SESAME: Semantic Editing of Scenes by Adding, Manipulating or Erasing Objects

Evangelos Ntavelis^{1,2}, Andrés Romero¹, Iason Kastanis², Luc Van Gool^{1,3}, and Radu Timofte¹

¹ Computer Vision Lab, ETH Zurich, Switzerland

² Robotics and Machine Learning, CSEM SA, Switzerland

³ PSI, ESAT, KU Leuven, Belgium

Fig. 1. We assess SESAME on three tasks (a) image editing with free form semantic drawings (first row) (b) semantic layout driven semantic editing (second row) (c) layout to image generation with SESAME discriminator (third row)

Abstract. Recent advances in image generation gave rise to powerful tools for semantic image editing. However, existing approaches can either operate on a single image or require an abundance of additional information. They are not capable of handling the complete set of editing operations, that is addition, manipulation or removal of semantic concepts. To address these limitations, we propose SESAME, a novel generator-discriminator pair for **S**emantic **E**dit**I**ng of **S**cen**E**s by **A**dding, **M**anipulating or **E**rasing objects. In our setup, the user provides the semantic labels of the areas to be edited and the generator synthesizes the corresponding pixels. In contrast to previous methods that employ a discriminator that trivially concatenates semantics and image as an input, the SESAME discriminator is composed of two input streams that independently process the image and its semantics, using the latter to manipulate the results of the former. We evaluate our model on a diverse set of datasets and report state-of-the-art performance on two tasks: (a) image manipulation and (b) image generation conditioned on semantic labels.

Keywords: Generative Adversarial Networks, Interactive Image Editing, Image Synthesis

1 Introduction

Image editing is a challenging task that has received increasing attention in the media, movies and social networks. Since the early 90s, tools like Gimp [38] and Photoshop [36] have been extensively utilized for this task. Yet, both require high level expertise and are labour intensive. Generative Adversarial Networks (GANs) [10] provide a learning-based alternative able to assist non-experts to express their creativity when retouching photographs. GANs have been able to produce results of high photo-realistic quality [19,20]. Despite their success in image synthesis, their applicability on image editing is still not fully explored. Being able to manipulate images is a crucial task for many applications such as autonomous driving [15] and industrial imaging [7], where data augmentation boosts the generalization capabilities of neural networks [1,49,9].

Image manipulation has been used in the literature to refer to various tasks. In this paper, we follow the formulation of Bau *et al.* [3], and define the task of semantic image editing as the process of adding, altering and removing instances of certain classes or *semantic concepts* in a scene. Examples of such manipulations include but are not limited to: removing a car from a road scene, changing the size of the eyes of a person, adding clouds in the sky, etc. We use the term *semantic concepts* to refer to various class labels that can not be identified as objects, *e.g.*, mountains, grass, etc.

Training neural networks for visual editing is not a trivial task. It requires a high level of understanding of the scene, the objects, and their interconnections [45]. Any region of an image added or removed should look realistic and should also fit harmoniously with the rest of the scene. In contrast to image generation, the co-existence of real and fake pixels make the fake pixels more detectable, as the network cannot take the "easy route" of generating simple textures and shapes or even omit a whole class of objects [4]. Moreover, the lack of natural image datasets, where a scene is captured with and without an object, makes it impossible to train such models in a supervised manner.

One way to circumvent this problem is by inpainting the regions of an image we seek to edit. Following this scheme, we mask out and remove all the pixels we want to manipulate. Recent works [55,32,37,16] improve upon this approach by incorporating sketch and color inputs to further guide the generation of the missing areas and thus provide higher level control. However, inpainting can only tackle some aspects of semantic editing. To address this limitation, Hong *et al.* [12] manipulate the semantic layout of an image, and subsequently, they utilize it for inpainting the image. Yet, this approach requires access to the full semantic information of the image, which is costly to acquire.

To this end, we propose SESAME, a novel semantic editing architecture based on adversarial learning, able to manipulate images based on a semantic input. In particular, our method is able to edit images with pixel-level guidance of semantic

labels, permitting full control over the output. Note that our method requires the semantics *only* for regions to be edited. The generator seeks to synthesize an altered image such that the synthesized pixels comply with both the context and the user input. Moreover, we propose a new approach for semantics-conditioned discrimination, by utilizing two independent streams to process the input image and the corresponding semantics. We use the output of the semantics stream to manipulate the output of the image stream. We employ visual results along with quantitative analysis and a human study to validate the performance and flexibility of the proposed approach.

2 Related Work

Generative Adversarial Networks [10] have completely revolutionized a great variety of computer vision tasks such as image generation [20,19,30], super resolution [48,27], image attribute manipulation [28,40] and image editing [12,3].

While in their original formulation GANs were only capable of generating samples drawn from a random distribution [10], soon multiple models emerged able to perform **conditional image synthesis** [29,33]. This gave rise to approaches that employ a generative model for producing outputs conditioned on different types of information. There are many approaches targeting multiple levels of abstraction and locality of features that we seek to encapsulate in the output. For example, [29,31,57,5] focus on representing images characterized by a single label. In a different setting, [39,58,59,52] employ a text to image pipeline to provide a high-level description of the corresponding image. Recently, many methods utilize information of a scene graph [18,2] and sketches with color [42] to represent where objects should be positioned on the output image.

