

SEMANTIC NOISE MODELING FOR BETTER REPRESENTATION LEARNING

Hyo-Eun Kim* and Sangheum Hwang

Lunit Inc.
Seoul, South Korea
{hekim, shwang}@lunit.io

Kyunghyun Cho

Courant Institute of Mathematical Sciences and Centre for Data Science
New York University
New York, NY 10012, USA
kyunghyun.cho@nyu.edu

ABSTRACT

Latent representation learned from multi-layered neural networks via hierarchical feature abstraction enables recent success of deep learning. Under the deep learning framework, generalization performance highly depends on the learned latent representation. In this work, we propose a novel latent space modeling method to learn better latent representation. We designed a neural network model based on the assumption that good base representation for supervised tasks can be attained by maximizing the sum of hierarchical mutual informations between the input, latent, and output variables. From this base model, we introduce a *semantic noise modeling* method which enables semantic perturbation on the latent space to enhance the representational power of learned latent feature. During training, latent vector representation can be stochastically perturbed by a modeled additive noise while preserving its original semantics. It implicitly brings the effect of semantic augmentation on the latent space. The proposed model can be easily learned by back-propagation with common gradient-based optimization algorithms. Experimental results show that the proposed method helps to achieve performance benefits against various previous approaches. We also provide the empirical analyses for the proposed latent space modeling method including t-SNE visualization.

1 INTRODUCTION

Enhancing the generalization performance against unseen data given some sample data is the main objective in machine learning. Under that point of view, deep learning has been achieved many breakthroughs in several domains such as computer vision (Krizhevsky et al., 2012; Simonyan & Zisserman, 2015; He et al., 2016), natural language processing (Collobert & Weston, 2008; Bahdanau et al., 2015), and speech recognition (Hinton et al., 2012; Graves et al., 2013). Deep learning is basically realized on deep layered neural network architecture, and it learns appropriate task-specific latent representation based on given training data. Better latent representation learned from training data results in better generalization over the future unseen data. Representation learning or latent space modeling becomes one of the key research topics in deep learning. During the past decade, researchers focused on unsupervised representation learning and achieved several remarkable landmarks on deep learning history (Vincent et al., 2010; Hinton et al., 2006; Salakhutdinov & Hinton, 2009). In terms of utilizing good base features for supervised learning, the base representation learned from unsupervised learning can be a good solution for supervised tasks (Bengio et al., 2007; Masci et al., 2011).

The definition of ‘good’ representation is, however, different according to target tasks. In unsupervised learning, a model is learned from unlabelled examples. Its main objective is to build a model

*Corresponding author

to estimate true data distribution given examples available for training, so the learned latent representation normally includes broadly-informative components of the raw input data (e.g., mutual information between the input and the latent variable can be maximized for this objective). In supervised learning, however, a model is learned from labelled examples. In the case of classification, a supervised model learns to discriminate input data in terms of the target task using corresponding labels. Latent representation is therefore obtained to maximize the performance on the target supervised tasks.

Since the meaning of good representations vary according to target tasks (unsupervised or supervised), pre-trained features from the unsupervised model are not be guaranteed to be useful for subsequent supervised tasks. Instead of the two stage learning strategy (unsupervised pre-training followed by supervised fine-tuning), several works focused on a joint learning model which optimizes unsupervised and supervised objectives concurrently, resulting in better generalization performance (Goodfellow et al., 2013; Larochelle & Bengio, 2008a; Rasmus et al., 2015; Zhao et al., 2015; Zhang et al., 2016; Cho & Chen, 2014).

In this work, we propose a novel latent space modeling method for supervised learning as an extension of the joint learning approach. We define a good latent representation of standard feed-forward neural networks under the basis of information theory. Then, we introduce a *semantic noise modeling* method in order to enhance the generalization performance. The proposed method stochastically perturbs the latent representation of a training example by injecting a modeled semantic additive noise. Since the additive noise is randomly sampled from a pre-defined probability distribution every training iteration, different latent vectors from a single training example can be fully utilized during training. The multiple different latent vectors produced from a single training example are semantically similar under the proposed latent space modeling method, so we can expect semantic augmentation effect on the latent space.

Experiments are performed on two datasets; MNIST and CIFAR-10. The proposed model results in better classification performance compared to previous approaches through notable generalization effect (stochastically perturbed training examples well cover the distribution of unseen data).

2 METHODOLOGY

The proposed method starts from the existing joint learning viewpoint. This section first explains the process of obtaining a good base representation for supervised learning which is the basis of the proposed latent space modeling method. And then, we will describe how the proposed semantic noise modeling method perturbs the latent space while maintaining the original semantics.

2.1 BASE JOINT LEARNING MODEL

In a traditional feed-forward neural network model (Figure 1(a)), output Y of input data X is compared with its true label, and the error is propagated backward from top to bottom, which implicitly learns a task-specific latent representation Z of the input X . As an extension of a joint learning approach, an objective to be optimized can be described in general as below (Larochelle & Bengio, 2008b):

$$\min_{\theta} \lambda L_{unsup} + L_{sup} \quad (1)$$

where L_{unsup} and L_{sup} are respectively an unsupervised loss and a supervised loss, and θ and λ are model parameters to be optimized during training and a loss weighting coefficient, respectively. In terms of modeling L_{unsup} in Eq. (1), we assume that good latent representation Z is attained by maximizing the sum of hierarchical mutual informations between the input, latent, and output variables; i.e. the sum of the mutual information between the input X and the Z and the mutual information between the Z and the output Y . Each mutual information is decomposed into an entropy and a conditional entropy terms, so the sum of hierarchical mutual informations is expressed as follows:

$$\mathbf{I}(X; Z) + \mathbf{I}(Z; Y) = \mathbf{H}(X) - \mathbf{H}(X|Z) + \mathbf{H}(Z) - \mathbf{H}(Z|Y) \quad (2)$$




Figure 1: (a) Standard feed-forward neural network model, (b) feed-forward neural network model with reconstruction paths, and (c) feed-forward neural network model with reconstruction and stochastic perturbation paths.

where $\mathbf{I}(\cdot; \cdot)$ is the mutual information between random variables, and $\mathbf{H}(\cdot)$ and $\mathbf{H}(\cdot|\cdot)$ are the entropy and the conditional entropy of random variables, respectively. Note that the sum of those mutual informations becomes equivalent to the total correlation of X , Z , and Y under the graphical structure of the general feed-forward model described in Figure 1(a); $P(X, Z, Y) = P(Y|Z)P(Z|X)P(X)$. The total correlation is equal to the sum of all pairwise mutual informations (Watanabe, 1960).

