Controlling Style and Semantics in Weakly-Supervised Image Generation

Dario Pavllo, Aurelien Lucchi, and Thomas Hofmann

Department of Computer Science, ETH Zurich



Fig. 1: Our approach enables control over the style of a scene and its objects via highlevel attributes or textual descriptions. It also allows for image manipulation through the mask, including moving, deleting, or adding object instances. The decomposition of the background and foreground (top-right corner) facilitates local changes in a scene.

Abstract. We propose a weakly-supervised approach for conditional image generation of complex scenes where a user has fine control over objects appearing in the scene. We exploit sparse semantic maps to control object shapes and classes, as well as textual descriptions or attributes to control both local and global style. In order to condition our model on textual descriptions, we introduce a semantic attention module whose computational cost is independent of the image resolution. To further augment the controllability of the scene, we propose a two-step generation scheme that decomposes background and foreground. The label maps used to train our model are produced by a large-vocabulary object detector, which enables access to unlabeled data and provides structured instance information. In such a setting, we report better FID scores compared to fully-supervised settings where the model is trained on groundtruth semantic maps. We also showcase the ability of our model to manipulate a scene on complex datasets such as COCO and Visual Genome.

1 Introduction

Deep generative models such as VAEs [23] and GANs [9] have made it possible to learn complex distributions over various types of data, including images and text. For images, recent technical advances [1,13,20,28,29,49] have enabled GANs to

produce realistically-looking images for a large number of classes. However, these models often do not provide high-level control over image characteristics such as appearance, shape, texture, or color, and they fail to accurately model multiple (or compound) objects in a scene, thus limiting their practical applications. A related line of research aims at disentangling factors of variation [21]. While these approaches can produce images with varied styles by injecting noise at different levels, the style factors are learned without any oversight, leaving the user with a loose handle on the generation process. Furthermore, their applicability has only been demonstrated for single-domain images (e.g. faces, cars, or birds). Some conditional approaches allow users to control the style of an image using either attributes [12,46] or natural language [45,50,51], but again, these methods only show compelling results on single-domain datasets.

One key aspect in generative modeling is the amount of required semantic information: i) weak conditioning (e.g. a sentence that describes a scene) makes the task underconstrained and harder to learn, potentially resulting in incoherent images on complex datasets. On the other hand, ii) rich semantic information (e.g. full segmentation masks) yields the best generative quality, but requires more effort from an artist or annotator. The applications of such richly-conditioned models are numerous, including art, animation, image manipulation, and realistic texturing of video games. Existing works in this category [4, 17, 31, 32, 44] typically require hand-labeled segmentation masks with per-pixel class annotations. Unfortunately, this is not flexible enough for downstream applications such as image manipulation, where the artist is faced with the burden of modifying the semantic mask coherently. Common transformations such as moving, deleting, or replacing an object require instance information (usually not available) and a strategy for infilling the background. Moreover, these models present little-to-no high-level control over the style of an image and its objects.

Our work combines the merits of both weak conditioning and strong semantic information, by relying on both mask-based generation – using a variant we call sparse masks – and text-based generation – which can be used to control the style of the objects contained in the scene as well as its global aspects. Fig. 1 conceptualizes our idea. Our approach uses a large-vocabulary object detector to obtain annotations, which are then used to train a generative model in a weaklysupervised fashion. The input masks are sparse and retain instance information - making them easy to manipulate - and can be inferred from images or videos in-the-wild. We additionally contribute a conditioning scheme for controlling the style of the scene and its instances, either using high-level attributes or natural language with an attention mechanism. Unlike prior approaches, our attention model is applied directly to semantic maps (making it easily interpretable) and its computational cost does not depend on the image resolution, enabling its use in high-resolution settings. This conditioning module is general enough to be plugged into existing architectures. We also tackle another issue of existing generative models: local changes made to an object (such as moving or deleting) can affect the scene globally due to the learned correlations between classes. While these entangled representations improve scene coherence, they do not allow the user to modify a local part of a scene without affecting the rest. To this end, our approach relies on a multi-step generation process where we first generate a *background image* and then we generate foreground objects conditioned on the former. The background can be frozen while manipulating foreground objects.

Finally, we evaluate our approach on COCO [2,5,26] and Visual Genome [25], and show that our weakly-supervised setting can achieve better FID scores [13] than fully-supervised counterparts trained on ground-truth masks, and weaklysupervised counterparts where the model is trained on dense maps obtained from an off-the-shelf semantic segmentation model, while being more controllable and scalable to large unlabeled datasets. We show that this holds both in presence and in absence of style control.

Code is available at https://github.com/dariopavllo/style-semantics.

2 Related work

The recent success of GANs has triggered interest for conditional image synthesis from categorical labels [1,28,29,49], text [33,45,50,51], semantic maps [17,31,44], and conditioning images from other domains [17,53].