A more fine-grained approach aims to translate semantic maps, which carry pixel-wise information, to realistic-looking images [14,47,35,22]. For all the aforementioned models, the user can control the output image by altering the conditional information. Nonetheless, they are not suitable for manipulating an existing image, as they do not consider an image as an input.

User-guided semantic image editing is the task where the user is able to semantically edit an image by adding, manipulating or removing semantic concepts [3]. Both GANPaint [3] and SinGAN [43] are able to perform such operations. GANPaint [3] achieves this by manipulating the neuron activations and SinGAN [43] by learning the internal batch statistics of an image. However, both are trained on a single image and require retraining in order to be applied to another, while our model is able to handle manipulation of multiple images without retraining.

Another line of work is inpainting [13,56,25], where the user masks a region of the image for removal and the network fills it accordingly to the image context. This can be interpreted as a simple form of editing, but the user does not have control over the generated pixel. To address this, other research works guide the generation of the missing areas using edges [55,32] and/or color [37,16] information.

Recently, researchers shifted their attention to more semantic aware approaches for inpainting, by focusing on object addition and removal. Shetty *et al.* [44], for instance, propose a two stage architecture to address removal operations, by using an auxiliary network that predicts the masks of the objects during training, while, at inference, users provide their own masks. Note that their model cannot handle new objects generation. Another line of work is tackling object synthesis by utilizing semantic layout information, which provides a fine-grained guidance over the manipulation of an image. Yet, a subset of them is limited by generating objects from a single class [34,51] or placing prior fixed objects on the semantics plane [23]. Hong *et al.* [12] are able to handle both addition and removal, but require full semantic information of the scene to produce even the smallest change to an image. In contrast, our method requires only the semantics of the region to be edited.

The majority of the aforementioned works rely on adversarial learning to tackle the problem of image editing, by primarily focusing on adjusting the generator. Most recent models use a PatchGAN variant [14] which is able to discriminate on the high frequencies of the image. This is a desired attribute as conventional losses like *Mean Squared Error* and *Mean Absolute Error* can only convey information about the lower frequencies to the generator. PatchGAN can also be used for conditional generation of images on semantic maps, similar to our case study. Previous works targeting a similar problem concatenate the semantic information to the image and use it as an input to the discriminator. However, conventional conditional generation literature suggests that concatenation is not the optimal approach for conditional discrimination [39,33,31]. In this work, we extend PatchGAN to better incorporate conditional information by processing it separately from the image input. In a later stage of the network the two processed streams are merged to produce the final output of the discriminator.

3 SESAME

In this work we describe a deep learning pipeline for semantically editing images, using conditional Generative Adversarial Networks (cGANs). Given an image I_{real} and a semantic guideline of the regions that should be altered by the network, denoted by M_{sem} , we want to produce a realistic output I_{out} . The real pixels values corresponding to M_{sem} are removed from the input image. The generated pixels in their place should be both true to the semantics dictated by the mask and coherent with the rest of the pixels of I_{real} . In order to achieve this, our network is trained end-to-end in an adversarial manner. The generator is a Encoder-Decoder architecture, with dilated convolutions[53] and SPADE[35] layers, explained in section 3 and the discriminator is a two-stream patch discriminator, described in section 3.

SESAME Generator. Semantically editing a scene is an Image to Image translation problem. We want to transform an image where we substituted the RGB pixels of the regions with an one-hot semantics vector. From the generator’s

Fig. 2. The SESAME Generator aims to generate the pixels designated by the semantic mask so they are both (1) true to their label and (2) fit naturally to the rest of the picture. It is an encoder-decoder architecture with dilated convolutions to increase the receptive field as well as SPADE layers in the decoder to guide in-class generation

output, only the pixels on the masked out regions are retained, while the rest are retrieved from the original image:

$$I_{gen} = G(I_m, M, M_{sem}), \quad (1)$$

$$I_{out} = I_{gen} \cdot M + I_{real} \cdot (1 - M). \quad (2)$$

This architecture should accomplish two goals: generated pixels should 1) be coherent with their real neighboring ones as well as 2) be true to the semantic input. To achieve these goals we adapt our generator from the network proposed by Johnson *et al.* [17] to fill the gaps: two down-sampling layers, a semantic core made of multiple residual layers and two up-sampling ones.

We conceptually divide our architecture into two parts: the encoder and the decoder. In the encoder we aim to extract the contextual information of the pixels we want to synthesize. In the decoder part we combine the semantic information using Spatially Adaptive De-Normalization [35] blocks to every layer. As the area to be edited can span over a large region, we would like the receptive field of our network to be relatively large. Thus, we use dilated convolutions in the last and first layers of the encoder and the decoder respectively. A scheme of our SESAME generator can be seen in the Fig. 2, and for further details refer to the supplementary materials.

SESAME Discriminator. Layout to image editing can be seen as a sub-task of label to image translation. Inspired by Pix2Pix [14], more recent approaches [47,35] employ a variation of the PatchGAN discriminator. The Markovian discriminator, as it is also called, was a paradigm shift that made the discriminator focus on the higher frequencies by limiting the attention of the discriminator into local patches, producing a different fake/real prediction value for each of them. The subsequent methods added a multi-scale discrimination

Fig. 3. The SESAME discriminator, in contrast to the commonly used PatchGAN, is handling the RGB Image and its Semantics independently. Before the last convolutional layer the two streams, D_{RGB} and D_{sem} , are merged. The semantics stream is reduced via a *Sum Global Pooling* operation to a 2D matrix of spatial dimensions equal to the number of output patches. The feature vector of D_{RGB} at each path is scaled by D_{sem} and a residual is added to product

approach, the Feature Matching-Loss [47] and the use of Spectral Normalization [30] instead of Instance Normalization [35], which stabilized training and further improved the quality of the generated samples. However, the way the conditional information was provided to the discriminator remained unchanged.