Our objective is to find the model parameters which maximize $\mathbf{I}(X; Z) + \mathbf{I}(Z; Y)$. Since $\mathbf{H}(X)$ and $\mathbf{H}(Z)$ are non-negative, and $\mathbf{H}(X)$ is constant in this case, the lower bound on $\mathbf{I}(X; Z) + \mathbf{I}(Z; Y)$ can be reduced to ¹:

$$\mathbf{I}(X; Z) + \mathbf{I}(Z; Y) \geq -\mathbf{H}(X|Z) - \mathbf{H}(Z|Y) . \quad (3)$$

It is known that maximizing $-\mathbf{H}(X|Z)$ can be formulated as minimizing the reconstruction error between the input $x^{(i)}$ (i -th example sampled from X) and its reconstruction $x_R^{(i)}$ under the general auto-encoder framework (Vincent et al., 2010). Since $\mathbf{H}(X|Z) + \mathbf{H}(Z|Y)$ is proportional to the sum of reconstruction errors of $x^{(i)}$ (with its reconstruction $x_R^{(i)}$) and $z^{(i)}$ (with its reconstruction $z_R^{(i)}$), the target objective can be expressed as follows (refer to Appendix (A1) for the details of mathematical derivations):

$$\min_{\theta} \sum_i L_{rec}(x^{(i)}, x_R^{(i)}) + L_{rec}(z^{(i)}, z_R^{(i)}) \quad (4)$$

where L_{rec} is a reconstruction loss.

Figure 1(b) shows the target model obtained from the assumption that good latent representation Z can be obtained by maximizing the sum of hierarchical mutual informations. Given an input sample x , feed-forward vectors and their reconstructions are attained deterministically by:

$$\begin{aligned} z &= f_{\theta_1}(x) \\ y &= f_{\theta_2}(f_{\theta_1}(x)) \\ x_R &= g_{\theta'_1}(z) = g_{\theta'_1}(f_{\theta_1}(x)) \\ z_R &= g_{\theta'_2}(y) = g_{\theta'_2}(f_{\theta_2}(f_{\theta_1}(x))) . \end{aligned} \quad (5)$$

¹Although $\mathbf{H}(Z)$ is an upper bound of $\mathbf{H}(Z|Y)$, $\mathbf{H}(Z)$ is anyway affected by the process of $\mathbf{H}(Z|Y)$ being minimized in Eq. (3). In Section 4, we experimentally show that we can obtain good base model even from the relatively loose lower bound defined in Eq. (3).

Given a set of training pairs $(x^{(i)}, t^{(i)})$ where $x^{(i)}$ and $t^{(i)}$ are the i -th input example and its label, target objective in Eq. (1) under the model described in Figure 1(b) can be organized as below (with real-valued input samples, L2 loss L_{L2} is a proper choice for the reconstruction loss L_{rec}):

$$\min_{\theta: \{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \sum_i \lambda \left(L_{L2}(x^{(i)}, x_R^{(i)}) + L_{L2}(z^{(i)}, z_R^{(i)}) \right) + L_{NLL}(y^{(i)}, t^{(i)}) \quad (6)$$

where L_{NLL} is a negative log-likelihood loss for the target supervised task. Note that Eq. (6) represents the ‘*proposed-base*’ in our experiment (see Section 4.3).

2.2 SEMANTIC NOISE MODELING

Based on the architecture shown in Figure 1(b) with the target objective in Eq. (6), we conjecture that stochastic perturbation on the latent space during training helps to achieve better generalization performance for supervised tasks. Figure 1(c) shows this strategy which integrates the stochastic perturbation process during training. Suppose that Z_P is a perturbed version of Z , and Y_P is an output which is feed-forwarded from Z_P . Given a latent vector $z = f_{\theta_1}(x)$ from an input sample x ,

$$z' = z + z_e \text{ and } \hat{y} = f_{\theta_2}(z') \quad (7)$$

where z' and \hat{y} are a perturbed latent vector and its output respectively, and z_e is an additive noise used in the perturbation process of z . Based on the architecture shown in Figure 1(c), target objective can be modified as:

$$\min_{\theta: \{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \sum_i \lambda_1 \left(L_{L2}(x^{(i)}, x_R^{(i)}) + L_{L2}(z^{(i)}, z_R^{(i)}) \right) + \lambda_2 L_{NLL}(y^{(i)}, t^{(i)}) + L_{NLL}(\hat{y}^{(i)}, t^{(i)}) . \quad (8)$$

Using random additive noise directly on z_e is the most intuitive approach (‘*proposed-perturb (random)*’ in Section 4.3). However, preserving the semantics of the original latent representation z cannot be guaranteed under the direct random perturbation on the latent space. While the latent space is not directly interpretable in general, the output logit y of the latent representation z is interpretable, because the output logit is tightly coupled to the prediction of the target label. In order to preserve the semantics of the original latent representation after perturbation, we indirectly model a semantic noise on the latent space by adding small random noise directly on the output space.

