

Controllable Image Synthesis via SegVAE

Yen-Chi Cheng^{1,2}, Hsin-Ying Lee¹, Min Sun², Ming-Hsuan Yang^{1,3}

¹University of California, Merced ²National Tsing Hua University ³Google Research

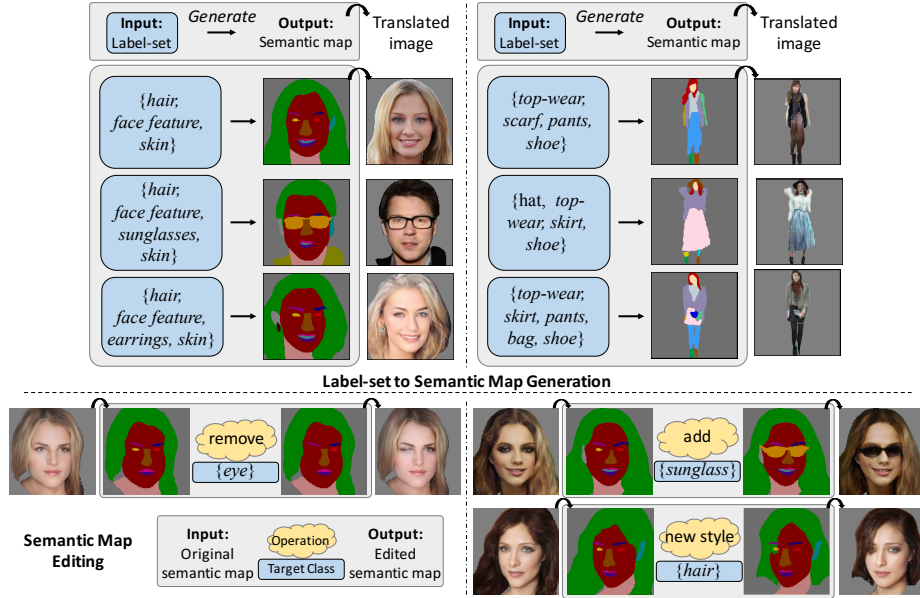


Fig. 1: **Label-set to Semantic map generation.** (*Top*) Given a label-set, SegVAE can generate diverse and realistic semantic maps. (*Bottom*) The proposed model enables several real-world flexible image editing.

Abstract. Flexible user controls are desirable for content creation and image editing. A semantic map is commonly used intermediate representation for conditional image generation. Compared to the operation on raw RGB pixels, the semantic map enables simpler user modification. In this work, we specifically target at generating semantic maps given a label-set consisting of desired categories. The proposed framework, SegVAE, synthesizes semantic maps in an iterative manner using conditional variational autoencoder. Quantitative and qualitative experiments demonstrate that the proposed model can generate realistic and diverse semantic maps. We also apply an off-the-shelf image-to-image translation model to generate realistic RGB images to better understand the quality of the synthesized semantic maps. Finally, we showcase several

real-world image-editing applications including object removal, insertion, and replacement.

1 Introduction

The recent success of deep generative models has made breakthroughs in a wide range of tasks such as image and video synthesis [2,9,11]. In addition to conventional generative models that aims to generate images from noise vectors sampled from prior distributions, conditional generative models have getting attention to handle various tasks including image-to-image translation (I2I) [6,15,43], text-to-image synthesis [5,41], and audio-to-video generation [18], to name a few. One major goal of these conditional generation tasks is to enable flexible user control, image editing, and content creation. Conditional generative models can greatly shorten the distance between professional creators and general users.

Among all forms of conditional context, semantic maps are recently getting attention. Semantic maps can be used as an intermediate representation or directly as inputs. As an intermediate representation, semantic maps serve as mediums that facilitate the original tasks such as image generation from text and scene graph [5,7]. As inputs, semantic maps can be translated to realistic images via I2I models [27]. These I2I models enable editing on the semantic maps, which is easier and more flexible than operating on the RGB space. However, in terms of image editing, users need to create semantic maps that are realistic in order to generate realistic RGB images. It is crucial to provide users greater flexibility and less overhead on image editing.

In this work, we focus on generating semantic maps given a label-set. The task is challenging for two reasons. First, the shape of components in the semantic maps not only need to be realistic, but also have to be mutually compatible. Second, the generation is inherently multimodal, that is, one label-set can correspond to multiple semantic maps. To handle these issues, we propose **SegVAE**, a VAE-based framework that can generate semantic maps in an iterative manner. For *compatibility*, the proposed model performs generation at each iteration conditioned on the label-set as well as previously generated components. For *diversity*, the proposed method learns a shape prior distribution that can be randomly sampled during the inference stage.