Label to image generation is a sub-task of conditional image generation. In this more general category of methods, we can observe how the discriminator has evolved from the cGAN’s input concatenation [29], to concatenating the class information with a hidden layer [39], and lastly, to take the form of the projection discriminator [31]. In the latter approach, the inner product of the embedding of the conditional information and a feature extracted from the hidden layers of the discriminator are summed with the output of the discriminator to produce the final prediction. Each step of the conditional discriminator evolution improved the results over the naïve concatenation at its input [31]. On these methods the discriminator produces, nonetheless, a scalar output for the whole image.

We aim to design a discriminator for label to image generation that combines the aforementioned attributes. On the one hand, it should preserve the ability of PatchGAN to discriminate on high-frequencies. On the other hand, we want to enforce the semantic information guidance on the discriminator’s decision. If the pixels of the whole image shared semantic class, the projection discriminator would be easily extended to PatchGAN. In contrast, our case is characterized by fine-grained per pixel semantics: each output patch encompasses a variety of classes and different compositions of them.

Our proposed SESAME discriminator is comprised by two independent streams that handle the RGB and Semantic Labels inputs. As Fig. 3 depicts, the two

streams have identical architectures. Before the information is merged a *Sum Global Pooling* operation is applied on the outputs of the Semantics Stream. The output of the semantics stream is used to scale each output coming from the RGB stream. The resulted feature map is passed as input to a last 3×3 convolutional layer, which produces the final output. The process can be written as follows:

$$D(I, Sem) = Conv_{3 \times 3}(D_{RGB}(I_{out}) \cdot (1 + \sum_{channels} D_{sem}(Sem))), \quad (3)$$

where the D_{RGB} is the output of the RGB stream and D_{sem} of the semantics stream before the *Global Sum Pooling*. We also integrate the changes made to PatchGAN by Pix2PixHD [47] and SPADE [35]. We use a multi-scale discrimination scheme with squared patches and two different edge-sizes of 70 and 140 pixels, in order to provide also discrimination at a coarser level and Spectral Normalization. The input to the semantic stream is the same for both fake and real images discrimination, so we only need to calculate D_{sem} once. Moreover, it makes sense to apply the Feature Matching Loss only to the Feature Maps produced by the RGB stream.

Training Losses. We train the Generator in an adversarial manner using the following losses: Perceptual Loss [17], Feature Matching Loss [41] and Hinge Loss [24,46,30] as the Adversarial Loss. Early experiments with Style Loss [17] did not improve the results. Accordingly:

$$L_G = \lambda_{percept} \cdot L_{perc} + \lambda_{feat} \cdot L_{FM} - E_{z \sim p(z)}[D_k(I_{out}, M, M_{sem})]. \quad (4)$$

For the discriminator at each scale the Hinge Loss takes the following form:

$$L_{D_k} = E_{z \sim q_{data}(x)}[\min(0, -1 + D_k(I_{real}, M, M_{sem}))] + E_{z \sim p(z)}[\min(1, -1 - D_k(I_{real}, M, M_{sem}))], \quad (5)$$

which is then combined to form the full discrimination loss,

$$L_D = L_{D_1} + L_{D_2}. \quad (6)$$

4 Experiments

In Section 3 we described how the SESAME Generator can be used to semantic edit images for addition, manipulation and removal and how we designed the SESAME discriminator to tackle both image editing. To elucidate the merits of our approach we conducted a series of different experiments:

- In order to quantify the performance of our network we follow the data preparation and evaluation steps of Hong *et al.* [12], for generating and removing objects based on a give semantic layout.

Fig. 4. Visual results of addition on Cityscapes: (a) input image (b) edited semantics (c) SESAME (d) Hong *et al.* [12]. Note that Hong *et al.* [12] require the whole semantics while we use only the semantics of the box

- We train our model to permit free form semantic input from users to manipulate scenes and show qualitative results.
- We train our SESAME discriminator along with SPADE Generator for Label to Image Generation.

Implementation Details. For training we are using the Two Time-Scale Update Rule [11] to determine the scale between the learning rate of the generator and the discriminators, with $lr_{gen} = 0.0001$ and $lr_{disc} = 0.0004$. We train for 200 epochs. After 100 we start to linearly decay the learning rates to 0. For our generator losses we multiplied the Feature Matching Loss and Perceptual loss by a factor of 10 before adding them to the adversarial loss. We use the Adam optimizer [21] with coefficient values of $b_1 = 0$ and $b_2 = 0.999$, similar to [35].

Datasets. In line with the literature we conduct experiments:

- **Cityscapes** [8]. The dataset contains 3,000 street-level view images of 50 different cities in Europe for the training set and 500 images for the validation set. The images are accompanied by fine grained information of the per-pixel semantics and instance segmentation with original resolution of the images is 2048×1024 pixels. For addition and removal we down-sample to 1024×512 pixels before patches of 256×256 pixels are extracted. Following Hong *et al.* [12], we choose 10 of the 30 available semantics classes as foreground objects, *e.g.*, *pedestrians*, *cars*, *bicycles*, *etc.*. For generation we resize the image to 512×256 pixels.
- **ADE20K** [60,61] ADE20K has over 20,000 images together with their detailed semantics for 150 different semantic classes. In addition, 2,000 more images are offered for validation. We use this dataset for Image Generation.