Based on the output (pre-softmax) logit y , the semantic-preserving variation of y (i.e. y') can be modeled by $y' = y + y_e$, where y_e is a random noise vector stochastically sampled from a zero-mean *Gaussian* with small standard deviation σ ; $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Now, the semantic perturbation z' can be reconstructed from the random perturbation y' through the decoding path $g_{\theta'_2}$ in Figure 1(c). From the original output logit y and the randomly perturbed output logit y' , semantic additive noise z_e on the latent space can be approximately modeled as below:

$$\begin{aligned} z_R &= g_{\theta'_2}(y) \\ z'_R &= g_{\theta'_2}(y') = g_{\theta'_2}(y + y_e) \\ z_e &\simeq z'_R - z_R = g_{\theta'_2}(y + y_e) - g_{\theta'_2}(y) \end{aligned} \quad (9)$$

By using the modeled semantic additive noise z_e and the original latent representation z , we can obtain the semantic perturbation z' as well as its output \hat{y} via Eq. (7) for our target objective Eq. (8).

From the described *semantic noise modeling* process (‘*proposed-perturb (semantic)*’ in Section 4.3), we expect to achieve better representation on the latent space. The effect of the proposed model in terms of learned latent representation will be explained in more detail in Section 4.4.




Figure 2: Previous works for supervised learning; (a) traditional feed-forward model, and (b) joint learning model with both supervised and unsupervised losses.

3 RELATED WORKS

Previous works on deep neural networks for supervised learning can be categorized into two types as shown in Figure 2; (a) a general feed-forward neural network model (LeCun et al., 1998; Krizhevsky et al., 2012; Simonyan & Zisserman, 2015; He et al., 2016), and (b) a joint learning model which optimizes unsupervised and supervised objectives at the same time (Zhao et al., 2015; Zhang et al., 2016; Cho & Chen, 2014). Here are the corresponding objective functions:

$$\min_{\theta:\{\theta_1,\theta_2\}} \sum_i L_{NLL}(y^{(i)}, t^{(i)}) \tag{10}$$

$$\min_{\theta:\{\theta_1,\theta'_1,\theta_2\}} \sum_i \lambda L_{L2}(x^{(i)}, x_R^{(i)}) + L_{NLL}(y^{(i)}, t^{(i)}) \tag{11}$$

where λ is a loss weighting coefficient between unsupervised and supervised losses.

Since the feed-forward neural network model is normally implemented with multiple layers in a deep learning framework, the joint learning model can be sub-classified into two types according to the type of reconstruction; reconstruction only with the input data x (Eq. (11)) and reconstruction with all the intermediate features including the input data x as follows:

$$\min_{\theta} \sum_i \left(\lambda_0 L_{L2}(x^{(i)}, x_R^{(i)}) + \sum_j \lambda_j L_{L2}(h_j^{(i)}, h_{jR}^{(i)}) + L_{NLL}(y^{(i)}, t^{(i)}) \right). \tag{12}$$

where $h_j^{(i)}$ and $h_{jR}^{(i)}$ are the j -th hidden representation of the i -th training example and its reconstruction.

Another type of the joint learning model, a *ladder network* (Figure 3), was introduced for semi-supervised learning (Rasmus et al., 2015). The key concept of the ladder network is to obtain robust features by learning de-noising functions ($g_{\theta'}$) of the representations at every layer of the model via reconstruction losses, and the supervised loss is combined with the reconstruction losses in order to build the semi-supervised model. The ladder network achieved the best performance in semi-supervised tasks, but it is not appropriate for supervised tasks with small-scale training set (experimental analysis for supervised learning on permutation-invariant MNIST is briefly summarized




Figure 3: Ladder network; a representative model for semi-supervised learning (Rasmus et al., 2015).

in Appendix (A2)). The proposed model in this work can be extended to semi-supervised learning, but our main focus is to enhance the representational power on latent space given labelled data for supervised learning. We leave the study for semi-supervised learning scenario based on the proposed methodology as our future research.

4 EXPERIMENTS

For quantitative analysis, we compare the proposed methodology with previous approaches described in Section 3; a traditional feed-forward supervised learning model and a joint learning model with two different types of reconstruction losses (reconstruction only with the first layer or with all the intermediate layers including the first layer). The proposed methodology includes a baseline model in Figure 1(b) as well as a stochastic perturbation model in Figure 1(c). Especially in the stochastic perturbation model, we compare the random and semantic perturbations and present some qualitative analysis on the meaning of the proposed perturbation methodology.

4.1 DATASETS

We experiment with two public datasets; MNIST (including a permutation-invariant MNIST case) and CIFAR-10. MNIST (10 classes) consists of 50k, 10k, and 10k 28×28 gray-scale images for training, validation, and test datasets, respectively. CIFAR-10 (10 classes) consists of 50k and 10k 32×32 3-channel images for training and test sets, respectively. We split the 50k CIFAR-10 training images into 40k and 10k for training and validation. Experiments are performed with different sizes of training set (from 10 examples per class to the entire training set) in order to verify the effectiveness of the proposed model in terms of generalization performance under varying sizes of training set.

4.2 IMPLEMENTATION

Figure 4 shows the architecture of the neural network model used in this experiment. W 's are convolution or fully-connected weights (biases are excluded for visual brevity). Three convolution (3×3 (2) 32, 3×3 (2) 64, 3×3 (2) 96, where each item means the filter kernel size and (stride) with the number of filters) and two fully-connected (the numbers of output nodes are 128 and 10, respectively) layers are used for MNIST. For the permutation-invariant MNIST setting, 784-512-256-256-128-10 nodes of fully-connected layers are used. Four convolution (5×5 (1) 64, 3×3 (2) 64, 3×3 (2) 64, and 3×3 (2) 96) and three fully-connected (128, 128, and 10 nodes) layers are used for CIFAR-10. Weights on the decoding (reconstruction) path are tied with corresponding weights on the encoding path as shown in Figure 4 (transposed convolution for the tied convolution layer and transposed matrix multiplication for the tied fully-connected layer).

In Figure 4, z' is perturbed directly from z by adding *Gaussian* random noise for random perturbation. For semantic perturbation, z' is indirectly generated from y' which is perturbed by adding *Gaussian* random noise on y based on Eq. (9). For perturbation, base activation vector (z is the base




Figure 4: Target network architecture; 3 convolution and 2 fully-connected layers were used for MNIST, 5 fully-connected layers were used for permutation-invariant MNIST, and 4 convolution and 3 fully-connected layers were used for CIFAR-10.