We evaluate the proposed methods through extensive qualitative and quantitative experiments. We conduct experiments on two diverse datasets, the CelebAMaskHQ [13] and HumanParsing [22] datasets, to demonstrate the general effectiveness of the proposed framework. We leverage Fréchet Inception Distance (FID) [3] and conduct a user study to evaluate realism. For diversity, we measure feature distances similar to the Learned Perceptual Image Patch Similarity (LPIPS) [42] metric. Furthermore, we demonstrate several real-world editing scenarios to showcase the superb controllability of the proposed method. We also apply an I2I model, SPADE [27], to synthesize realistic images based on the generated semantic maps to better visualize the quality of the proposed model.

We make the following contributions:

- We propose a new approach that can generate semantic maps from label-sets. Components in the generated semantic maps are mutually compatible and the overall semantic maps are realistic. SegVAE can also generate diverse results.
- We validate that our method performs favorably against existing methods and baselines in terms of realism and diversity on the CelebAMask-HQ and HumanParsing datasets.
- We demonstrate several real-world image editing applications using the proposed method. Our model enables flexible editing and manipulation. Our code and more results is available at <https://github.com/yccyenchicheng/SegVAE>.

2 Related Work

Generative models. Generative models aim to model a data distribution given a set of samples from that distribution. The mainstream of generative models approximates the distribution through maximum likelihood. There are two branches of the maximum likelihood method. One stream of work explicitly models the data distribution. PixelRNN [25] and PixelCNN [24] are auto-regressive models that perform sequential generation where the next pixel value is predicted conditioned on all the previously generated pixel values. Variational autoencoder [11] model the real data distribution by maximizing the lower bound of the data log-likelihood. The other stream of work learns implicit data distribution. Generative adversarial networks [2] model the data distribution by a two-player game between a generator and a discriminator.

Based on conventional generative models, conditional generative models synthesize images based on various contexts. Conditional generative models can formulate a variety of topics in image editing and content creations, including super-resolution [12], image-to-image translation [16,43], text-to-image generation [37,41], video generation [20,39], and music-to-dance translation [18].

Generative models with semantic maps. Semantic map is an important modality in generative modeling. There are two major ways of using semantic maps. First, semantic maps can be used as the conditional context for conditional image synthesis. Image-to-image translation models can learn the mapping from semantic maps to realistic RGB images [6,26,27,38]. Second, semantic maps can serve as an intermediate representation during the training of image synthesis conditioned on text or scene graph. [7,19,31,33,35]. Using semantic layouts provides rough guidance of the location and appearance of objects and further facilitate the training. In this work, we focus on generating semantic maps directly from a label-set. The proposed model enables flexible user editing and the generated results can be further used for photorealistic image generation.

Image editing. Image editing is a process of altering images. Traditional image editing mostly focuses on low-level changing like image enhancement, image retouching, and image colorization. In this work, we target at generating and altering the content of images. For generation, previous work can synthesize desired output given instructions like a set of categories [8,17] or a sentence [34], or change the semantic information in a user-specified region of the




Fig. 2: **Overview.** (a) Given a label-set as input, we iteratively predict the semantic map of each category starting from a blank canvas. we initialize the initial canvas as a blank semantic map x_0 . The generation conditions on the embedding of (*Current canvas*, *Target class*, *Label-set*). (b) In each iteration, the context encoder E^c encodes the input into the conditional context which is copied as the input to E_ϕ , E_ψ , and D . During training, the posterior encoder $\{E_\phi, LSTM_\phi\}$ takes the ground-truth semantic map of the target class as additional input to output $N(\mu_\phi, \sigma_\phi)$, which is then used to reconstruct the ground-truth semantic map. The prior encoder $\{E_\psi, LSTM_\psi\}$ encodes the conditional context to output $N(\mu_\psi, \sigma_\psi)$, which enables the stochastic sampling during test time.

target image [32]. Other stream of work can perform operations like object insertion [14, 23], object removal [40], and object attribute editing [36]. The proposed method can achieve both generations: given specified label-sets, and editing including object insertion and removal.