Fig. 5. Visual results of addition on ADE20k: (a) input image (b) edited semantics (c) SESAME (d) Hong *et al.* [12]. In this setting we use the full semantic information to guide the editing

- **ADE20K-bedroom** Following Hong *et al.* [12] we experiment on a subset of the ADE20K dataset comprised of bedroom scenes. Similarly to Cityscapes, 31 objects are chosen as foreground objects. In total, we consider 49 semantic categories for training and evaluation.
- **Flickr-Landscapes Datasets [35]** Similar to SPADE [35], we first scrapped 200,000 images from Flickr with only landscape constraint. As our main purpose is to show image editing over significant areas within a landscape, we use a DeepLab-v2 [6] network trained on COCO-Stuff in order to extract images that contain at least 80% pixels of clouds, mountains, water, grass, etc. After post-processing, our curated dataset consists of 7367 training and 500 validation images with their corresponding segmentation for 17 different semantic classes.

Data Pre-processing. Free-form semantic editing is not trivial to achieve. The model can easily overfit on mask shapes used during training. In order to train our free-form semantic editing experiments, we randomly draw a box mask in conjunction with random strokes [55] with 70% chance, otherwise we drop all the pixels belonging to a semantic class of the training image.

For layout driven editing, we extend the data pre-processing scheme introduced by Hong *et al.* [12]. A rectangular area is removed from the input image and we try to inpaint it using the semantic labels. To train the addition operation, they extract the boundary boxes based on the instances of the foreground classes. For the removal sub-task, they randomly choose and remove blocks to train the network to inpaint background classes. While this makes sense for a dataset like ADE20k where the foreground objects can be found anywhere in the pictures, in Cityscapes the foreground objects placement follow certain dis-

Fig. 6. Free-form image manipulation. The user can select a semantic brush and paint over the image to adjust as they see fit

Table 1. Results for the Cityscapes dataset [8]. For the SSIM and mIoU higher is better, while for FID, lower is better. We follow [35] and use the DRN [53,54] for segmenting

Method	Addition				Removal			
	SSIM↑	accu↑	mIoU↑	FID↓	SSIM↑	accu↑	mIoU↑	FID↓
Hong <i>et al.</i> [12]	0.377	83.8%	60.7%	12.11	0.584	83.9%	65.3%	10.34
SESAME	0.410	86.0%	65.3%	11.03	0.797	85.0%	67.6%	7.43

tributions [23]. Thus we only extract a randomly chosen rectangular area if it contains at least a pixel of *ground, road, sidewalk and parking*.

Quantitative Results. For measuring the performance of our network we combine the evaluation approach of previous methods [12,35]. To assess the visual quality of our synthesis we use the Frechet Inception Distance (FID) [11]. Then we compare the mean Intersection over Union (mIoU) and the pixel-accuracy loss between the ground truth semantic layout and the inferred one. For all our experiments, we chose every time the same object in order to maintain consistency. Additionally, in our comparison with Hong *et al.* [12], we compute the Structural Similarity Index (SSIM) [50] between each $\langle I_{real}, I_{out} \rangle$ pair, taking into account only the generated pixels. Naturally, as in the case of Editing only a small percentage of image pixels are changed we expect better results than in the Generation experiments, but also better methods yield larger performance gains when tackling the latter.

Table 2. Quantitative Results for the proposed SESAME in comparison to Hong *et al.* [12] on ADE20K dataset [61,60]. Note that the majority of the pixels used to calculate these metrics are real, thus reducing the actual numerical improvement of SESAME over our baseline

Method	Addition				Removal			
	SSIM↑	accu↑	mIoU↑	FID↓	SSIM↑	accu↑	mIoU↑	FID↓
Hong <i>et al.</i> [12]	0.236	92.5%	33.9%	27.53	0.456	91.7%	40.0%	24.98
SESAME	0.267	93.4%	36.7%	24.88	0.491	92.3%	41.6%	23.30

Table 3. Comparison in number of parameters

Method	Parameters in millions	
	Generator	Discriminator
Hong <i>et al.</i> [12]	190m	5.6m
SESAME	20.5m	11.1m

Our Baselines. For our semantic image editing baseline we are using the work of Hong *et al.* [12]. They introduced a hierarchical model to tackle the task of image editing. In the first stage, they inpaint the semantic classes of an image with a missing region. Then they combine the predicted output with the ground truth and after concatenating the real image with the missing pixels, they use their second stage model to fill the image. Similar to their work, we focus on the *mask to image generation* task and compare our model against their image generator trained on the ground truth labels. Their approach consists of an encoder and a decoder. The encoder has two input streams where the image and the semantics are processed separately and are then *fused* based on the mask of the object location. The result of the fusion is then passed to an image decoder which produces the end result. The generator is trained in conjunction with a PatchGAN discriminator. We use different architectures and largely decrease the number of parameters for the generator and have a larger discriminator as shown in Table 3. However, during inference time only the generator is used. Reduced number of parameters for the generator is clearly beneficial during execution. For Image Generation our Baseline is SPADE [35].