Table 1: Error rate (%) on the test set using the model with the best performance on the validation set. Numbers on the first row of each sub-table are the number of randomly chosen per-class training examples. The average performance and the standard deviation of three different random-split datasets (except for the case using the entire training set in the last column) are described in this table (error rate on each random set is summarized in Appendix (A3)). Performance of three previous approaches (with gray background; previous-1, 2, 3 are feed-forward model Figure 2(a), joint learning model with recon-one Figure 2(b), joint learning model with recon-all Figure 2(b), respectively) and the proposed methods (proposed-1, 2, 3 are baseline Figure 1(b), random perturbation Figure 1(c), semantic perturbation Figure 1(c), respectively) is summarized.

dataset	number of per-class examples chosen from 50k entire MNIST training examples								entire set
MNIST	10	20	50	100	200	500	1k	2k	50k
previous-1	24.55 (3.04)	16.00 (1.33)	10.35 (0.66)	6.58 (0.42)	4.71 (0.28)	2.94 (0.23)	1.90 (0.27)	1.45 (0.08)	1.04
previous-2	21.67 (3.19)	13.60 (0.99)	7.85 (0.10)	5.44 (0.37)	4.14 (0.08)	2.50 (0.15)	1.84 (0.07)	1.45 (0.07)	1.12
previous-3	20.11 (2.81)	13.69 (0.62)	9.15 (0.15)	6.77 (0.25)	5.39 (0.11)	3.89 (0.27)	2.91 (0.17)	2.28 (0.10)	1.87
proposed-1	21.35 (1.16)	11.65 (1.15)	6.33 (0.10)	4.32 (0.31)	3.07 (0.11)	1.98 (0.11)	1.29 (0.09)	0.94 (0.02)	0.80
proposed-2	20.17 (1.52)	11.68 (0.81)	6.24 (0.29)	4.12 (0.24)	3.04 (0.13)	1.88 (0.05)	1.24 (0.03)	0.96 (0.08)	0.65
proposed-3	20.11 (0.81)	10.59 (0.74)	5.92 (0.12)	3.79 (0.23)	2.72 (0.09)	1.78 (0.05)	1.15 (0.01)	0.88 (0.03)	0.62

dataset	number of per-class examples chosen from 40k entire CIFAR-10 training examples								entire set
CIFAR-10	10	20	50	100	200	500	1k	2k	40k
previous-1	73.82 (1.43)	68.99 (0.54)	61.30 (0.83)	54.93 (0.56)	46.97 (0.59)	33.69 (0.43)	26.63 (0.39)	20.97 (0.09)	17.80
previous-2	75.68 (1.56)	69.05 (1.13)	61.44 (0.63)	55.02 (0.34)	46.18 (0.51)	33.62 (0.38)	26.78 (0.48)	21.25 (0.40)	17.68
previous-3	73.33 (1.06)	67.63 (0.56)	62.59 (0.76)	56.37 (0.20)	50.51 (0.61)	41.26 (0.73)	32.55 (1.20)	26.38 (0.08)	22.71
proposed-1	71.63 (0.69)	66.17 (0.40)	58.91 (0.86)	52.65 (0.28)	43.46 (0.30)	31.86 (0.54)	25.76 (0.31)	21.06 (0.18)	17.45
proposed-2	71.69 (0.25)	66.75 (0.54)	58.95 (0.63)	53.01 (0.26)	43.71 (0.19)	31.80 (0.18)	25.50 (0.33)	20.81 (0.27)	17.43
proposed-3	71.50 (1.14)	66.87 (0.17)	58.30 (0.62)	52.32 (0.08)	42.98 (0.34)	30.91 (0.23)	24.81 (0.26)	20.19 (0.25)	16.16

vector for the random perturbation and y is the base vector for the semantic perturbation) is scaled to $[0.0, 1.0]$, and the zero-mean *Gaussian* noise with 0.2 of standard deviation is added (via element-wise addition) on the normalized base activation. This perturbed scaled activation is de-scaled with the original min and max activations of the base vector.

Initial learning rates are 0.005 and 0.001 for MNIST and permutation-invariant MNIST, and 0.002 for CIFAR-10, respectively. The learning rates are decayed by a factor of 5 every 40 epochs until the 120-th epoch. For both datasets, the minibatch size is set to 100, and the target objective is optimized using *Adam* optimizer (Kingma & Ba, 2015) with a momentum 0.9. All the λ 's for reconstruction losses in Eq. (11) and Eq. (12) are 0.03 and 0.01 for MNIST and CIFAR-10, respectively. The same weighting factors for reconstruction losses (0.03 for MNIST and 0.01 for CIFAR-10) are used for λ_1 in Eq (8), and 1.0 is used for λ_2 .

Input data is first scaled to $[0.0, 1.0]$ and then whitened by the average across all the training examples. In CIFAR-10, random cropping (24×24 image is randomly cropped from the original 32×32 image) and random horizontal flipping (mirroring) are used for data augmentation. We selected the network that performed best on the validation dataset for evaluation on the test dataset. All the experiments are performed with TensorFlow (Abadi et al., 2015).

4.3 QUANTITATIVE ANALYSIS

Three previous approaches (a traditional feed-forward model, a joint learning model with the input reconstruction loss, and a joint learning model with reconstruction losses of all the intermediate layers including the input layer) are compared with the proposed methods (the baseline model in Figure 1(b), and the stochastic perturbation model in Figure 1(c) with two different perturbation methods; random and semantic). We measure the classification performance according to varying sizes of training set (examples randomly chosen from the original training dataset). Performance is averaged over three different random trials.




Figure 5: Examples reconstructed from the perturbed latent vectors via (a) random perturbation, and (b) semantic perturbation (top row shows the original training examples). More examples are summarized in Appendix (A4.1).

Table 1 summarizes the classification performance for MNIST and CIFAR-10. As we expected, the base model obtained by maximizing the sum of mutual informations (*proposed-base*) mostly performs better than previous approaches, and the model with the semantic perturbation (*proposed-perturb (semantic)*) performs best among all the comparison targets. Especially in MNIST, the error rate of ‘*proposed-perturb (semantic)*’ with 2k per-class training examples is less than the error rate of all types of previous works with the entire training set (approximately 5k per-class examples).

