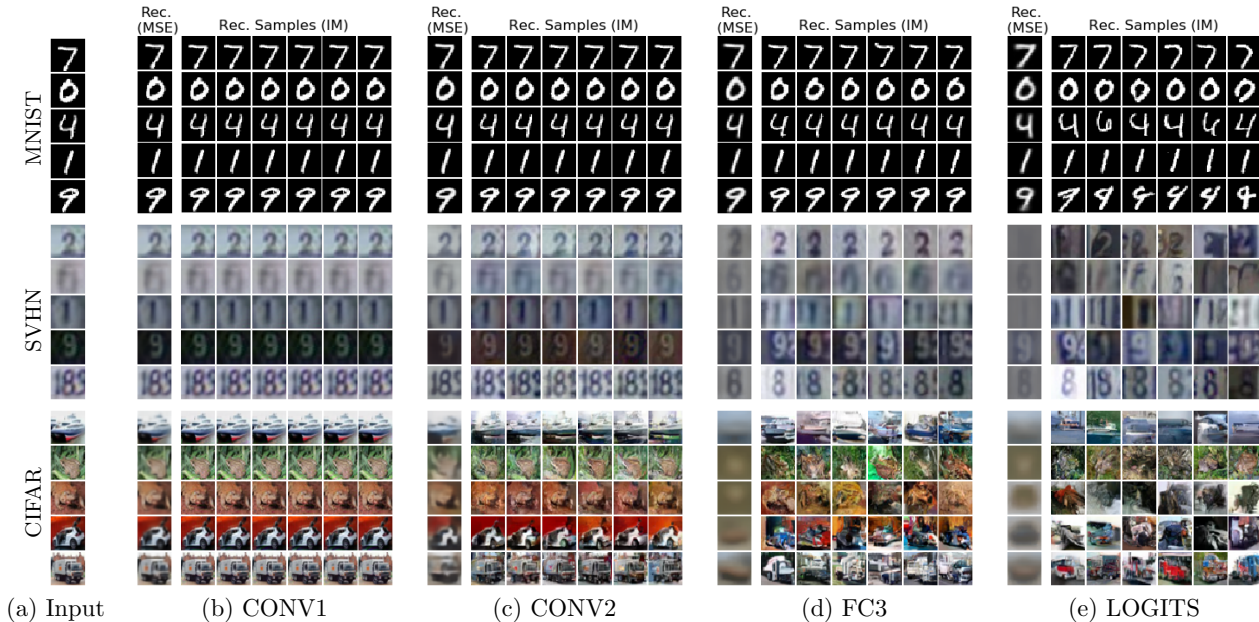


Figure 1: **Inverting network layers:** We compare both inversion model samples and MSE reconstructions across different layers of the neural network. For MNIST, even the logits retain much of the information about the original inputs. For SVHN and CIFAR the CONV1 and CONV2 layers appear to retain most of the information in the original images, while the FC3 layer becomes invariant to many low and mid-level changes.

5 Experiments

5.1 Inverting the layers of image classifiers

We first explore the use of inversion models to explore the invariances and abstractions learned at each layer in a convolutional neural network.

We trained classifiers on three image datasets: MNIST (LeCun et al., 1998), SVHN (Netzer et al., 2011) and CIFAR-10 (Krizhevsky and Hinton, 2009) achieving test accuracies of 99.6%, 93.3% and 81.6% respectively. Following Huang et al. (2017) we used ReLU-convolutional networks consisting of two convolutional layers with max pooling, one fully connected layer, and a linear layer that outputs the predicted logits for each class. We refer to these layers as CONV1, CONV2, FC3 and LOGITS respectively. For each dataset and each classifier layer we trained separate PixelCNN++ inversion models. The architectural and training details are described in appendix A.

Samples. Figure 1 shows samples from trained inversion models along with reconstructions from models trained using the MSE loss from Equation 3. For SVHN and CIFAR the sample variability increases significantly from lower to higher layers. This is consistent with the increasing blurriness of the MSE reconstructions and confirms the expectation that input information is increasingly discarded in successive network layers. For MNIST the reconstructions are visu-

ally similar right up to the output layer, indicating that strong invariance in intermediate layers is not a requirement for good performance on this dataset.