Addition and Removal of Objects To elucidate the ability of our network to perform well both on the addition and the removal part we compare on both tasks separately. The computed metrics for these cases can be found in Tables 1 and 2 for Cityscapes and ADE20k, respectively. For a fair comparison, we showcase a version of our model where a rectangular region of pixels is removed from the input image.

For Cityscapes and ADE20K we use the boundary box (BBox) semantics and the full semantics, respectively. In the visual results we can observe that objects look sharper and their features are more distinctive. Furthermore, as Figures 4

Fig. 7. Label to Image Generation results. For each triplet of images we are showing the semantic layout input (left), generation using PatchGAN (center) and SESAME (right) discriminator on top of SPADE generator

and 5 illustrate, our method generates different patterns for different clothes, and cars in which the windows are not mixed with the rest of the car. Besides our better numerical results, our user study (see Section 4) further illustrates the superiority of our approach. In the case of removal, artifacts of the *BBox* are commonly left in picture by the method of Hong *et al.* [12], whilst in our case this effect is difficult to notice.

Labels to Image Generation The SESAME Discriminator is designed to tackle the shortcomings of the naive concatenation of an image and its semantics label when generating images. We measure the performance on Labels to Image Generation against SPADE [35] using the same generator. The results for Cityscapes and ADE20k datasets can be found on Table 6 and Fig. 7.

Free-Form Semantic Image Editing The user selects a brush of a semantic class and paints over the image. The pixels that are painted over are removed from the image and SESAME is filling the gaps based on the painted semantic guidance. Examples of hand painted masks and corresponding results can be seen on Fig. 6. We should note that *snow* is a different semantic label from *mountain* and we can observe when *draw* with the mountain brush the model learned to

Table 4. We ablate on the semantics availability, the generator architecture and the discriminator architecture for adding objects on street scenes from Cityscapes w.r.t. FID score (lower is better)

Generator	Discriminator			
	Full semantics		BBox semantics	
	Patch	Ours	Patch	Ours
SPADEPix2PixHD	11.92	12.74	12.32	12.66
Ours	11.95	11.64	11.13	11.03

differentiate this from snow. Additionally, the context is very important for the label we want to add: a patch of grass cannot be drawn in the middle of the sky. More results can be found in the supplementary materials.

Ablation Study SESAME incorporates a Generation/Discrimination pair able to edit a scene by only considering the Semantics of regions in the image that the user seeks to edit. In order to showcase the benefits of our approach we ablate the performance of our architecture by varying (a) the generator architecture, (b) the discriminator architecture and (c) the available semantics, by utilizing either the *Full* semantic layout or the semantics of the rectangular region we want to edit, which we refer to as *BBox* Semantics. Regarding the generator architecture, we compare our generator against Pix2PixHD++ [47], with SPADE layers on the decoder part. Regarding discriminator architecture, we compare against the commonly used PatchGAN.

The performance assessment is reported in Table 4. The use of our discriminator over PatchGAN improved the results in almost all cases. However, we observe that while our generator-discriminator combination performs the best, the second combination is without any of our networks. We argue that our proposed method works better together as the large receptive field provided by the dilated convolutions in our generator synergizes well with the highly focused gradient flow coming from our discriminator.

In another series of experiments, we substituted our Semantics merging operation. Instead of applying *Sum Global Pooling*, we experiment with 1) concatenating the two streams of information and 2) calculating their element-wise product resulted in a lower FID score compared to the proposed approach, 11.96 and 12.02 respectively.

User Study We employed Amazon Mechanical Turk⁴ to conduct two experiments for the user study. For each of them, we sampled 100 images from our validation set and asked 20 Turkers: *Which among the images looks more photo-realistic?*

The first experiment presented the Turkers with three options: our method with access to only the *BBox* information and both ours and Hong *et al.* [12]

⁴ <https://www.mturk.com>

Table 5. User Study Results: *Which image is the most photo-realistic?* The first study invited the users to choose between Hong *et al.* [12] with full semantics information and ours with full and bbox semantics, respectively. The second study invited to choose between the results produced by our SESAME and the PatchGAN discriminator, for different availability of semantics

User Study I		User Study II		
Setting	Preference [%]	Discriminator:	SESAME	PatchGAN
Hong <i>et al.</i> [12]	22.50	FullContext	56.67	43.33
Ours w Full	35.83	BBoxContext	61.04	38.96
Ours w BBox	41.67			

Table 6. Layout to image generation results. For the SSIM, mIoU and accu, higher is better, while for FID, lower is better

Method	Cityscapes			ADE20k		
	mIoU	accu	FID	mIoU	accu	FID
Pix2PixHD	58.3	81.4	95.0	22.4	68.38	81.8
SPADE	62.3	81.9	71.8	38.5	81.9	33.9
Ours	66.0	82.5	54.2	49.0	85.5	31.9

model using the *Full* Semantics. As shown in Table 5 the results of our SESAME approach were clearly preferred by the users over the results of our baseline. Moreover, in agreement with our quantitative analysis, the proposed scaling scheme in our discriminator benefits from less irrelevant semantic information. Another group of settings compared the results when the PatchGAN is used instead of our SESAME discriminator. The results consistently show that the independent processing of the semantic information leads to better perceptual quality of the results; they are picked more often by the human subjects.