Image generation from semantic maps. In this setting, a semantic segmentation map is translated into a natural image. Non-adversarial approaches are typically based on perceptual losses [4,32], whereas GAN architectures are based on patch-based discriminators [17], progressive growing [20, 44], and conditional batch normalization where the semantic map is fed to the model at different resolutions [31]. Similarly to other state-of-the-art methods, our work is also based on this paradigm. Most approaches are trained on hand-labeled masks (limiting their application in the wild), but [31] shows one example where the model is weakly supervised on masks inferred using a semantic segmentation model [3]. Our model is also weakly supervised, but instead of a semantic segmentation model we use an object detector - which allows us to maintain instance information during manipulations, and results in *sparse masks*. While early work focused on class semantics, recent methods support some degree of style control. E.g. [44] trains an instance autoencoder and allows the user to choose a latent code from among a set of modes, whereas [31] trains a VAE to control the global style of a generated image by copying the style of a guide image. Both these methods, however, do not provide fine-grained style control (e.g. changing the color of an object to *red*). Another recent trend consists in generating images from structured layouts, which are transformed into semantic maps as an intermediate step to facilitate the task. In this regard, there is work on generation from bounding-box layouts [14, 15, 40, 52] and scene graphs [18]. Although these approaches tackle a harder task, they generate low-resolution images and are not directly relatable to our work, which tackles controllability among other aspects. Semantic control. Existing approaches do not allow for easy manipulation of the semantic map because they present no interface for encoding existing images. In principle, it is possible to train a weakly-supervised model on maps inferred from a semantic segmentation model, as [31] does for landscapes. However, as



Fig. 2: Left: when manipulating a ground-truth mask (e.g. deleting one bus), one is left with the problem of infilling the background which is prone to ambiguities (e.g. selecting a new class as either road or building). Furthermore, in existing models, local changes affect the scene globally due to learned correlations. Middle: in the wild, ground-truth masks are not available (neither are instance maps). One can infer maps using a semantic segmentation model, but these are often noisy and lack instance information (in the example above, we observe that the two buses are merged). Right: our weakly-supervised sparse mask setting, which combines fine-detailed masks with instance information. The two-step decomposition ensures that changes are localized.

we show in sec. 4.2, the results in this setting are notably worse than fullysupervised baselines. Furthermore, manipulations are still challenging because instance information is not available. Since the label masks are *dense*, even simple transformations such as deleting or moving an object would create holes in the semantic map that need to be adjusted by the artist (Fig. 2). *Dense* masks also make the task too constrained with respect to background aspects of the scene (e.g. sky, land, weather), which leaves less room for style control. Semantic control can also be framed as an unpaired image-to-image translation task [30], but this requires ground-truth masks for both source and target instances, and can only translate between two classes.

Text-based generation. Some recent models condition the generative process on text data. These are often based on autoregressive architectures [34] and GANs [33, 45, 50, 51]. Learning to generate images from text using GANs is known to be difficult due to the task being unconstrained. In order to ease the training process, [50, 51] propose a two-stage architecture named *StackGAN*. To avoid the instability associated with training a language model jointly with a GAN, they use a pretrained sentence encoder [24] that encodes a caption into a fixed-length vector which is then fed to the model. More advanced architectures such as AttnGAN [45] use an attention mechanism which we discuss in one of the next paragraphs. These approaches show interesting results on single-domain datasets (birds, flowers, etc.) but are less effective on complex datasets such as COCO [26] due to the intrinsic difficulty of generating coherent scenes from text alone. Some works [19, 48] have demonstrated that generative models can benefit from taking as input multiple diverse textual descriptions per image. Finally, we are not aware of any prior work that conditions the generative process on both text and semantic maps (our setting).

Multi-step generation. Approaches such as [38, 47] aim at disentangling background and foreground generation. While fully-unsupervised disentanglement is provably impossible [27], it is still achievable through some form of inductive bias – either in the model architecture or in the loss function. While [47] uses spatial transformers to achieve separation, [38] uses object bounding boxes. Both methods show compelling results on single-domain datasets that depict a centered object, but are not directly applicable to more challenging datasets. For composite scenes, [42] generates foreground objects sequentially to counteract merging effects. In our work, we are not interested in full disentanglement (i.e. we do not assume independence between background and foreground), but merely in separating the two steps while keeping them interpretable. Our model still exploits correlations among classes to maximize visual quality, and is applied to datasets with complex scenes. Finally, there has also been work on interactive generation using dialogue [6, 8, 36].