For all datasets, CONV1 samples are almost indistinguishable from the inputs. CONV2 reconstructions also preserve information about the locations, styles and colors of objects and digits, but are more variable with respect to finer details. There is a distinct increase in reconstruction variability from CONV2 to FC3, particularly on the CIFAR dataset. However color information is preserved in FC3, along with object structures and scene textures. A surprising amount of information is retained even in the networks’ logit predictions; this is particularly evident in the MNIST reconstructions for which style and orientation information is preserved. This is consistent with the dark knowledge hypothesis described by Hinton et al. (2015) that suggests that the particular output probabilities that a model assigns to its inputs provides a rich characterization of the similarity between examples.

Mutual information. Figure 5a shows lower bounds on the negative cross entropy at the different layers in the neural network, as discussed in section 4.1. In general, the mutual information bounds decrease through successive layers of the network, indicating a general loss of information about the input. The biggest reduction in mutual information by far is




Figure 2: **Comparing inversion model samples to nearest neighbors:** For each input image we sample 1024 images using an FC3 inversion model, and show the top-k by L_1 distance. We compare these to the top-k nearest neighbors from the training set calculated in the same way.




Figure 3: **SGAN samples conditioned on FC3:** The SGAN samples are less diverse than the inversion model samples and the nearest neighbors shown in Figure 2. In the SGAN model, this diversity is controlled by an entropy loss term with a tunable weight, making it difficult to calibrate it in a principled way.

	MNIST			SVHN			CIFAR		
	CONV1	CONV2	FC3	CONV1	CONV2	FC3	CONV1	CONV2	FC3
1NN	1.25e-2	4.95e-2	1.22e-1	5.35e-2	6.74e-2	2.92e-1	4.11e-2	9.56e-2	5.28e-1
IM-S	7.68e-4	1.65e-2	1.35e-1	1.02e-3	2.80e-2	3.22e-1	4.02e-3	6.43e-2	7.39e-1
IM-NN	6.29e-4	1.40e-2	9.52e-2	9.40e-4	2.43e-2	2.50e-1	3.47e-3	5.73e-2	5.79e-1

Table 1: **Comparing images in representation space:** For 500 test samples from each dataset we compute the average L_1 distance at various network layers between the input image and: (1NN) the nearest neighbor training set example, (IM-S) a random inversion sample or (IM-NN) the closest of 10 inversion model samples. We can see that the IM-NN distance is the closest for all datasets and layers except the FC3 layer of CIFAR10.

between CONV1 and CONV2, which is surprising as the biggest jump in variability for samples from the inversion models appears to occur in later layers. It is possible that our intuitions about visual information are poorly calibrated, and that we underestimate the amount of information present in high-frequency pixel detail. This is supported by the popularity of perceptual loss metrics in generative vision applications (Johnson et al., 2016; Dosovitskiy and Brox, 2016a), that aim to characterize the differences between images in a feature space better aligned with human perception.

Comparison to nearest neighbors and SGAN.

In order to evaluate how well the inversion models capture the distribution $p(\mathbf{x}|\mathbf{h})$ we pass the generated inversion samples, $\hat{\mathbf{x}}$, back through the image classifier to generate the hidden states $\hat{\mathbf{h}}$ for these generated images. Table 1 shows the results of doing this for 500 test set images and computing the average L_1 distance between \mathbf{h} and $\hat{\mathbf{h}}$. We also perform the same

calculation using the nearest neighbor images in the training set, instead of generated samples. We report the distances for a single random sample from an inversion model (IM-S) as well as for the closest of 10 random samples (IM-NN). We find that for a single random sample, the FC3 L_1 distance is smaller than that of the nearest training set example for CONV1 and CONV2, but not for FC3. However the IM-NN distances is smaller for all layers on the MNIST and SVHN datasets, but not on the CIFAR dataset. Figure 2 shows the 6 nearest training examples, and the 6 closest inversion model samples from a collection of 1024, for a selection of input images.