We further verify the proposed method on the permutation-invariant MNIST task with a standard feed-forward neural network. Classification performance is measured against three different sizes of training set (1k, 2k, and 5k per-class training examples). ‘*Proposed-perturb (semantic)*’ achieves the best performance among all the configurations; 2.57%, 1.82%, and 1.28% error rates for 1k, 2k, and 5k per-class training examples, respectively. The joint learning model with the input reconstruction loss performs best among three previous approaches; 2.72%, 1.97%, and 1.38% error rates for 1k, 2k, and 5k per-class training examples, respectively.

4.4 QUALITATIVE ANALYSIS

As mentioned before, random perturbation by adding unstructured noise directly to the latent representation cannot guarantee preserving the semantics of the original representation. We compared two different perturbation methods (random and semantic) by visualizing the examples reconstructed from the perturbed latent vectors (Figure 5). Top row is the original examples selected from training set (among 2k per-class training examples), and the rest are the reconstructions of their perturbed latent representations. Based on the architecture described in Figure 1(b), we generated five different perturbed latent representations according to the type of perturbation, and reconstructed the perturbed latent vectors through decoding path for reconstruction.

Figure 5(a) and (b) show the examples reconstructed from the random and semantic perturbations, respectively. For both cases, zero-mean *Gaussian* random noise (0.2 standard deviation) is used for perturbation. As shown in Figure 5(a), random perturbation partially destroys the original semantics; for example, semantics of ‘1’ is mostly destroyed under random perturbation, and some examples of ‘3’ are reconstructed as being similar to ‘8’ rather than its original content ‘3’. Figure 5(b) shows the examples reconstructed from the semantic perturbation. The reconstructed examples show subtle semantic variations while preserving the original semantic contents; for example, thickness difference in ‘3’ (example on the third row) or writing style difference in ‘8’ (openness of the top left corner).

Figure 6 shows the overall effect of the perturbation. In this analysis, 100 per-class MNIST examples are used for training. From the trained model based on the architecture described in Figure 1(b), latent representations z of all the 50k examples (among 50k examples, only 1k examples were used for training) are visualized by using t-SNE (Maaten & Hinton, 2008). Only the training examples of three classes (0, 1, and 9) among ten classes are depicted as black circles for visual discrimination in




Figure 6: Training examples (circles or crosses with colors described below) over the examples not used for training (depicted as background with different colors); (a) training examples (black circles), (b) training examples (yellow circles) with $3\times$ random-perturbed samples (blue crosses), and (c) training examples (yellow circles) with $3\times$ semantic-perturbed samples (blue crosses). Best viewed in color.

Figure 6(a). The rest of the examples which were not used for training (approximately 4.9k examples per class) are depicted as a background with different colors. We treat the colored background examples (not used for training) as a true distribution of unseen data in order to estimate the generalization level of learned representation according to the type of perturbation. Figure 6(b) and (c) show the training examples (100 examples per class with yellow circles) and their perturbed ones ($3\times$ sampled from each example with blue crosses) through random and semantic perturbations, respectively.

In Figure 6(b), perturbed samples are distributed near the original training examples, but some samples outside the true distribution cannot be identified easily with appropriate classes. This can be explained with Figure 5(a), since some perturbed samples are ambiguous semantically. In Figure 6(c), however, most of the perturbed samples evenly cover the true distribution. As mentioned before, stochastic perturbation with the semantic additive noise during training implicitly incurs the effect of augmentation on the latent space while resulting in better generalization. Per-class t-SNE results are summarized in Appendix (A4.2).

5 DISCUSSION

We introduced a novel latent space modeling method for supervised tasks based on the standard feed-forward neural network architecture. The presented model simultaneously optimizes both supervised and unsupervised losses based on the assumption that the better latent representation can be obtained by maximizing the sum of hierarchical mutual informations. Especially the stochastic perturbation process which is achieved by modeling the semantic additive noise during training enhances the representational power of the latent space. From the proposed *semantic noise modeling* process, we can expect improvement of generalization performance in supervised learning with implicit semantic augmentation effect on the latent space.

The presented model architecture can be intuitively extended to semi-supervised learning because it is implemented as the joint optimization of supervised and unsupervised objectives. For semi-supervised learning, however, logical link between features learned from labelled and unlabelled data needs to be considered additionally. We leave the extension of the presented approach to semi-supervised learning for the future.

REFERENCES

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vin-

- cent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <http://tensorflow.org/>. Software available from tensorflow.org.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015.
- Yoshua Bengio, Pascal Lamblin, Dan Popovici, Hugo Larochelle, et al. Greedy layer-wise training of deep networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- Kyunghyun Cho and Xi Chen. Classifying and visualizing motion capture sequences using deep neural networks. In *International Conference on Computer Vision Theory and Applications*, 2014.
- Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning (ICML)*, 2008.
- Ian Goodfellow, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Multi-prediction deep boltzmann machines. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *International conference on acoustics, speech and signal processing*, 2013.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine, IEEE*, 29(6):82–97, 2012.
- Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Hugo Larochelle and Yoshua Bengio. Classification using discriminative restricted boltzmann machines. In *International Conference on Machine Learning (ICML)*, 2008a.
- Hugo Larochelle and Yoshua Bengio. Classification using discriminative restricted boltzmann machines. In *International Conference on Machine Learning (ICML)*, 2008b.
- 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.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research (JMLR)*, 9(Nov):2579–2605, 2008.
- Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. Stacked convolutional auto-encoders for hierarchical feature extraction. In *International Conference on Artificial Neural Networks*, 2011.
- Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Ruslan Salakhutdinov and Geoffrey E Hinton. Deep boltzmann machines. In *Artificial Intelligence and Statistics Conference (AISTATS)*, 2009.

- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research (JMLR)*, 11:3371–3408, 2010.
- Satosi Watanabe. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82, 1960.
- Yuting Zhang, Kibok Lee, and Honglak Lee. Augmenting supervised neural networks with unsupervised objectives for large-scale image classification. In *International Conference on Machine Learning (ICML)*, 2016.
- Junbo Zhao, Michael Mathieu, Ross Goroshin, and Yann Lecun. Stacked what-where auto-encoders. In *International Conference on Learning Representations (ICLR)*, 2015.

APPENDIX

(A1) DERIVATION OF RECONSTRUCTION ERRORS FROM CONDITIONAL ENTROPY TERMS

Extended from Section 2. From the lower bound in Eq. (3), we consider the following optimization problem (refer to ‘Section 2. From mutual information to autoencoders’ in (Vincent et al., 2010)):

$$\max_{\{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \mathbb{E}_{q(X, Z, Y)} [\log q(X|Z)] + \mathbb{E}_{q(X, Z, Y)} [\log q(Z|Y)] . \quad (13)$$

Here, we denote $q(X, Z, Y)$ an unknown joint distribution. Note that Z and Y are respectively the variables transformed from parametric mappings $Z = f_{\theta_1}(X)$ and $Y = f_{\theta_2}(Z)$ (see Fig. 1). $q(X, Z, Y)$ then can be reduced to $q(X)$ from $q(Z|X; \theta_1) = \delta(Z - f_{\theta_1}(X))$ and $q(Y|Z; \theta_2) = \delta(Y - f_{\theta_2}(Z))$ where δ denotes Dirac-delta function.

From the Kullback-Leibler divergence that $D_{\text{KL}}(q||p) \geq 0$ for any two distributions p and q , the optimization in Eq. (13) corresponds to the following optimization problem where $p(\cdot)$ denotes a parametric distribution:

$$\max_{\{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \mathbb{E}_{q(X)} [\log p(X|Z; \theta'_1)] + \mathbb{E}_{q(X)} [\log p(Z|Y; \theta'_2)] . \quad (14)$$

By replacing $q(X)$ with a sample distribution $q^0(X)$ and putting all parametric dependencies between X , Z and Y , we will have

$$\max_{\{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \mathbb{E}_{q^0(X)} [\log p(X|Z = f_{\theta_1}(X); \theta'_1)] + \mathbb{E}_{q^0(X)} [\log p(Z|Y = f_{\theta_2}(f_{\theta_1}(X)); \theta'_2)] . \quad (15)$$

For a given input sample x of X , it is general to interpret x_R and z_R as the parameters of distributions $p(X|X_R = x_R)$ and $p(Z|Z_R = z_R)$ which reconstruct x and z with high probability (i.e. x_R and z_R are not exact reconstructions of x and z). Since x_R and z_R are real-valued, we assume Gaussian distribution for these conditional distributions, that is,

$$\begin{aligned} p(X|X_R = x_R) &= \mathcal{N}(x_R, \sigma_0^2 \mathbf{I}) \\ p(Z|Z_R = z_R) &= \mathcal{N}(z_R, \sigma_0^2 \mathbf{I}) . \end{aligned} \quad (16)$$

The assumptions yield $-\log p(\cdot|\cdot) \propto L_{L2}(\cdot, \cdot)$.

With the following relations for log terms in Eq. (15),

$$\begin{aligned} p(X|Z = f_{\theta_1}(x); \theta'_1) &= p(X|X_R = g_{\theta'_1}(f_{\theta_1}(x))) \\ p(Z|Y = f_{\theta_2}(f_{\theta_1}(x)); \theta'_2) &= p(Z|Z_R = g_{\theta'_2}(f_{\theta_2}(f_{\theta_1}(x)))) , \end{aligned} \quad (17)$$

the optimization problem in Eq. (15) corresponds to the minimization problem of reconstruction errors for input examples $x^{(i)}$ as below:

$$\min_{\{\theta_1, \theta'_1, \theta_2, \theta'_2\}} \sum_i L_{L2}(x^{(i)}, x_R^{(i)}) + L_{L2}(z^{(i)}, z_R^{(i)}) . \quad (18)$$

(A2) LADDER NETWORK, A REPRESENTATIVE SEMI-SUPERVISED LEARNING MODEL

Extended from Section 3. We performed experiments with a ladder network model (Rasmus et al., 2015) in order to estimate the performance on pure supervised tasks according to different sizes of training set. We used the code (<https://github.com/rinuboney/ladder.git>) for this experiment. The network architecture implemented on the source code is used as is; (784-1000-500-250-250-10). Based on the same network architecture, we implemented the proposed stochastic perturbation model described in Figure 1(c) and compared the classification performance with the ladder network as described in Table 2 (we did not focus on searching the optimal hyperparameters for the proposed model in this experiment). As summarized in the bottom of the table (mean over 3 random trials), the proposed semantic noise modeling method shows a fairly large performance gain compared to the ladder network model with small-scale datasets (e.g., in a case of 10 per-class training examples, the proposed method achieves 22.11% of error rate, while the ladder network shows 29.66%).

Table 2: Classification performance (error rate in %) of the ladder network and the proposed model on three different sets of randomly chosen training examples (MNIST).

set No.1 (# training examples per class)	10	20	50	100	200	500	1k	2k	(all) 5k
ladder network model; Figure 3	25.85	16.48	9.26	6.00	4.66	3.07	2.15	1.26	0.91
proposed-perturb (semantic); Figure 1(c)	19.76	12.33	8.77	6.06	4.59	2.93	1.87	1.31	0.93
set No.2 (# training examples per class)	10	20	50	100	200	500	1k	2k	
ladder network model; Figure 3	33.14	17.46	10.44	6.67	4.43	2.82	1.94	1.37	
proposed-perturb (semantic); Figure 1(c)	23.36	15.35	9.43	5.75	4.43	2.99	1.87	1.39	
set No.3 (# training examples per class)	10	20	50	100	200	500	1k	2k	
ladder network model; Figure 3	29.99	16.99	9.73	7.34	4.39	3.00	2.12	1.47	
proposed-perturb (semantic); Figure 1(c)	23.21	13.98	8.83	6.51	4.32	2.94	2.22	1.49	
mean over 3 random trials	10	20	50	100	200	500	1k	2k	(all) 5k
ladder network model; Figure 3	29.66	16.98	9.81	6.67	4.49	2.96	2.07	1.37	0.91
proposed-perturb (semantic); Figure 1(c)	22.11	13.89	9.01	6.11	4.45	2.95	1.99	1.40	0.93