3 Semantic Maps Generation

Our goal is to learn the mapping from label-sets \mathcal{X} to semantic maps $\mathcal{Y} \subset \mathbb{R}^{H \times W \times C}$. A label-set $\mathbf{x} = [x_1, x_2, \dots, x_c] \in \mathcal{X}$ is a binary vector with c classes, where $x_i \in \{0, 1\}$. A semantic map $y \in \mathcal{Y}$ denotes a c channels label map where $Y_{i,j,k} \in \{0, 1\}$ for class k , $k = 1, 2, \dots, C$. The proposed **SegVAE** consist of a context encoder E_c , a posterior encoders $\{E_\phi, LSTM_\phi\}$, a learned prior encoders $\{E_\psi, LSTM_\psi\}$, and a decoder D . The context encoder E^c aims to encode information including label-set, current canvas, and target label into a context for the following conditional generation. The posterior network $\{E_\phi, LSTM_\phi\}$, consisting of a convolutional encoder and an LSTM, aims to conditionally en-

code input each iteration into an inference prior. The learned prior network $\{E_\psi, LSTM_\psi\}$, consisting of a fully-connected encoder and an LSTM, targets at learning a conditional prior distribution to be sampled from during inference time. The decoder D then learns the mapping from both posterior and learned prior distribution to semantic maps.

In this section, we first introduce the background knowledge of conditional variational auto-encoder in Section 3.1. Then we detail the proposed model in Section 3.2. Finally, we provide implementation details in Section 3.3.

3.1 Conditional VAE

Variational autoencoder. A Variational Autoencoder (VAE) [11,29] attempts to explicitly model the real data distribution by maximizing the lower bound of the data log-likelihood. VAE leverages a simple prior $p_\theta(\mathbf{z})$ (e.g., Gaussian) and a complex likelihood $p_\theta(\mathbf{x}|\mathbf{z})$ on latent variable \mathbf{z} to maximize the data likelihood $p_\theta(\mathbf{x})$. An inference network $q_\phi(\mathbf{z}|\mathbf{x})$ is introduced to approximate the intractable latent posterior $p_\theta(\mathbf{z}|\mathbf{x})$. Here θ and ϕ denote the parameters of the generation and inference networks. We then jointly optimize over θ and ϕ ,

$$\begin{aligned} \log p_\theta(\mathbf{x}) &= \log \int_{\mathbf{z}} p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z}) d\mathbf{z} \\ &\geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}) - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \end{aligned} \quad (1)$$

With this inequality, the variational autoencoder aims to reconstruct data \mathbf{x} with latent variable \mathbf{z} sampled from the posterior $q_\phi(\mathbf{z}|\mathbf{x})$ while minimizing the KL-divergence between the prior $p(\mathbf{z})$ and posterior $q_\phi(\mathbf{z}|\mathbf{x})$.

Conditional variational autoencoder. A Conditional Variational Autoencoder (C-VAE) [30] is an extension of the VAE which condition on a prior information described by a variable or feature vector \mathbf{c} . The generation and inference network’s output will base on the conditional variable \mathbf{c} , and the optimization over θ and ϕ becomes,

$$\begin{aligned} \log p_\theta(\mathbf{x}, \mathbf{c}) &= \log \int_{\mathbf{z}} p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})p(\mathbf{z}) d\mathbf{z} \\ &\geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c}) - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p(\mathbf{z}|\mathbf{c})). \end{aligned} \quad (2)$$

3.2 Iterative Generation with Learned Prior

A label-set consists of various amount of categories. To handle the dynamic amount of categories as well as capture dependency among categories, we leverage Long-Short Term Memory (LSTM) [4] as our recurrent backbone. However, how to perform stochastic generation and model data variations remain a challenge. The complexity is twofold. First, there are various possible combinations of categories to construct a label-set. For example, $y_1 = \{t\text{-shirt}, short, shoe\}$ and $y_2 = \{sunglasses, dress, bag\}$ represent two completely different outfit style. Second, in addition to label-set-level variations, category-level variations also

need to be captured. For instance, each *bag* and *dress* have diverse possible shapes. Moreover, the label-set-level variations and the category-level variations are not independent. For example, different choices of *t-shirt* will affect the choices of *dress* considering compatibility.

To handle these issues, we leverage C-VAE as our generative model. However, the usage of $\mathcal{N}(\mathbf{0}, \mathbf{I})$ which is conventionally used in VAE ignores dependencies between iterations since the priors are drawn independently at each iteration. Therefore, we adopt the similar idea from [1] to use a parameterized network ψ other than $\mathcal{N}(\mathbf{0}, \mathbf{I})$ for inferring the shape prior distribution. In order to learn the conditional prior, we first define our conditional context at iteration t as:

$$c_t = E^c(y_t, y, x_{1:t-1}), \quad (3)$$

where y_t is the current category to be generated, y represents the given label set, and x_k denote the semantic map at iteration k . The learned prior can thus be modeled with $\{E_\psi, LSTM_\psi\}$. At iteration t ,

$$\begin{aligned} h_t &= E_\psi(c_t) \\ \mu_\psi^t, \sigma_\psi^t &= LSTM_\psi(h_t) \\ z_t &\sim \mathcal{N}(\mu_\psi^t, \sigma_\psi^t). \end{aligned} \quad (4)$$