For comparison, in Figure 3 we also show samples from an SGAN conditioned on the FC3 layer of an equivalent CNN (reproduced from (Huang et al., 2017)). We can see that the variability of these samples is much lower than both the samples from our model and from the nearest neighbors in Figure 2c. In fact, in order to prevent the SGAN from collapsing to a deterministic

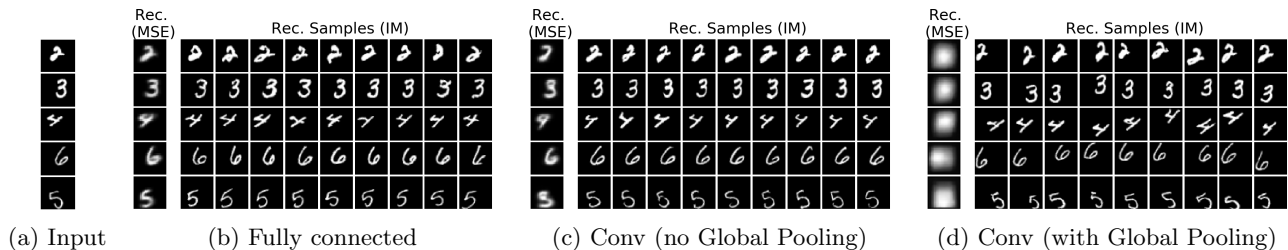


Figure 4: **Architecture comparison:** Inversion model samples and MSE reconstructions from the FC3 layers of three network architectures trained on affNIST. The network with global pooling learns translational invariance, while the others do not. Furthermore, the inversion samples indicate that the style and orientation of the digits are retained in the global pooling network, an observation that cannot be gleaned from the MSE reconstructions.

function, the authors added an explicit entropy term to the loss function. As the weight on this loss term is a tunable hyperparameter, there’s no principled way to set this to ensure that the variability is well calibrated. In contrast, since we are using an autoregressive model which is trained directly to maximize the likelihood of the data, we need no such tunable hyperparameter, and can expect the variance to be better calibrated. It’s also worth noting that our use of the autoregressive model enables estimation of the mutual information, as discussed above, a capability not provided by the SGAN model.

5.2 Comparing network architectures using inversion models

Another practical application of inversion models is to facilitate analysis of network architectures by revealing the invariances present in network layers. If we know that our networks are not learning the desired invariances, we can take steps to modify our architecture. As a case study, we analyze a design choice for convolutional architectures that has become increasingly popular: global spatial pooling. Global pooling layers aggregate information from all spatial locations, greatly reducing the number of parameters in network architectures. They have been found to help reduce overfitting, and have largely replaced fully connected layers in the final processing steps in modern image architectures (Lin et al., 2014; He et al., 2016). As an additional comparison we include a network that replaces convolutional layers with fully connected layers.

We train supervised models on the affNIST dataset¹. affNIST consists of MNIST digits with random affine transformations on a 40×40 canvas. These transformations increase the need for invariance in network representations in comparison to standard MNIST. We trained three supervised networks: the first is identical to that used on MNIST in Section 5.1, the sec-

ond applies global max pooling to the CONV2 feature maps and passes the resulting vector to FC3. The final network replaces convolutional layers with fully connected layers with 2048 units. The network with global pooling performs best, achieving 98.9% accuracy compared to 98.7% for the version without global pooling and 95.0% for the fully connected network. We trained PixelCNN++ inversion models to invert FC3 representations for the supervised networks, using the same architecture as the MNIST model described in Section 5.1.

For the fully connected network we obtain a relative mutual information lower bound of -715.17 nats. For the convolutional networks with and without global pooling we obtain -704.01 and -699.55 nats respectively. This indicates that more input information is preserved in the layers of the convolutional networks than the fully-connected network, and that the global pooling layer discards an estimated ~ 4.5 nats of information about the inputs.

Figure 4 shows samples from the inversion models along with MSE reconstructions. The samples indicate that for the network with global pooling the translation of the digit is not preserved in FC3, whereas for the network without global pooling and the fully connected network the digit’s location is preserved. These results are reinforced by the MSE reconstructions, which are very blurry for the network with global-pooling. It should be noted that it is not possible to tell from the MSE reconstructions that the global-pooling network preserves style and rotation information about the digit, whereas samples from the inversion model indicate that this is the case. It is surprising that the models without global pooling do not achieve translation invariance by FC3, and it may explain their relatively worse performance, as the networks must learn this invariance in the final linear layer.