(A3) QUANTITATIVE ANALYSIS

Extended from Section 4.3. Among the total 50k and 40k training examples in MNIST and CIFAR-10, we randomly select the examples for training. Classification performance according to three different randomly chosen training sets are summarized in Table 3 (MNIST) and Table 4 (CIFAR-10). Further experiments with denoising constraints are also included. Zero-mean *Gaussian* random noise with 0.1 standard deviation is used for noise injection. Denoising function helps to achieve slightly better performance on MNIST, but it results in performance degradation on CIFAR-10 (we did not focus on searching the optimal parameters for noise injection in this experiments).

Table 3: Classification performance (error rate in %) on three different sets of randomly chosen training examples (MNIST).

Set No.1 (# train examples per class)	10	20	50	100	200	500	1k	2k	(all) 5k
feed-forward model; Figure 2(a)	22.61	14.20	11.25	6.37	4.34	2.63	1.83	1.56	1.04
joint learning model with recon-one; Figure 2(b)	18.69	12.21	7.84	5.17	4.02	2.58	1.79	1.47	1.12
joint learning model with recon-one with denoising constraints	20.39	11.91	7.41	4.64	3.65	2.57	1.97	1.53	0.97
joint learning model with recon-all; Figure 2(b)	18.82	12.82	9.34	6.43	5.23	4.12	2.68	2.42	1.87
joint learning model with recon-all with denoising constraints	17.93	11.76	7.32	4.78	3.91	3.04	2.52	1.99	1.36
proposed-base; Figure 1(b)	20.23	10.18	6.47	3.89	3.04	1.89	1.33	0.91	0.80
proposed-base with denoising constraints	19.88	10.89	6.62	4.26	3.40	2.44	2.11	1.54	1.13
proposed-perturb (random); Figure 1(c)	18.38	10.58	6.64	3.78	3.14	1.90	1.21	0.89	0.65
proposed-perturb (semantic); Figure 1(c)	19.33	9.72	5.98	3.47	2.84	1.84	1.16	0.84	0.62
Set No.2 (# train examples per class)	10	20	50	100	200	500	1k	2k	
feed-forward model; Figure 2(a)	28.84	17.36	10.14	6.20	4.78	3.02	1.61	1.41	
joint learning model with recon-one; Figure 2(b)	26.09	14.40	7.98	5.18	4.17	2.29	1.94	1.52	
joint learning model with recon-one with denoising constraints	27.69	13.11	6.95	5.07	3.54	2.37	1.83	1.28	
joint learning model with recon-all; Figure 2(b)	24.01	14.13	8.98	6.84	5.44	3.51	2.98	2.18	
joint learning model with recon-all with denoising constraints	23.05	13.29	7.79	5.12	3.92	3.01	2.27	1.84	
proposed-base; Figure 1(b)	22.95	12.98	6.27	4.43	3.22	2.14	1.37	0.96	
proposed-base with denoising constraints	26.96	12.21	6.45	4.62	3.13	2.53	1.88	1.49	
proposed-perturb (random); Figure 1(c)	22.10	12.52	5.97	4.26	2.86	1.94	1.23	0.92	
proposed-perturb (semantic); Figure 1(c)	21.22	11.52	5.75	3.91	2.61	1.73	1.14	0.89	
Set No.3 (# train examples per class)	10	20	50	100	200	500	1k	2k	
feed-forward model; Figure 2(a)	22.20	16.43	9.67	7.16	5.02	3.17	2.25	1.39	
joint learning model with recon-one; Figure 2(b)	20.23	14.19	7.73	5.96	4.22	2.62	1.79	1.35	
joint learning model with recon-one with denoising constraints	19.32	12.25	7.44	5.39	3.58	2.37	1.49	1.56	
joint learning model with recon-all; Figure 2(b)	17.51	14.12	9.12	7.04	5.49	4.05	3.08	2.25	
joint learning model with recon-all with denoising constraints	17.07	12.50	7.86	5.48	4.05	2.97	2.02	1.98	
proposed-base; Figure 1(b)	20.86	11.79	6.25	4.63	2.96	1.91	1.16	0.96	
proposed-base with denoising constraints	19.89	11.30	6.26	4.57	3.50	2.63	1.61	1.47	
proposed-perturb (random); Figure 1(c)	20.02	11.94	6.12	4.32	3.13	1.81	1.28	1.08	
proposed-perturb (semantic); Figure 1(c)	19.78	10.53	6.03	4.00	2.70	1.76	1.14	0.92	

Table 4: Classification performance (error rate in %) on three different sets of randomly chosen training examples (CIFAR-10).