5 Conclusion

In this work, we introduce SESAME a novel method for semantic image editing covering the complete spectrum of adding, manipulating and erasing operations. Our generator is capable of manipulating an image by only conditioning on the semantics of the regions the user seeks to edit, namely without requiring the information about the full layout. Our discriminator processes the semantic and image information in separate streams and overcomes the limitations of the concatenating approach inherent in PatchGAN. SESAME produces state-of-the-art results on the tasks of (a) semantic image manipulation and (b) layout to image generation and permits the user to edit an image by intuitively painting over it. As a future research direction, we plan to extend this work on image generation conditioned on other types of information, *e.g.*, scene graphs, could also benefit from our two-stream discriminator. We refer to supplementary material for more

details. We will open-source the code and the models under the repository name **OpenSESAME**.

Acknowledgements This work was partly supported by CSEM, ETH Zurich Fund (OK) and by Huawei, Amazon AWS and Nvidia GPU grants. We are grateful to Despoina Paschalidou, Siavash Bigdeli and Danda Pani Paudel for fruitful discussions. We also thank Gene Kogan for providing guidance on how to prepare the Flickr Landscapes Dataset.

References

1. Antoniou, A., Storkey, A., Edwards, H.: Data augmentation generative adversarial networks. arXiv preprint arXiv:1711.04340 (2017)
2. Ashual, O., Wolf, L.: Specifying object attributes and relations in interactive scene generation. In: Proceedings of the International Conference Computer Vision (ICCV) (2019)
3. Bau, D., Strobel, H., Peebles, W., Wulff, J., Zhou, B., Zhu, J., Torralba, A.: Semantic photo manipulation with a generative image prior. ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH) **38**(4) (2019)
4. Bau, D., Zhu, J.Y., Wulff, J., Peebles, W., Strobel, H., Zhou, B., Torralba, A.: Seeing what a gan cannot generate. In: Proceedings of the International Conference Computer Vision (ICCV) (2019)
5. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: Proceedings of the International Conference on Learning Representations (ICLR) (2019)
6. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence **40**(4), 834–848 (2017)
7. Cognex: Visionpro vidi: Deep learning-based software for industrial image analysis. <https://www.cognex.com/products/machine-vision/vision-software/visionpro-vidi>, accessed: 2019-03-05
8. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
9. Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., Greenspan, H.: Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification. Neurocomputing **321**, 321331 (Dec 2018)
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems. pp. 2672–2680 (2014)
11. 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: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. pp. 6626–6637 (2017)

12. Hong, S., Yan, X., Huang, T.E., Lee, H.: Learning hierarchical semantic image manipulation through structured representations. In: *Advances in Neural Information Processing Systems*. pp. 2713–2723 (2018)
13. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)* **36**(4), 1–14 (2017)
14. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
15. Janai, J., Gney, F., Behl, A., Geiger, A.: Computer vision for autonomous vehicles: Problems, datasets and state of the art. *arXiv preprint arXiv:1704.05519* (2017)
16. Jo, Y., Park, J.: Sc-fegan: Face editing generative adversarial network with user’s sketch and color. In: *Proceedings of the International Conference Computer Vision (ICCV)* (2019)
17. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2016)
18. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
19. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018)
20. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
21. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014)
22. Lee, C.H., Liu, Z., Wu, L., Luo, P.: Maskgan: Towards diverse and interactive facial image manipulation. *arXiv preprint arXiv:1907.11922* (2019)
23. Lee, D., Liu, S., Gu, J., Liu, M.Y., Yang, M.H., Kautz, J.: Context-aware synthesis and placement of object instances. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. pp. 10393–10403 (2018)
24. Lim, J.H., Ye, J.C.: Geometric gan. *arXiv preprint arXiv:1705.02894* (2017)
25. Liu, G., Reda, F.A., Shih, K.J., Wang, T.C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 85–100 (2018)
26. Liu, X., Yin, G., Shao, J., Wang, X., Li, H.: Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In: *Advances in Neural Information Processing Systems* (2019)
27. Lugmayr, A., Danelljan, M., Timofte, R., Fritsche, M., Gu, S., Purohit, K., Kandula, P., Suin, M., Rajagopalan, A., Joon, N.H., et al.: Aim 2019 challenge on real-world image super-resolution: Methods and results. In: *Proceedings of the International Conference Computer Vision (ICCV), Advances in Image Manipulation Workshop* (2019)
28. Mao, Q., Lee, H.Y., Tseng, H.Y., Ma, S., Yang, M.H.: Mode seeking generative adversarial networks for diverse image synthesis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1429–1437 (2019)
29. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014)

30. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018)
31. Miyato, T., Koyama, M.: cGANs with projection discriminator. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018)
32. Nazeri, K., Ng, E., Joseph, T., Qureshi, F.Z., Ebrahimi, M.: Edgeconnect: Generative image inpainting with adversarial edge learning. arXiv preprint arXiv:1901.00212 (2019)
33. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier gans. arXiv preprint arXiv:1610.09585 (2016)
34. Ouyang, X., Cheng, Y., Jiang, Y., Li, C.L., Zhou, P.: Pedestrian-synthesis-gan: Generating pedestrian data in real scene and beyond. arXiv preprint arXiv:1804.02047 (2018)
35. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
36. Photoshop: version 21.1.0. Adobe Inc., San Jose, California, U.S. (2020)
37. Portenier, T., Hu, Q., Szabó, A., Bigdeli, S.A., Favaro, P., Zwicker, M.: Faceshop: Deep sketch-based face image editing. ACM Trans. Graph. **37**(4) (2018)
38. Program, G.I.M.: version 2.10.18. The GIMP Development Team (2018)
39. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text-to-image synthesis. In: Proceedings of the International Conference on Machine Learning (ICML) (2016)
40. Romero, A., Arbeláez, P., Van Gool, L., Timofte, R.: Smit: Stochastic multi-label image-to-image translation. In: Proceedings of the International Conference Computer Vision (ICCV), Workshops (2019)
41. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training gans. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems. pp. 2234–2242 (2016)
42. Sangkloy, P., Lu, J., Fang, C., Yu, F., Hays, J.: Scribbler: Controlling deep image synthesis with sketch and color. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
43. Shaham, T.R., Dekel, T., Michaeli, T.: Singan: Learning a generative model from a single natural image. In: Proceedings of the International Conference Computer Vision (ICCV) (2019)
44. Shetty, R., Fritz, M., Schiele, B.: Adversarial scene editing: Automatic object removal from weak supervision. In: Bengio, S., Wallach, H., Larochelle, H., Graumann, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. pp. 7716–7726. Curran Associates, Montréal, Canada (2018)
45. Shetty, R., Schiele, B., Fritz, M.: Not using the car to see the sidewalk: Quantifying and controlling the effects of context in classification and segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2019)
46. Tran, D., Ranganath, R., Blei, D.: Hierarchical implicit models and likelihood-free variational inference. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. pp. 5523–5533 (2017)
47. 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:

- Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
48. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 0–0 (2018)
  49. Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2018)
  50. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., et al.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
  51. Wu, S., Lin, S., Wu, W., Azzam, M., Wong, H.S.: Semi-supervised pedestrian instance synthesis and detection with mutual reinforcement. In: Proceedings of the International Conference Computer Vision (ICCV) (2019)
  52. Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. *arXiv preprint arXiv:1711.10485* (2017)
  53. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: Proceedings of the International Conference on Learning Representations (ICLR) (2016)
  54. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
  55. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. *arXiv preprint arXiv:1806.03589* (2018)
  56. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5505–5514 (2018)
  57. Zhang, H., Goodfellow, I.J., Metaxas, D.N., Odena, A.: Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318* (2018)
  58. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of the International Conference Computer Vision (ICCV) (2017)
  59. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *arXiv preprint arXiv:1710.10916* (2017)
  60. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. *arXiv preprint arXiv:1608.05442* (2016)
  61. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)

A Supplementary Material

Model Architectures The detailed architectures of the SESAME generator and discriminator are depicted in Tables 7 and 8 respectively. Please note that for the discriminator we use this architecture twice, once for each scale.

Table 7. SESAME generator architecture. We depict the number of **F**ilters, the **K**ernel size, the **S**tride and the **D**ilation factor

Layer	Normalization	Activation
ConvBlock F = 64, K = 7, S = 1, D = 1	Instance	ReLU
ConvBlock F = 128, K = 3, S = 2, D = 1	Instance	ReLU
ConvBlock F = 256, K = 3, S = 2, D = 1	Instance	ReLU
ResBlock F = 256, K = 3, S = 1, D = 1	Instance	ReLU
ResBlock F = 256, K = 3, S = 1, D = 2	Instance	ReLU
ResBlock F = 256, K = 3, S = 1, D = 2	Instance	ReLU
ResBlock F = 256, K = 3, S = 1, D = 2	Instance	ReLU
ResBlock F = 256, K = 3, S = 1, D = 2	SPADE	LeakyReLU(0.02)
ResBlock F = 256, K = 3, S = 1, D = 2	SPADE	LeakyReLU(0.02)
ResBlock F = 256, K = 3, S = 1, D = 2	SPADE	LeakyReLU(0.02)
ResBlock F = 256, K = 3, S = 1, D = 2	SPADE	LeakyReLU(0.02)
ResBlock F = 256, K = 3, S = 1, D = 1	SPADE	LeakyReLU(0.02)
<i>Nearest Neighbour Upsampling $\times 2$</i>	-	-
ResBlock F = 128, K = 3, S = 1, D = 1	SPADE	LeakyReLU(0.02)
<i>Nearest Neighbour Upsampling $\times 2$</i>	-	-
ResBlock F = 64, K = 3, S = 1, D = 1	SPADE	LeakyReLU(0.02)
ConvBlock F = 3, K = 3, S = 1, D = 1	-	TanH

Replacement of objects Apart from adding and removing objects, SESAME can also be used to replace an instance of an object, given that we know its class and its outline. SESAME can be utilized in this manner for dataset augmentation. We conduct experiments on replacing objects in street scenes of Cityscapes and we ablate on the usage of our SESAME discriminator against the PatchGAN. In Table 9 we measure the FID score of the image results, the SSIM of the generated regions and we also devoted a part of our user study, described in Section 4 of the main paper, to test which of the two discriminators produces the most *photo-realistic* results. Visual results can be found in Figure 8.

Visual Results We show more edited and generated images produced by our method:

- Figure 8 contains visual results of our ablation analysis on various access levels of semantics and different discriminators.
- Figure 9 contains visual results for removing objects under different configurations.