The learned prior is trained with the help of the posterior network $\{E_\phi, LSTM_\phi\}$ and the decoder D . At iteration t , given x_t , the semantic map of current category, we perform reconstruction by

$$\begin{aligned} h_t &= E_\phi(c_t, x_t) \\ \mu_\phi^t, \sigma_\phi^t &= LSTM_\phi(h_t) \\ z_t &\sim \mathcal{N}(\mu_\phi^t, \sigma_\phi^t) \\ \hat{x}_t &= D(z_t, c_t). \end{aligned} \quad (5)$$

Therefore, the model can be trained with a reconstruction loss and a KL loss at iteration t :

$$\begin{aligned} L_{\text{recon}}^t &= \|\hat{x}_t - x_t\|_1 \\ L_{\text{KL}}^t &= -D_{\text{KL}}(\mathcal{N}(\mu_\phi^t, \sigma_\phi^t) \parallel \mathcal{N}(\mu_\psi^t, \sigma_\psi^t)). \end{aligned} \quad (6)$$

We express the final objective as:

$$L = \mathbb{E}_{x,y} \left[\sum_{i=1}^{|y|} \lambda_{\text{recon}} L_{\text{recon}}^i + \lambda_{\text{KL}} L_{\text{KL}}^i \right], \quad (7)$$

where the hyper-parameters λ s control the importance of both term.

In the inference time, given a label-set as input, we initialize an empty label map as x_0 . We then generate x_k autoregressively by setting the inputs for p_ψ with the generated shapes $\hat{x}_{1:t-1}$.

Table 1: **Realism and Diversity.** We use the FID and a feature space distance metric to measure the realism and diversity (\pm indicates the 95% C.I.).

Method	HumanParsing		CelebAMask-HQ	
	FID \downarrow	Diversity \uparrow	FID \downarrow	Diversity \uparrow
C-GAN	171.1293 \pm .3359	N/A	76.0115 \pm .1981	N/A
C-VAE _{sep}	85.0505 \pm .3052	.1781 \pm .040	39.7445 \pm .2254	.1566 \pm .033
C-VAE _{global}	83.8214 \pm .6747	.1730 \pm .034	36.2903 \pm .2084	.1582\pm.031
sg2im-orig [7]	207.0786 \pm .3324	N/A	208.8142 \pm .1876	N/A
sg2im [7]	56.7421 \pm .2206	.2064 \pm .050	34.7316 \pm .3071	.1476 \pm .045
Ours w/o LSTM	50.8830 \pm .2374	.2024 \pm .045	34.5073 \pm .2156	.1535 \pm .034
Ours w/o Learned Prior	44.6217 \pm .2881	.1625 \pm .054	33.8060 \pm .3167	.1492 \pm .038
Ours	39.6496\pm.3543	.2072\pm.053	28.8221\pm.2732	.1575\pm.043
GT	33.1562 \pm .3801	.2098 \pm .050	22.5981 \pm .0870	.1603 \pm .045

3.3 Implementation Details

We implement our model in PyTorch [28]. For all experiments, we use the resolution of 128×128 for the input image and semantic map. For the context encoder E^c , we use two multilayer perceptrons to produce the embeddings of the label-set and target label, and fuse them with the current canvas. Then a six convolution layers will encode the embeddings and the current canvas into the conditional context. For the E_ψ , we apply a fully connected layers with LSTM followed by two multilayer perceptrons to output the mean and log variance. The E_ϕ is similar to E_ψ by replacing the fully connected layers with the convolution layers to encoder the ground truth semantic map. We use the latent code size of $z_\psi, z_\phi \in \mathbb{R}^{384}$ for all experiments. Finally, the D consists of five fractionally strided convolution layers. We apply the instance norm and spectral norm for all the architectures. For more details of the network architecture, please refer to the supplementary material.

For training, we use the Adam optimizer [10] with a learning rate of $5e^{-5}$, a batch size of 24, and $(\beta_1, \beta_2) = (0.5, 0.999)$. For the HumanParsing dataset, we use $\lambda_{\text{recon}} = 1$, $\lambda_{\text{KL}} = 1e^{-4}$ and $\lambda_{\text{recon}} = 1$, $\lambda_{\text{KL}} = 1e^{-7}$ for the CelebAMask-HQ dataset. For the order of the iterative generation with learned prior, we use $\{\text{face, hair, left arm, right arm, left leg, right leg, upper clothes, dress, skirt, pants, left shoe, right shoe, hat, sunglasses, belt, scarf, bag}\}$ for the HumanParsing dataset. In addition, $\{\text{skin, neck, hair, left eyebrow, right eyebrow, left ear, right ear, left eye, right eye, nose, lower lip, upper lip, mouth, hat, cloth, eyeglass, earrings, necklace}\}$ for the CelebAMask-HQ dataset.