Attention models in GANs. For unconditional models (or models conditioned on simple class labels), self-attention GANs [1, 49] use visual-visual attention to improve spatial coherence. For generation from text, [45] employ sentence-visual attention coupled with an LSTM encoder, but only in the generator. In the discriminator, the caption is enforced through a supervised loss based on features extracted from a pretrained Inception [41] network. We introduce a new form of attention (*sentence-semantic*) which is applied to semantic maps instead of convolutional feature maps, and whose computational cost is independent of the image resolution. It is applied both to the generator and the discriminator, and on the sentence side it features a transformer-based [43] encoder.

3 Approach

3.1 Framework

Our main interest is conditional image generation of complex scenes where a user has fine control over the objects appearing in the scene. Prior work has focused on generating objects from ground-truth masks [17, 31, 44, 53] or on generating outdoor scenes based on simple hand-drawn masks [31]. While the former approach requires a significant labeling effort, the latter is not directly suitable for complex datasets such as COCO-Stuff [2], whose images consist of a large number of classes with complex (hard to draw) shapes. We address these problems by introducing a new model that is conditioned on sparse masks – to control object shapes and classes - and on text/attributes to control style and textures. This gives the ability to a user to produce scenes through a variety of image manipulations (such as moving, scaling or deleting an instance, adding an instance from another image or from a database of shapes) as well as style manipulations controlled using either high-level attributes on individual instances (e.g. red, green, wet, shiny) or using text that refers to objects as well as global context (e.g. "a red car at night"). In the latter case, visual-textual correlations are not explicitly defined but are learned in an unsupervised way.

Sparse masks. Instead of training a model on precise segmentation masks as in [17, 31, 44], we use a mask generated automatically from a large-vocabulary object detector. Compared to a weakly-supervised setting based on semantic segmentation, this process introduces less artifacts (see Appendix A.4 in the supplementary material) and has the benefit of providing information about each instance (which may not always be available otherwise), including parts of objects which would require significant manual effort to label in a new dataset. In general, our set of classes comprises countable objects (person, car, etc.), parts of objects (light, window, door, etc.), as well as uncountable classes (grass, water, snow), which are typically referred to as "stuff" in the COCO terminology [2]. For the latter category, an object detector can still provide useful sparse information about the background, while keeping the model autonomous to fill-in the gaps. We describe the details of our object detection setup in sec. 4.1.

Two-step generation. In the absence of constraints, conditional models learn class correlations observed in the training data. For instance, while dogs typically stand on green grass, zebras stand on yellow grass. While this feature is useful for maximizing scene coherence, it is undesirable when only a local change in the image is wanted. We observed similar global effects on other local transformations, such as moving an object or changing its attributes, and generally speaking, small perturbations of the input can result in large variations of the output. We show a few examples in the Appendix A.4. To tackle this issue, we propose a variant of our architecture which we call *two-step* model and which consists of two concatenated generators (Fig. 3, right). The first step (generator G_1) is responsible for generating a *background* image, whereas the second step (generator G_2) generates a *foreground* image conditioned on the background image. The definition of what constitutes background and foreground is arbitrary: our choice is to separate by class: static/uncountable objects (e.g. buildings, roads, grass, and other surfaces) are assigned to *background*, and moving/countable objects are assigned to *foreground*. Some classes can switch roles depending on the parent class, e.g. window is background by default, but it becomes foreground if it is a child of a foreground object such as a car.

When applying a local transformation to a foreground object, the background can conveniently be frozen to avoid global changes. As a side benefit, this also results in a lower computational cost to regenerate an image. Unlike work on disentanglement [38, 47] which enforces that the background is independent of the foreground without necessarily optimizing for visual quality, our goal is to enforce separation while maximizing qualitative results. In our setting, G_1 is exposed to both background and foreground objects, but its architecture is designed in a way that foreground information is not rendered, but only used to induce a bias in the background (see sec. 3.2).

Attributes. Our method allows the user to control the style of individual instances using high-level attributes. These attributes refer to appearance factors such as colors (e.g. white, black, red), materials (wood, glass), and even modifiers that are specific to classes (leafless, snowy), but not shape or size, since these two are determined by the mask. An object can also combine multiple attributes (e.g. black and white) or have none – in this case, the generator would pick a predefined mode. This setup gives the user a lot of flexibility to manipulate a scene, since the attributes need not be specified for every object.

Captions. Alternatively, one can consider conditioning style using natural language. This has the benefit of being more expressive, and allows the user to control global aspects of the scene (e.g. time of the day, weather, landscape) in addition to instance-specific aspects. While this kind of conditioning is harder to learn than plain attributes, in sec. 3.2 we introduce a new attention model that shows compelling results without excessively increasing the model complexity.