¹Available at <https://www.cs.toronto.edu/~tijmen/affNIST/>

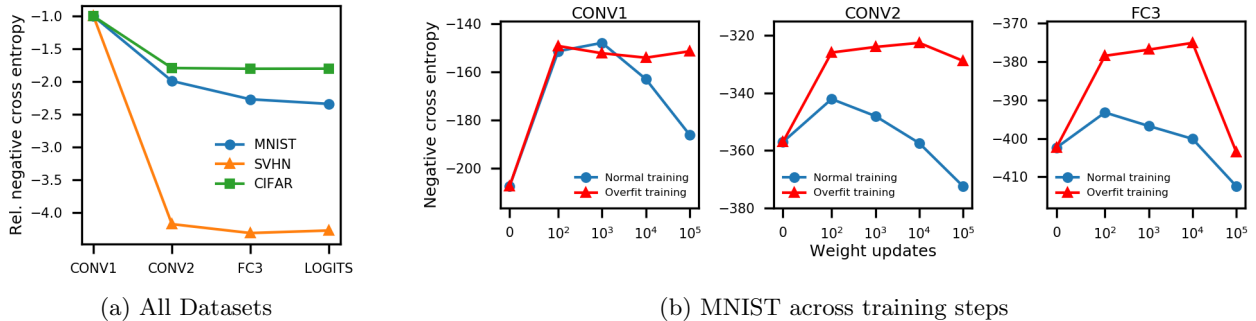


Figure 5: **Mutual information analysis:** (a) Negative Cross Entropy (NCE) across datasets and layers, expressed relative to the magnitude of the CONV1 cross entropy. (b) Absolute NCE in nats for network layers over the course of training for regular (blue) and intentionally overfitted (red) classifiers on MNIST. The initial increase and subsequent decrease in NCE over the course of training for the regular model is consistent with the expansion and compression phases predicted in previous work (Shwartz-Ziv and Tishby, 2017).

5.3 Training dynamics

Inversion models can also be used to better understand the compression dynamics of neural network training. These dynamics were recently studied by Schwartz-Ziv and Tishby (2017) who examined the mutual information between inputs and intermediate layers over the course of SGD optimization. Their findings suggested that there exist two distinct phases of training: an expansion phase in which networks increase the mutual information between inputs and hidden layers and a compression phase in which the mutual information is reduced as information is filtered out. More recently Saxe et al. (2018) provided evidence that the compression phase only occurs when saturating non-linearities such as tanh’s are used, and that no compression is present for ReLU networks. In both cases, the mutual information was estimated using either discretization or non-parametric methods, each of which have issues in terms of scalability to high dimensionality features as discussed in Section 3.3. Instead, we can use inversion models to examine these claims for the larger networks that we employ.

Using the MNIST classifier described in Section 5.1 we trained inversion models on representations established after 0, 10², 10³, 10⁴ and 10⁵ weight updates. We use this relatively coarse view due to the computational expense of training inversion models. For the MNIST network, which takes about 5×10^4 weight updates to converge, this range is fairly representative of the training process. In order to investigate the connections between mutual information and generalization we additionally trained an overfitted network by using only 100 training examples and removing dropout from the classifier. We use equivalent inversion models to the ones described in section 5.1.

Figure 5b shows the negative cross entropy of inver-

sion models trained at different network layers over the course of training. As discussed in Section 4.1 this is equal to a lower bound on the mutual information up to a constant. The results for normal training are shown in blue, and the results for the overfitted network are shown in red. Our main observation is that in the normal training regime for all layers of the network the lower bound on the mutual information initially increases, and then decreases significantly over the course of training. We therefore see a reproduction of the main findings of Schwartz-Ziv and Tishby for ReLU-convolutional networks. This apparent contradiction of the findings of Saxe et al. (2018) can potentially be explained by their use of non-parametric methods to estimate the mutual information. Our results additionally indicate that the mutual information is considerably higher for the overfitted network than for the well-regularized network, which supports the notion that compression in network layers has an important role in a model’s generalization performance.