4 Experimental Results

Datasets. We perform the evaluation on two datasets:

- **HumanParsing.** The HumanParsing dataset [22], extended from [21], contains $\sim 17,000$ of street fashion images and 18 classes of semantic maps with pixel-annotations. It is composed of diverse appearances and multiple combinations of fashion items and human parts. We first clean the dataset by inspecting the aspect ratio of the ground truth semantic maps and remove those over ± 1 standard deviation. For each example, we compute the bounding box of each class on the semantic map and crop them accordingly.
- **CelebAMask-HQ.** CelebAMask-HQ [13] is a large-scale face image dataset which includes 30,000 high-resolution paired face images and semantic maps. The semantic maps contain 19 classes of facial attributes and accessories. The various shapes of facial attributes and accessories make it suitable for testing. We split the data into 80%, 10%, 10% for training, validation and testing.

Evaluated Method. We evaluate the following algorithms.

- **C-GAN.** We implement a conditional GAN which takes a label-set as input and generates the corresponding images. This method has to handle the class-dependency, compatibility between shapes, and the image quality at the same time, which imposes a great burden on a single model.
- **C-VAE_{sep}.** This baseline generates semantic maps for each category independently. The generated semantic maps are then aggregated to the final output.
- **C-VAE_{global}** This baseline takes global context into consideration while generating semantic maps for each category. We encode the label-set into a global feature, which serves as the conditional context for the generation. With reference to the global feature, the generated shapes will be more compatible.
- **sg2im [7].** Sg2im is a conditional generative model that takes a scene graph as input and output the corresponding image. To facilitate the learning, a semantic map is first predicted from the object embedding during the optimization process. We compare the predicted semantic maps from sg2im given a label-set as input, which can be seen as the simplest form of a scene graph. We use the official implementation provided by [7] for training. For a fair comparison, we provide ground truth bounding boxes for sg2im when generating the masks. We report both the metrics from the images translated from its predicted semantic maps using SPADE [27] and the images generated by sg2im (denoted by “sg2im” and “sg2im-orig” respectively in Table 1).

4.1 Quantitative Evaluation

Visual quality. We evaluate the realism of the generated semantic maps using the Fréchet Inception Distance (FID) [3] metric. We first generate a semantic map given a label-set, then we use an off-the-shelf I2I model [27] to output the translated image. We compute FID with the official implementation from [3] on all compared methods to measure the realism of the generated semantic maps. Table 1 shows that the proposed model performs favorably against the baselines and the existing method.

Diversity. We measure the diversity by computing distance between the generated semantic maps using the distance metric similar to the LPIPS [42] metric.

Table 2: **Compatibility and reconstructability.** We train a shape predictor to measure the compatibility error (abbreviated as Compat. error) over the generated shapes. We adopt an auto-encoder to measure the quality of our generated results by computing the reconstruction error (denoted as Recon. error).

Method	HumanParsing		CelebAMask-HQ	
	Compat. error ↓	Recon. error ↓	Compat. error ↓	Recon. error ↓
C-VAE _{sep}	.7823 \pm .0161	.6857 \pm .010	.1029 \pm .0017	.1165 \pm .003
C-VAE _{global}	.7345 \pm .0141	.6186 \pm .018	.1015 \pm .0040	.1142 \pm .003
sg2im [7]	.6983 \pm .0176	.5434\pm.012	.0844 \pm .0020	.1334 \pm .003
Ours	.6174\pm.0147	.5663 \pm .011	.0754\pm.0013	.0840\pm.001

However, there are no general feature extractors for semantic maps. Therefore, we trained an auto-encoder on the ground truth semantic maps. We use the intermediate representation as the extracted features for semantic maps. We can then measure the feature distance between two semantic maps. We measure the diversity between 3000 pairs of generated semantic maps by sampling from the test-set. Table 1 shows that the proposed method generates diverse semantic maps without sacrificing visual quality.

Compatibility and reconstruction error. To evaluate the compatibility of a generated semantic maps, we design a metric measure the performance of all compared methods. A shape is compatible if it can be easily inferred from all the other class’s semantic maps. We measure this quantitatively by training a shape predictor which takes a semantic map and a target label as input, and outputs the shape of the target class. The training pair of the data is created by excluding one of the class in the ground truth semantic map as the prediction target, and the remaining classes’ semantic maps along with the class of the excluded target form the input. Then we use this shape predictor to compute the compatibility error for all the compared methods. Meanwhile, we also train an auto-encoder on the ground-truth semantic maps to measure if one generated result is close to the real data distribution. Given a generated semantic map as input, we calculate the reconstruction error between the input and the reconstructed result output by the auto-encoder. Table 2 shows that the proposed method generates compatible and reasonable semantic maps.