3.2 Architecture.

We design our conditioning mechanisms to have sufficient generality to be attached to existing conditional generative models. In our experiments, we choose SPADE [31] as the backbone for our conditioning modules, which to our knowledge represents the state of the art. As in [31], we use a multi-scale discriminator [44], a perceptual loss in the generator using a pretrained VGG network [37], and a feature matching loss in the discriminator [44].

One-step model. Since this model (Fig. 3, left) serves as a baseline, we keep its backbone as close as possible to the reference model of [31]. We propose to insert the required information about attributes/captions in this architecture by modifying the input layer and the conditional batch normalization layers of the generator, which is where semantic information is fed to the model. We name these S-blocks (short for *semantic-style* block).

Semantic-style block. For class semantics, the input sparse mark is fed to a pixel-wise embedding layer to convert categorical labels into 64D embeddings (including the empty space, which is a special class "no class"). To add style information, we optionally concatenate another 64D representation to the class embedding (pixel-wise); we explain how we derive this representation in the next two paragraphs. The resulting feature map is convolved with a 3×3 kernel, passed through a ReLU non-linearity and convolved again to produce two feature maps γ and β , respectively, the conditional batch normalization gain and bias. The



Fig. 3: Left: One-step model. Right: two-step model. The background generator G_1 takes as input a *background mask* (processed by *S*-blocks) and the full mask (processed by S_{avg} -blocks, where positional information is removed). The foreground generator takes as input the output of G_1 and a *foreground mask*. Finally, the two outputs are alpha-blended. For convenience, we do not show attributes/text in this figure.



Fig. 4: Left: Conditioning block with attributes. Class and attribute embeddings are concatenated and processed to generate the conditional batch normalization gain and bias. In the attribute mask, embeddings take the contour of the instance to which they refer. In G_1 of the two-step model, where S and S_{avg} are both used, the embedding weights are shared. Right: Attention mechanism for conditioning style via text. The sentence (of length n = 7 including delimiters) is fed to a pretrained attention encoder, and each token is transformed into a key and a value using two trainable linear layers. The queries are learned for each class, and the attention yields a set of contextualized class embeddings that are concatenated to the regular semantic embeddings.

normalization is then computed as $\mathbf{y} = BN(\mathbf{x}) \odot (1 + \gamma) + \beta$, where $BN(\mathbf{x})$ is the parameter-free batch normalization. The last step is related to [31] and other architectures based on conditional batch normalization. Unlike [31], however, we do not use 3×3 convolutions on one-hot representations in the input layer. This allows us to scale to a larger number of classes without significantly increasing the number of parameters. We apply the same principle to the discriminators. Conditioning on attributes. For attributes, we adopt a bag-of-embeddings approach where we learn a 64D embedding for each possible attribute, and all attribute embeddings assigned to an instance are broadcast to the contour of the instance, summed together, and concatenated to the class embedding. Fig. 4 (left) (S-block) depicts this process. To implement this efficiently, we create a multi-hot attribute mask (1 in the locations corresponding to the attributes assigned to the instance, 0 elsewhere) and feed it through a 1×1 convolutional layer with N_{attr} input channels and 64 output channels. Attribute embeddings are shared among classes and are not class-specific. This helps the model generalize better (e.g. colors such as "white" apply both to vehicles and animals), and we empirically observe that implausible combinations (e.g. leafless person) are simply ignored by the generator without side effects.

Conditioning on text. While previous work has used fixed-length vector representations [50, 51] or one-layer attention models coupled with RNNs [45], the diversity of our scenes led us to use a more powerful encoder entirely based on self-attention [43]. We encode the image caption using a pretrained BERT_{base}

model [7] (110M parameters). It is unreasonable to attach such a model to a GAN and fine-tune it, both due to excessive memory requirements and due to potential instabilities. Instead, we freeze the pretrained model and encode the sentence, extract its hidden representation after the last or second-to-last layer (we compare these in sec. 4.2), and train a custom multi-head attention layer for our task. This paradigm, which is also suggested by [7], has proven successful on a variety of NLP downstream tasks, especially when these involve small datasets or limited vocabularies. Furthermore, instead of storing the language model in memory, we simply pre-compute the sentence representations and cache them.