Set No.1 (# train examples per class)	10	20	50	100	200	500	1k	2k	(all) 4k
feed-forward model; Figure 2(a)	73.30	69.25	62.42	55.65	47.71	34.30	27.04	21.06	17.80
joint learning model with recon-one; Figure 2(b)	75.19	70.38	62.25	55.30	46.89	34.12	26.63	21.05	17.68
joint learning model with recon-one with denoising constraints	73.72	68.20	61.99	55.23	46.64	36.37	29.78	25.53	21.73
joint learning model with recon-all; Figure 2(b)	74.79	68.33	62.92	56.24	51.37	40.30	30.91	26.49	22.71
joint learning model with recon-all with denoising constraints	76.56	69.67	64.53	57.88	52.74	42.24	36.90	30.93	27.41
proposed-base; Figure 1(b)	70.79	66.57	59.91	52.98	43.29	32.25	26.19	20.92	17.45
proposed-base with denoising constraints	71.03	67.49	60.37	53.52	44.28	33.40	28.00	25.06	21.34
proposed-perturb (random); Figure 1(c)	71.89	67.12	59.22	52.79	43.87	31.82	25.04	20.97	17.43
proposed-perturb (semantic); Figure 1(c)	71.59	66.90	58.64	52.34	42.74	30.94	24.45	20.10	16.16
Set No.2 (# train examples per class)	10	20	50	100	200	500	1k	2k	
feed-forward model; Figure 2(a)	72.39	69.49	60.45	54.85	46.91	33.39	26.73	21.00	
joint learning model with recon-one; Figure 2(b)	74.06	69.14	60.71	54.54	45.70	33.54	27.43	20.90	
joint learning model with recon-one with denoising constraints	76.40	69.33	60.28	55.38	47.40	36.29	29.31	24.60	
joint learning model with recon-all; Figure 2(b)	72.28	67.60	61.53	56.65	49.99	42.08	32.99	26.33	
joint learning model with recon-all with denoising constraints	73.90	69.23	61.90	57.99	52.35	45.12	37.23	30.14	
proposed-base; Figure 1(b)	72.49	65.62	57.82	52.66	43.20	32.24	25.60	21.32	
proposed-base with denoising constraints	72.99	66.75	57.78	53.81	44.33	33.56	28.40	25.03	
proposed-perturb (random); Figure 1(c)	71.84	65.98	58.08	53.37	43.44	31.56	25.69	21.03	
proposed-perturb (semantic); Figure 1(c)	72.85	66.65	57.44	52.21	42.74	31.17	24.99	20.54	
Set No.3 (# train examples per class)	10	20	50	100	200	500	1k	2k	
feed-forward model; Figure 2(a)	75.78	68.24	61.02	54.29	46.28	33.38	26.11	20.85	
joint learning model with recon-one; Figure 2(b)	77.79	67.62	61.37	55.22	45.96	33.21	26.29	21.81	
joint learning model with recon-one with denoising constraints	76.60	69.27	61.13	55.10	47.50	37.12	29.63	24.88	
joint learning model with recon-all; Figure 2(b)	72.92	66.97	63.31	56.23	50.16	41.41	33.75	26.31	
joint learning model with recon-all with denoising constraints	76.83	68.53	65.58	58.29	52.43	45.42	39.01	32.32	
proposed-base; Figure 1(b)	71.60	66.31	58.99	52.30	43.88	31.10	25.48	20.95	
proposed-base with denoising constraints	72.39	67.20	60.60	52.64	44.62	33.52	28.01	25.25	
proposed-perturb (random); Figure 1(c)	71.34	67.15	59.55	52.86	43.81	32.01	25.78	20.42	
proposed-perturb (semantic); Figure 1(c)	70.06	67.07	58.83	52.41	43.47	30.61	25.00	19.94	

(A4.1) QUALITATIVE ANALYSIS

Extended from Section 4.4. Figure 7 shows reconstructed examples from perturbed (random or semantic) latent representations (refer to Figure 5 and the analysis described in Section 4.4).

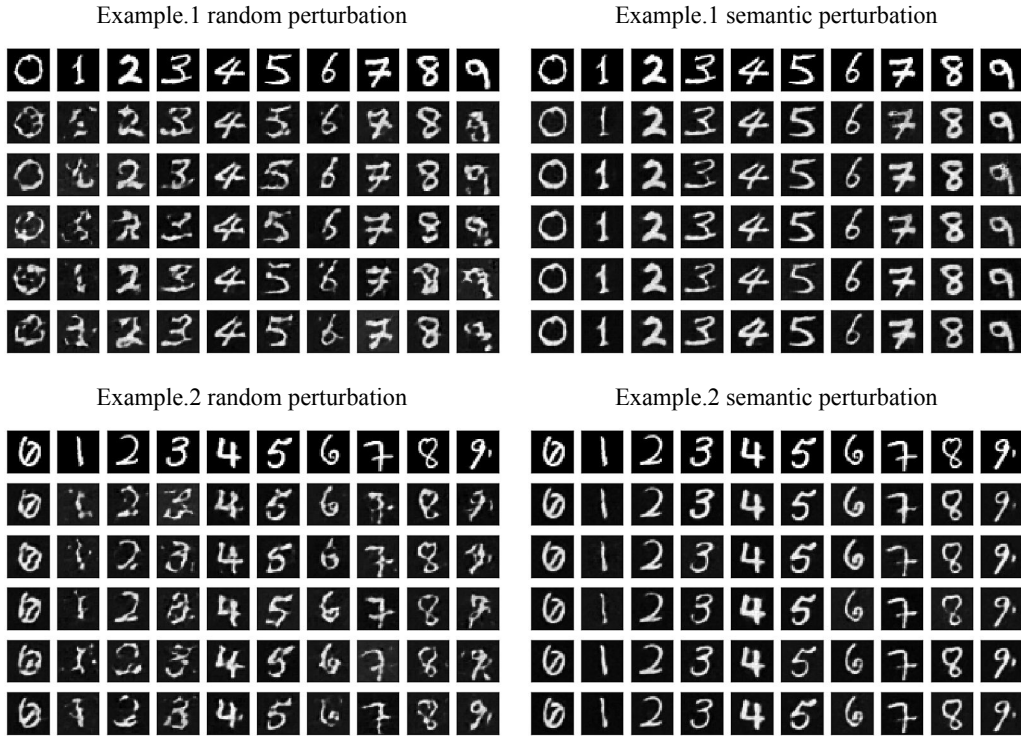



Figure 7: For each example, top row is the original examples selected from the training set, and the rest are reconstructed from the perturbed representations via random (left) and semantic (right) perturbations.

(A4.2) QUALITATIVE ANALYSIS

Extended from Section 4.4. Figure 8 shows the t-SNE results per class on MNIST. The overall tendency is similar to the description in Section 4.4.






Figure 8: From top to bottom: 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9. From left to right: training examples (circle), training examples (circle) + random-perturbed samples (cross), and training examples (circle) + semantic-perturbed samples (cross). Best viewed in color.