Prediction order analysis. We analyze how the prediction order will affect the performance of the proposed model at Tables 3. First, we organize the classes in each dataset into three major categories for two datasets where the class in each category has similar properties. Then, we train the proposed methods with different permutations of these orders.

HumanParsing. We organize the classes into *Body*, *Clothes*, *Accessories*. The classes in each category are

- *Body*: {face, hair, left arm, right arm, left leg, right leg}.
- *Clothes*: {upper clothes, dress, skirt, pants, left shoe, right shoe}.

Table 3: **Prediction order analysis.** (Upper) On the HumanParsing dataset, order-1 and order-2 perform similarly since both *Body*, *Clothes* provide great contexts for the model. In contrast, order-3 degrades the generation results and the diversity since *Accessories* offers limited contexts. (Lower) On the CelebAMask-HQ dataset, *Face* is crucial for serving as the initial canvas for the subsequent generation. Therefore order-1 outperforms order-2 and order-3. in FID. Similarly, order-3 largely constrains the possible generation for the remaining class and causes lower diversity.

Order (HumanParsing)	FID↓	Diversity↑
1 <i>Body</i> → <i>Clothes</i> → <i>Accessories</i> (Ours)	39.6496^{±.3543}	.2072^{±.053}
2 <i>Clothes</i> → <i>Body</i> → <i>Accessories</i>	39.9008 ^{±.5263}	.2062 ^{±.0494}
3 <i>Accessories</i> → <i>Body</i> → <i>Clothes</i>	40.2909 ^{±.2195}	.2043 ^{±.0521}
Order (CelebAMask-HQ)	FID↓	Diversity↑
1 <i>Face</i> → <i>Face features</i> → <i>Accessories</i> (Ours)	28.8221^{±.2732}	.1575^{±.043}
2 <i>Face features</i> → <i>Face</i> → <i>Accessories</i>	30.6547 ^{±.1267}	.1517 ^{±.0376}
3 <i>Accessories</i> → <i>Face</i> → <i>Face features</i>	32.0325 ^{±.1294}	.1489 ^{±.0363}

- *Accessories*: {*hat*, *sunglasses*, *belt*, *scarf*, *bag*}.

Three orders are as follows. Order-1: *Body*, *Clothes*, *Accessories* (Ours), order-2: *Clothes*, *Body*, *Accessories*, and order-3: *Accessories*, *Body*, *Clothes*.

CelebAMask-HQ. We organize the classes into *Face*, *Face features*, *Accessories*. The classes in each category are

- *Face*: {*skin*, *neck*, *hair*}.
- *Face features*: {*left eyebrow*, *right eyebrow*, *left ear*, *right ear*, *left eye*, *right eye*, *nose*, *lower lip*, *upper lip*, *mouth*}.
- *Accessories*: {*hat*, *cloth*, *eyeglass*, *earrings*, *necklace*}.

We experiment on three prediction orders. Order-1: *Face*, *Face features*, *Accessories* (Ours), order-2: *Face features*, *Face*, *Accessories*, and order-3: *Accessories*, *Face*, *Face features*.

Table 3 shows how the prediction order affects the performance of realism and diversity. For the HumanParsing dataset, order-1 and order-2 perform similarly on FID and diversity since both *Body* and *Clothes* provide great contexts for the model. While in order-3, *Accessories* do not deliver good information as the other two categories and will constrain the possible shapes and locations when generating the semantic maps of *Body* and *Clothes*. This results in a degradation in FID and diversity. For the CelebAMask-HQ dataset, the *Face* category is essential. Acting as a canvas, it ensures the generation of the subsequent class is located in the semantically meaningful position, and gives the most degree of freedom when generating compatible shapes compared to other orders.

Ablation studies. We analyze the contribution of each component of our method with ablation studies.




Fig. 3: **User study.** We conduct a user study to evaluate the realism of the generated semantic maps. We have two sets of comparison: the comparison with other algorithms, and all methods to real maps. Results showed that users favored SegVAE (gray bar) against all other compared methods in both settings.

- **Ours w/o LSTM.** This model omits the LSTM when outputting the distribution of the learned prior and the posterior.
- **Ours w/o Learned Shape Prior.** This version adopts a fixed Gaussian distribution for the shape prior. The proposed method with fixed prior.