Next, we describe the design of our trainable attention layer (Fig. 4, right). Our attention mechanism is different from the commonly-used sentence-visual attention [45], where attention is directly applied to convolutional feature maps inside the generator. Instead, we propose a form of sentence-semantic attention which is computationally efficient, interpretable, and modular. It can be concatenated to conditioning layers in the same way as we concatenate attributes. Compared to sentence-visual attention, whose cost is $\mathcal{O}(nd^2)$ (where n is the sentence length and $d \times d$ is the feature map resolution), our method has a cost of $\mathcal{O}(nc)$ (where c is the number of classes), i.e. it is independent of the image resolution. We construct a set of c queries (i.e. one for each class) of size h = 64 (where h is the attention head size). We feed the hidden representations of each token of the sentence to two linear layers, one for the keys and one for the values. Finally, we compute a scaled dot-product attention [43], which yields a set of *c* values. To allow the conditioning block to attend to multiple parts of the sentence, we use 6 or 12 attention heads (ablations in sec. 4.2), whose output values are concatenated and further transformed through a linear layer. This process can be thought of as generating *contextualized class embeddings*, i.e. class embeddings customized according to the sentence. For instance, given a semantic map that depicts a car and the caption "a red car and a person". the query corresponding to the visual class *car* would most likely attend to "red car", and the corresponding value will induce a bias in the model to add redness to the position of the car. Finally, the *contextualized class embeddings* are applied to the semantic mask via pixel-wise matrix multiplication with one-hot vectors. and concatenated to the class embeddings in the same way as attributes. In the current formulation, this approach is unable to differentiate between instances of the same class. We propose a possible mitigation in sec. 5.

Two-step model. It consists of two concatenated generators. G_1 generates the background, i.e. it models $p(x_{bg})$, whereas G_2 generates the foreground conditioned on the background, i.e. $p(x_{fg}|x_{bg})$. One notable difficulty in training such a model is that background images are never observed in the training set (we only observe the final image), therefore we cannot use an intermediate discriminator for G_1 . Instead, we use a single, final discriminator and design the architecture in a way that the gradient of the discriminator (plus auxiliary losses) is redirected to the correct generator. The convolutional nature of G_1 would then ensure that the background image does not contain visible holes. A natural choice is *alpha blending*, which is also used in [38, 47]. G_2 generates an RGB foreground image

plus a transparency mask (*alpha* channel), and the final image is obtained by pasting the foreground onto the background via linear blending:

$$x_{\text{final}} = x_{\text{bg}} \cdot (1 - \alpha_{\text{fg}}) + x_{\text{fg}} \cdot \alpha_{\text{fg}} \tag{1}$$

where x_{final} , x_{bg} , and x_{fg} are RGB images, and α_{fg} is a 1-channel image bounded in [0, 1] by a sigmoid. Readers familiar with highway networks [39] might notice a similarity to this approach in terms of gradients dynamics. If $\alpha_{\text{fg}} = 1$, the gradient is completely redirected to x_{fg} , while if $\alpha_{\text{fg}} = 0$, the gradient is redirected to x_{bg} . This scheme allows us to train both generators in an end-to-end fashion using a single discriminator, and we can also preserve auxiliary losses (e.g. VGG loss) which [31] has shown to be very important for convergence. To incentivize separation between classes as defined in sec. 3.1, we supervise α_{fg} using a binary cross-entropy loss, and decay this term over time (see sec. 4.1).

 G_2 uses the same S-blocks as the ones in the one-step model, but here they take a foreground mask as input (Fig. 3, right). G_1 , on the other hand, must exploit foreground information without rendering it. We therefore devise a further variation of input conditioning that consists of two branches: (i) the first branch (S-block) takes a background mask as input and processes it as usual to produce the batch normalization gain γ and bias β . (ii) the second branch (S_{avg} -block, Fig. 4 left) takes the full mask as input (background plus foreground), processes it, and applies global average pooling to the feature map to remove information about localization. This way, foreground information is only used to bias G_1 and cannot be rendered at precise spatial locations. After pooling, it outputs γ_{avg} and β_{avg} . (iii) The final conditional batch normalization is computed as:

$$\mathbf{y} = BN(\mathbf{x}) \odot (1 + \gamma + \gamma_{avg}) + \beta + \beta_{avg}$$
(2)

Finally, the discriminator D takes the full mask as input (background plus foreground). Note that, if G_1 took the full mask as input without information reduction, it would render visible "holes" in the output image due to gradients never reaching the foreground zones of the mask, which is what we are trying to avoid. The Appendix A.1 provides more details about our architectures, and A.2 shows how G_2 can be used to generate one object at a time to fully disentangle foreground objects from each other (although this is unnecessary in practice).

4 Experiments

For consistency with [31], we always evaluate our model on the COCO-Stuff validation set [2], but we train on a variety of training sets:

COCO-Stuff (COCO2017) [2, 26] contains 118k training images with captions [5]. We train with and without captions. COCO-Stuff extends COCO2017 with ground-truth semantic maps, but for our purposes the two datasets are equivalent since we do not exploit ground-truth masks.

Visual Genome (VG) [25] contains 108k images that partially overlap with COCO ($\approx 50\%$). VG does not have a standard train/test split, therefore we leave



Fig. 5: Left: the larger set of labels in our sparse masks improves fine details. These masks are easy to obtain with a semi-supervised object detector, and would otherwise be too hard to hand-label. **Right:** sparse masks are also easy to sketch by hand.

out 10% of the dataset to use as a validation set (IDs ending with 9), and use the rest as a training set from which we remove images that overlap with the COCO-Stuff validation set. We extract the attributes from the scene graphs.