Table 8. SESAME discriminator architecture **per scale**. We depict the number of **F**ilters, the **K**ernel size, the **S**tride and the **D**ilation factor

Layer	Normalization	Activation
Image Stream		
ConvBlock F = 64, K = 4, S = 2, D = 1	-	LeakyReLU(0.02)
ConvBlock F = 128, K = 4, S = 2, D = 1	SpectralInstance	LeakyReLU(0.02)
ConvBlock F = 256, K = 4, S = 2, D = 1	SpectralInstance	LeakyReLU(0.02)
ConvBlock F = 512, K = 4, S = 1, D = 1	SpectralInstance	LeakyReLU(0.02)
Semantics Stream		
ConvBlock F = 64, K = 4, S = 2, D = 1	-	LeakyReLU(0.02)
ConvBlock F = 128, K = 4, S = 2, D = 1	SpectralInstance	LeakyReLU(0.02)
ConvBlock F = 256, K = 4, S = 2, D = 1	SpectralInstance	LeakyReLU(0.02)
ConvBlock F = 512, K = 4, S = 1, D = 1	SpectralInstance	LeakyReLU(0.02)
<i>Sum Global Pooling</i>	-	-
Common Head		
ConvBlock F = 1, K = 4, S = 1, D = 1	-	-

Table 9. Object Replacement - Cityscapes. We show the performance of our SESAME model with our discriminator and the PatchGAN[14,35] discriminator as well as the percentage of user answers to the question: *Which image looks more photo-realistic?*

Discriminator	SSIM $\uparrow$	FID $\downarrow$	User Preference(%)
PatchGAN	0.390	10.63	45.03
SESAME	0.433	9.3	54.97

- Figure 10 shows results for editing ADE20k-Bedroom scenes[61,60].
- Figure 11 showcases examples of free-from semantic editing.
- On Figures 12 and 13 we can observe layout to image generation results for Cityscapes and ADE20k.

Label to Image Generation: Comparison with CC-FPSE: Concurrently to our work, Liu *et al.* [26] developed an approach to tackle image generation conditioned on semantic layouts. They propose a generator architecture that learns to predict convolutional kernel weights conditioned on the semantic input. Moreover, they propose a feature pyramid semantics-embedding (FPSE) discriminator using a U-Net architecture. Each upsampling layer outputs two per-patch score maps, one trying to measure the *realness* and one to gauge the *semantic matching* with the labels; the later is derived after a patch-wise inner product operation with the down-sampled semantic embeddings.

Their FPSE discriminator, while it also addresses the shortcomings of previous models, follows a different approach to our SESAME discriminator. Although they similarly aim to short-circuit the guidance of the semantic labels to the discrimination, they choose to do so by embedding the patch with a 1×1 convolution and down-sampling via average pooling the semantic layout to match the size of their image processing pipeline. As we explained in the main paper, the

Fig. 8. Visual results for addition on Cityscapes[8]. We ablate on: (a) using the Full context, using only the labels of the rectangular areas to be edited and only replacing an object given its mask (b) generations due to training with the PatchGAN Discriminator and SESAME. Finally, we show the results produced by the method of Hong *et al.* [12], using the full semantics information

Fig. 9. Visual results for removal on Cityscapes[8]. We ablate on: (a) using the Full context and using only the labels of the rectangular areas to be edited (b) generations due to training with the PatchGAN Discriminator and SESAME. In the first row we show the results produced by the method of Hong *et al.* [12] , using the full semantics information

Fig. 10. Visual results for editing Bedroom scenes from ADE20K dataset. Here we are using the Full semantic information to alter the gray area

Fig. 11. Examples of free form editing using semantic brushes. Note that in the last, row instead of the *mountain* brush, we used *rock* and *tree*, on the left and right side respectively. The model fails to correctly depict a *semantic concept* out of context or with an unexpected shape

Fig. 12. Cityscapes[8]: Visual results for image generation conditioned on Semantic Labels. We showcase the results using the generator from SPADE[35] with the PatchGAN Discriminator (SPADE) and ours (SESAME)

Fig. 13. Ade20k[61,60]: Visual results for image generation conditioned on Semantic Labels. We showcase the results using the generator from SPADE[35] with the PatchGAN Discriminator(SPADE) and ours(SESAME)

Table 10. Layout to image generation results. For the mIoU and accu, higher is better, while for FID, lower is better.

Method	Cityscapes			ADE20k		
	mIoU↑	accu↑	FID↓	mIoU↑	accu↑	FID↓
Pix2PixHD[47]	58.3	81.4	95.0	<i>20.3</i>	<i>69.2</i>	81.8
SPADE[35]	62.3	81.9	71.8	38.5	<i>79.9</i>	33.9
CC-FPSE[26]	66.5	82.3	54.3	43.7	82.9	31.7
SESAME	66.0	82.5	54.2	49.0	85.5	31.9

receptive field of a patch may contain a multitude of different semantic classes with a variety of compositions. Trivially down-sampling the semantic label can result into loss of information. Thus, we proposed a dedicated part of the discriminator to derive a meaningful representation for such an intricate semantic patch. Moreover, our model independently processes the semantic information for each scale.

In this work, we develop a solution to tackle the problem of semantic image editing and we apply our SESAME discriminator on top of SPADE’s generator to produce our *label to image generation* results. However, Liu *et al.* [26] focus solely on this task and achieve results similar to ours when incorporating all components to their model, as seen in Table 10. Note that their performance deteriorates when using SPADE’s generator, a setup comparable to our approach.

The values in italics of Table 10 are corrections of the ones reported on the main paper.