The results in Table 1 show the necessity of the proposed model’s architecture design. LSTM serves as a crucial role to handle the dynamic among the classes as well as capture the dependency among them. On both datasets, the model yields significant improvements over FID in contrast to the model with fixed prior. Meanwhile, adopting the learned shape prior mainly contributes to the better diversity of the generated outputs and improves the realism.

User study. To better explicitly evaluate the visual quality of the generated semantic maps without relying on the off-the-shelf image-to-image translation model, we conduct a user study. We perform two sets of comparisons. First, we compare the proposed method to baselines. Second, we compare all methods to ground-truth semantic maps. In each test, users are given two semantic maps generated from different methods and asked “which semantic map looks more realistic?”. We ask each user to compare 20 pairs and collect results from a total of 60 subjects. Figure 3 presents the results of our user study.

4.2 Qualitative Evaluation

Visual quality. We compare the visual quality of the generated semantic maps from the proposed method and the baseline methods. As shown in Figure 4, the proposed model can generate semantic maps that are more realistic than others. The **C-VAE_{sep}** model generates each category independently. The generation of each semantic map is not only unaware of the shape of other categories, but even unknowing of what other categories are in this label-set. Although the individual semantic maps are reasonable, they are not mutually compatible after combination. The **C-VAE_{global}** model performs slightly better with taking the label-set information as an additional input. However, the generation process still disregards the appearance of each other. Finally, the proposed SegVAE out-



Fig. 4: **Qualitative comparison.** We present the generated semantic maps given label-sets on the CelebAMask-HQ (left) and the HumanParsing (right) datasets. The proposed model generates images with better visual quality compared to other methods. We also present the translated realistic images via SPADE. Please refer to Figure 5 for the color mapping for each category.

performs sg2im from several perspectives. First, the iterative generation allows us to use teacher-forcing in training, which helps capture the dependency among components. Second, we can better capture the appearance distribution with the help of VAE. Finally, the learned latent space enables random sampling during inference time, which supports generating diverse outputs.

Multi-modality. We demonstrate the capability of the proposed method to generate diverse outputs in Figure 5. SegVAE can generate diverse appearances of objects given the same label-set. For example, the eyeglasses (the first row of CelebAMask-HQ), the hat (the second row of CelebAMask-HQ), the skirt (the second row of HumanParsing), and the scarf (the fourth row of HumanParsing). The generated semantic maps are mutually compatible thanks to the iterative conditional generation process. In the fourth row of CelebAMask-HQ, the hair and face change jointly to different orientations. In the first row of HumanParsing, both left- and right-leg changes jointly to different poses.



Fig. 5: **Multi-modality.** We demonstrate the ability of SegVAE to generate diverse results given a label-set on both datasets.

4.3 Editing

The proposed method enables flexible user-control for image editing. We demonstrate three real-world applications, including *remove*, *add*, and *new style*. For all applications, we first obtain the semantic maps from the input RGB images. Then we perform editing on the semantic maps with SegVAE. Finally, we translate the edited semantic maps back to RGB images with an I2I model. We show the results of image editing in Figure 6. For *Remove*, we remove a category from the label-set. For *Add*, we add an additional category to the existing label-set to perform object insertion. For *New style*, we could alter the style of the targeted category due to the diverse outputs of the purposed method.



Fig. 6: **Editing.** We present three real-world image editing applications: *add*, *remove*, and *new style*. We show results of three operations on both datasets.

5 Conclusion

In this paper, we present a novel VAE-based framework with iterative generation process for label-set to semantic maps generation. To generate realistic appearance and mutually compatible shapes among the classes, we propose to generate the target shape in an autoregressive fashion by conditioning on the previously generated components. To handle the diverse output, we adopt a parameterized network for learning a shape prior distribution. Qualitative and quantitative results show that the generated semantic maps are realistic and diverse. We also apply an image-to-image translation model to generate RGB images to better understand the quality of the synthesized semantic maps. Finally, we showcase several real-world image-editing scenarios including removal, insertion, and replacement, which enables better user-control over the generated process.

6 Acknowledgement

This work is supported in part by the NSF CAREER Grant #1149783, MOST 108-2634-F-007-016-, and MOST 109-2634-F-007-016-.