Visual Genome augmented (VG+) VG augmented with the 123k images from the COCO unlabeled set. The total size is 217k images after removing exact duplicates. The goal is to evaluate how well our method scales to large unlabeled datasets. We train without attributes and without captions.

For all experiments, we evaluate the *Fréchet Inception Distance* (FID) [13] (precise implementation details of the FID in the Appendix A.3). Furthermore, we report our results in sec. 4.2 and provide additional qualitative results in A.4.

4.1 Implementation details

Semantic maps. To construct the input semantic maps, we use the semisupervised implementation of Mask R-CNN [11,35] proposed by [16]. It is trained on bounding boxes from Visual Genome (3000 classes) and segmentation masks from COCO (80 classes), and learns to segment classes for which there are no ground-truth masks. We discard the least frequent classes, and, since some VG concepts overlap (e.g. car, vehicle) leading to spurious detections, we merge these classes and end up with a total of c = 280 classes (plus a special class for "no class"). We set the threshold of the object detector to 0.2, and further refine the predictions by running a class-agnostic non-maximum-suppression (NMS) step on the detections whose mask intersection-over-union (IoU) is greater than 0.7. We also construct a transformation hierarchy to link children to their parents in the semantic map (e.g. headlight of a car) so that they can be manipulated as a whole; further details in the Appendix A.1. We select the 256 most frequent attributes, manually excluding those that refer to shapes (e.g. *short*, *square*). **Training.** We generate images at 256×256 and keep our experimental setting and hyperparameters as close as possible to [31] for a fair comparison. For the two-step model, we provide supervision on the alpha blending mask and decay

two-step model, we provide supervision on the alpha blending mask and decay this loss term over time, observing that the model does not re-entangle background and foreground. This gives G_2 some extra flexibility in drawing details



Fig. 6: Qualitative results (256×256) . **Top-left** and **top-middle**: two-step generation with manipulation of attributes and instances. **Top-right:** manipulating style (both context and instances) via text. **Bottom:** manipulating global style via text.

that are not represented by the mask (reflections, shadows). Hyperparameters and additional training details are specified in the Appendix A.1.

4.2 Results

Quantitative. We show the FID scores for the main experiments in Table 1 (left). While improving FID scores is not the goal of our work, our weaklysupervised sparse mask baseline (#3) interestingly outperforms both the fullysupervised baseline on SPADE [31] (#1) and the weakly-supervised baseline (#2) trained on dense semantic maps. These experiments adopt an identical architecture and training set, no style input, and differ only in the type of input mask. For #2 we obtain the semantic maps from DeepLab-v2 [3], a state-of-theart semantic segmentation model pretrained on COCO-Stuff. Our improvement is partly due to masks better representing fine details (such as windows, doors, lights, wheels) in compound objects, which are not part of the COCO class set. In Fig. 5 (left) we show some examples. Moreover, the experiment on the augmented Visual Genome dataset highlights that our model benefits from extra unlabeled images (#4). Rows #5-9 are trained with style input. In particular, we observe that these outperform the baseline even when they use a two-step architecture (which is more constrained) or are trained on a different training set (VG instead of COCO). Row #6-7 draw their text embeddings from the last BERT layer and adopt 12 attention heads (the default), whereas #5 draws its embeddings from the 2nd-last layer, uses 6 heads, and performs slightly better. Qualitative. In Fig. 6 we show qualitative results as well as examples of manipulations, either through attributes or text. Additional examples can be seen in Table 1: Left: FID scores for the main experiments; lower is better. The first line represents the SPADE baseline [31]. For the models trained on VG, we also report FID scores on our VG validation set. (†) indicates that the model is weakly-supervised, (6h) denotes "6 attention heads", L_{n-1} indicates that the text embeddings are drawn from the second-to-last BERT layer. **Right:** ablation study with extra experiments.