6 Discussion

We present a method for the inversion of supervised representations that uses flexible autoregressive neural density models to express a distribution over inputs given an intermediate representation. Our method has two benefits: it facilitates visualisation of model invariances, thus enabling analysis of architectural choices. Secondly it provides a scalable quantitative estimate of the amount of information preserved by a network. One difficulty is that density estimation is challenging in higher dimensions, and that we don’t know how well a given model represents the true density. However, as neural density models improve, so too does our method, and we will be able to achieve tighter mutual information bounds and more representative samples.

Acknowledgements

The work of Christopher Williams is supported by EP-SRC grant EP/N510129/1 to the Alan Turing Institute. Charlie Nash is supported by the Centre for Doctoral Training in Data Science, funded by EPSRC grant EP/L016427/1 and by a Microsoft Azure Research Award.

References

- Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. *Neural Information Processing Systems*, 2016a.
- Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. *Computer Vision and Pattern Recognition*, 2016b.
- Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: masked autoencoder for distribution estimation. *International Conference on Machine Learning*, 2015.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Neural information processing systems*, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Computer Vision and Pattern Recognition*, 2016.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *Deep Learning and Representation Learning Workshop, Neural Information Processing Systems*, 2015.
- Xun Huang, Yixuan Li, Omid Poursaeed, John E. Hopcroft, and Serge J. Belongie. Stacked generative adversarial networks. *Computer Vision and Pattern Recognition*, 2017.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. *European Conference on Computer Vision*, 2016.
- Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. *International Conference on Learning Representations*, 2014.
- Artemy Kolchinsky and Brendan D. Tracey. Estimating mixture entropy with pairwise distances. *Entropy*, 19(7):361, 2017.
- Artemy Kolchinsky, Brendan D. Tracey, and David H. Wolpert. Nonlinear information bottleneck. *CoRR*, abs/1705.02436, 2017.
- Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- 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.
- Sukhan Lee and Rhee Man Kil. Inverse mapping of continuous functions using local and global information. *IEEE Trans. Neural Networks*, 5(3):409–423, 1994.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *International conference on learning representations*, 2014.
- A. Linden and J. Kindermann. Inversion of multilayer nets. *International Joint Conference on Neural Networks*, 1989.
- Bao-Liang Lu, Hajime Kita, and Yoshikazu Nishikawa. Inverting feedforward neural networks using linear and nonlinear programming. *IEEE Trans. Neural Networks*, 10(6):1271–1290, 1999.
- Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. *Computer Vision and Pattern Recognition*, 2015.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *Deep Learning and Unsupervised Feature Learning Workshop, Neural Information Processing Systems*, 2011.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Neural Information Processing Systems*, 2017.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *International Conference on Machine Learning*, 2014.
- Tim Salimans and Diederik P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Neural Information Processing Systems*, 2016.
- Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. PixelCNN++: Improving the PixelCNN with discretized logistic mixture likelihood and other modifications. *International Conference on Learning Representations*, 2017.

- Andrew Michael Saxe, Yamini Bansal, Joel Dapello, Madhu Advani, Artemy Kolchinsky, Brendan Daniel Tracey, and David Daniel Cox. On the information bottleneck theory of deep learning. *International Conference on Learning Representations*, 2018.
- Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *CoRR*, abs/1703.00810, 2017.
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. LSTM neural networks for language modeling. *Conference of the International Speech Communication Association*, 2012.
- L. Theis, Aaron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *International Conference on Learning Representations*, 2016.
- Benigno Uria, Iain Murray, and Hugo Larochelle. RNADE: the real-valued neural autoregressive density-estimator. *Advances in Neural information processing systems*, 2013.
- 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. *ISCA Speech Synthesis Workshop*, 2016a.
- Aäron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Koray Kavukcuoglu, Oriol Vinyals, and Alex Graves. Conditional image generation with pixelcnn decoders. *Neural Information Processing Systems*, 2016b.
- Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *International Conference on Machine Learning*, 2016c.
- Matthew D. Zeiler and Rob” Fergus. Visualizing and understanding convolutional networks. *European Conference on Computer Vision*, 2014.