References

1. Denton, E., Fergus, R.: Stochastic video generation with a learned prior. In: ICML (2018) 6
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014) 2, 3
3. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NIPS (2017) 2, 8
4. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* 9(8), 1735–1780 (1997) 5
5. Hong, S., Yang, D., Choi, J., Lee, H.: Inferring semantic layout for hierarchical text-to-image synthesis. In: CVPR (2018) 2
6. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017) 2, 3
7. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: CVPR (2018) 2, 3, 7, 8, 9
8. Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: Layoutvae: Stochastic scene layout generation from a label set. In: ICCV (2019) 3
9. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019) 2
10. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015) 7
11. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014) 2, 3, 5
12. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., Shi, W.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR (2017) 3
13. Lee, C.H., Liu, Z., Wu, L., Luo, P.: Maskgan: Towards diverse and interactive facial image manipulation. *arXiv preprint arXiv:1907.11922* (2019) 2, 8
14. Lee, D., Liu, S., Gu, J., Liu, M.Y., Yang, M.H., Kautz, J.: Context-aware synthesis and placement of object instances. In: NeurIPS (2018) 4
15. Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M.K., Yang, M.H.: Diverse image-to-image translation via disentangled representations. In: ECCV (2018) 2
16. Lee, H.Y., Tseng, H.Y., Mao, Q., Huang, J.B., Lu, Y.D., Singh, M.K., Yang, M.H.: Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision* pp. 1–16 (2020) 3
17. Lee, H.Y., Yang, W., Jiang, L., Le, M., Essa, I., Gong, H., Yang, M.H.: Neural design network: Graphic layout generation with constraints. In: ECCV (2020) 3
18. Lee, H.Y., Yang, X., Liu, M.Y., Wang, T.C., Lu, Y.D., Yang, M.H., Kautz, J.: Dancing to music. In: NeurIPS (2019) 2, 3
19. Li, W., Zhang, P., Zhang, L., Huang, Q., He, X., Lyu, S., Gao, J.: Object-driven text-to-image synthesis via adversarial training. In: CVPR (2019) 3
20. Li, Y., Min, M.R., Shen, D., Carlson, D., Carin, L.: Video generation from text. In: AAAI (2018) 3
21. Liang, X., Liu, S., Shen, X., Yang, J., Liu, L., Dong, J., Lin, L., Yan, S.: Deep human parsing with active template regression. *TPAMI* (2015) 8
22. Liang, X., Xu, C., Shen, X., Yang, J., Liu, S., Tang, J., Lin, L., Yan, S.: Human parsing with contextualized convolutional neural network. In: ICCV (2015) 2, 8

23. Lin, C.H., Yumer, E., Wang, O., Shechtman, E., Lucey, S.: St-gan: Spatial transformer generative adversarial networks for image compositing. In: CVPR (2018) [4](#)
24. Van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al.: Conditional image generation with pixelcnn decoders. In: NIPS (2016) [3](#)
25. Oord, A.v.d., Kalchbrenner, N., Kavukcuoglu, K.: Pixel recurrent neural networks. In: ICML (2016) [3](#)
26. Pan, J., Wang, C., Jia, X., Shao, J., Sheng, L., Yan, J., Wang, X.: Video generation from single semantic label map. In: CVPR (2019) [3](#)
27. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: CVPR (2019) [2, 3, 8](#)
28. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS (2019) [7](#)
29. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: ICML (2014) [5](#)
30. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: NIPS (2015) [5](#)
31. Sun, W., Wu, T.: Image synthesis from reconfigurable layout and style. In: ICCV (2019) [3](#)
32. Suzuki, R., Koyama, M., Miyato, T., Yonetsuji, T.: Spatially controllable image synthesis with internal representation collaging. arXiv preprint arXiv:1811.10153 (2018) [4](#)
33. Talavera, A., Tan, D.S., Azcarraga, A., Hua, K.: Layout and context understanding for image synthesis with scene graphs. In: ICIP (2019) [3](#)
34. Tan, F., Feng, S., Ordonez, V.: Text2scene: Generating compositional scenes from textual descriptions. In: CVPR (2019) [3](#)
35. Tripathi, S., Bhiwandiwalla, A., Bastidas, A., Tang, H.: Heuristics for image generation from scene graphs. In: ICLR workshop (2019) [3](#)
36. Tseng, H.Y., Fisher, M., Lu, J., Li, Y., Kim, V., Yang, M.H.: Modeling artistic workflows for image generation and editing. In: ECCV (2020) [4](#)
37. Tseng, H.Y., Lee, H.Y., Jiang, L., Yang, W., Yang, M.H.: Retrievegan: Image synthesis via differentiable patch retrieval. In: ECCV (2020) [3](#)
38. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: CVPR (2018) [3](#)
39. Wang, T.H., Cheng, Y.C., Lin, C.H., Chen, H.T., Sun, M.: Point-to-point video generation. In: ICCV (2019) [3](#)
40. Yang, J., Hua, K., Wang, Y., Wang, W., Wang, H., Shen, J.: Automatic objects removal for scene completion. In: INFOCOM WKSHPS (2014) [4](#)
41. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stack-gan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017) [2, 3](#)
42. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) [2, 8](#)
43. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networkss. In: ICCV (2017) [2, 3](#)