#	Training set	Test set(s)	Type	Mask input	Style input	FID				
1	COCO-train	COCO-val	1-step [31]	Ground truth	None	22.64		In c	In	
2	COCO-train	COCO-val	1-step [†]	Semantic seg.	None	23.97	-	Ref.	Experiment	$FID(\Delta)$
3	COCO-train	COCO-val	1-step [†]	Sparse (ours)	None	20.02	1	#1	COCO "things" only	32.31 (+9.67)
4	VG+ (aug.)	COCO-val/VG-val	1-step [†]	Sparse (ours)	None	18.93/13.23	II	$^{\#6}_{\#6}$	12h, L_n , attr. in D	20.44 (-0.19)
5	COCO-train	COCO-val	1-step [†]	Sparse (ours)	Text (6h. L_{n-1})	19.65	111		12h, L_{n-1}	19.77 (-0.86)
6	COCO _s train	COCO-val	1-step [†]	Sparse (ours)	Text (12h L_{π})	20.63	IV	#6	$6h, L_{n-1}$	19.65 (-0.98)
7	COCO-train	COCO-val	2-step [†]	Sparse (ours)	Text (12h, L_n)	20.60	V	#9	No f.g. info in S_{avg}	25.16(+4.33)
-	VC	COCO-vai	2-300p	Sparse (ours)	10A0 (12H, Dn)	20.04	VI	#9	Attr. randomization	20.64 (-0.19)
8	VG	COCO-vai/vG-vai	1-step	Sparse (ours)	Attributes	21.13/15.12			•	
9	VG	COCO-val/VG-val	2-step [†]	Sparse (ours)	Attributes	20.83/14.88				
_										



Fig. 7: Random styles by sampling attributes from a per-class empirical distribution.

the Appendix A.4, including latent space interpolation [22]. In A.5, we visualize the attention mechanism. Finally, we observe that sketching sparse masks by hand is very practical (Fig. 5, right) and provides an easier interface than dense semantic maps (in which the class of every pixel must be manually specified). The supplementary video (see Appendix A.7) shows how these figures are drawn. Style randomization. Since we represent style explicitly, at inference we can randomize the style of an image by drawing attributes from a per-class empirical distribution. This is depicted in Fig. 7, and has the additional advantage of being interpretable and editable (attributes can be refined manually after sampling). The two-step decomposition also allows users to specify different sampling strategies for the background and foreground; more details in the Appendix A.2. Ablation study. While Table 1 (left) already includes a partial ablation study where we vary input conditioning and some aspects of the attention module. in Table 1 (right) we make this more explicit and include additional experiments. First, we train a model on a sparsified COCO dataset by only keeping the "things" classes and discarding the "stuff" classes. This setting (I) performs significantly worse than #1 (which uses all classes), motivating the use of a large class vocabulary. Next, we ablate conditioning via text (baseline #6, which adopts the default hyperparameters of BERT). In (II), we augment the discriminator with ground-truth attributes to provide a stronger supervision signal for the generator (we take the attributes from Visual Genome for the images that overlap between the two datasets). The improvement is marginal, suggesting

that our model can learn visual-textual correlations without explicit supervision. In (III), we draw the token representations from the second-to-last layer instead of the last, and in (IV) we further reduce the number of attention heads from 12 to 6. Both III and IV result in an improvement of the FID, which justifies the hyperparameters chosen in #5. Finally, we switch to attribute conditioning (baseline #9). In (IV), we remove foreground information at inference from the S_{avg} block of the first generator G_1 (we feed the background mask twice in S and S_{avg}). The FID degrades significantly, suggesting that G_1 effectively exploits foreground information to bias the result. In (V) we show that randomizing style at inference (previous paragraph) is not detrimental to the FID, but in fact seems to be slightly beneficial, probably due to the greater sample diversity.

Robustness and failure cases. Input masks can sometimes be noisy due to spurious object detections on certain classes. Since these are also present at train time, weakly-supervised training leads to some degree of noise robustness, but sometimes the artifacts are visible in the generated images. We show some positive/negative examples in the Appendix Fig. 14. In principle, mask noise can be reduced by using a better object detector. We also observe that our setup tends to work better on outdoor scenes and sometimes struggles with fine geometric details in indoor scenes or photographs shot from a close range.

5 Conclusion

We introduced a weakly-supervised approach for the conditional generation of complex images. The generated scenes can be controlled through various manipulations on the sparse semantic maps, as well as through textual descriptions or attribute labels. Our method enables a high level of semantic/style control while benefiting from improved FID scores. From a qualitative point-of-view, we have demonstrated a wide variety of manipulations that can be applied to an image. Furthermore, our weakly supervised setup opens up opportunities for large-scale training on unlabeled datasets, as well as generation from hand-drawn sketches.

There are several ways one could pursue to further enrich the set of tools used to manipulate the generation process. For instance, the current version of our attention mechanism cannot differentiate between instances belonging to the same class and does not have direct access to positional information. While incorporating such information is beyond the scope of this work, we suggest that this can be achieved by appending a *positional embedding* to the attention queries. In the NLP literature, the latter is often learned according to the position of the word in the sentence [7,43], but images are 2D and therefore do not possess such a natural order. Additionally, this would require captions that are more descriptive than the ones in COCO, which typically focus on actions instead of style. Finally, in order to augment the quality of sparse maps, we would like to train the object detector on a higher-quality, large-vocabulary dataset [10].

Acknowledgments. This work was partly supported by the Swiss National Science Foundation (SNF) and Research Foundation Flanders (FWO), grant #176004. We thank Graham Spinks and Sien Moens for helpful discussions.

15

References

- 1. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: International Conference on Learning Representations (ICLR) (2019)
- 2. Caesar, H., Uijlings, J., Ferrari, V.: COCO-Stuff: Thing and stuff classes in context. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- 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 (TPAMI) 40(4), 834–848 (2018)
- Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: IEEE International Conference on Computer Vision (ICCV). pp. 1511–1520 (2017)
- Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft COCO captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325 (2015)
- Cheng, Y., Gan, Z., Li, Y., Liu, J., Gao, J.: Sequential attention GAN for interactive image editing via dialogue. arXiv preprint arXiv:1812.08352 (2018)
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
- El-Nouby, A., Sharma, S., Schulz, H., Hjelm, D., El Asri, L., Ebrahimi Kahou, S., Bengio, Y., Taylor, G.W.: Tell, draw, and repeat: Generating and modifying images based on continual linguistic instruction. In: IEEE International Conference on Computer Vision (ICCV) (2019)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Neural Information Processing Systems. pp. 2672–2680 (2014)
- Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5356–5364 (2019)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: IEEE International Conference on Computer Vision (ICCV). pp. 2961–2969 (2017)
- 12. He, Z., Zuo, W., Kan, M., Shan, S., Chen, X.: Attgan: Facial attribute editing by only changing what you want. IEEE Transactions on Image Processing (2019)
- 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: Neural Information Processing Systems. pp. 6626–6637 (2017)
- 14. Hinz, T., Heinrich, S., Wermter, S.: Generating multiple objects at spatially distinct locations. In: International Conference on Learning Representations (ICLR) (2019)
- Hong, S., Yang, D., Choi, J., Lee, H.: Inferring semantic layout for hierarchical text-to-image synthesis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7986–7994 (2018)
- Hu, R., Dollár, P., He, K., Darrell, T., Girshick, R.: Learning to segment every thing. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4233–4241 (2018)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1125–1134 (2017)

- 16 D. Pavllo et al.
- Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Joseph, K., Pal, A., Rajanala, S., Balasubramanian, V.N.: C4Synth: Cross-caption cycle-consistent text-to-image synthesis. In: IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 358–366. IEEE (2019)
- Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (ICLR) (2018)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4401–4410 (2019)
- Kilcher, Y., Lucchi, A., Hofmann, T.: Semantic interpolation in implicit models. In: International Conference on Learning Representations (ICLR) (2018)
- 23. Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. In: International Conference on Learning Representations (ICLR) (2014)
- Kiros, R., Zhu, Y., Salakhutdinov, R.R., Zemel, R., Urtasun, R., Torralba, A., Fidler, S.: Skip-thought vectors. In: Neural Information Processing Systems. pp. 3294–3302 (2015)
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. International Journal of Computer Vision (IJCV) **123**(1), 32–73 (2017)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: European Conference on Computer Vision (ECCV). pp. 740–755. Springer (2014)
- 27. Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations. In: International Conference on Machine Learning (ICML) (2019)
- Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: International Conference on Learning Representations (ICLR) (2018)
- 29. Miyato, T., Koyama, M.: cGANs with projection discriminator. In: International Conference on Learning Representations (ICLR) (2018)
- Mo, S., Cho, M., Shin, J.: Instance-aware image-to-image translation. In: International Conference on Learning Representations (ICLR) (2019)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2337–2346 (2019)
- 32. Qi, X., Chen, Q., Jia, J., Koltun, V.: Semi-parametric image synthesis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. arXiv preprint arXiv:1605.05396 (2016)
- Reed, S., van den Oord, A., Kalchbrenner, N., Bapst, V., Botvinick, M., de Freitas, N.: Generating interpretable images with controllable structure. OpenReview (2016)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: Neural Information Processing Systems. pp. 91–99 (2015)
- Sharma, S., Suhubdy, D., Michalski, V., Kahou, S.E., Bengio, Y.: Chatpainter: Improving text to image generation using dialogue. arXiv preprint arXiv:1802.08216 (2018)

- 37. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- Singh, K.K., Ojha, U., Lee, Y.J.: FineGAN: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Srivastava, R.K., Greff, K., Schmidhuber, J.: Highway networks. In: ICML Deep Learning Workshop (2015)
- Sun, W., Wu, T.: Image synthesis from reconfigurable layout and style. In: IEEE International Conference on Computer Vision (ICCV). pp. 10531–10540 (2019)
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2818–2826 (2016)
- 42. Turkoglu, M.O., Thong, W., Spreeuwers, L., Kicanaoglu, B.: A layer-based sequential framework for scene generation with GANs. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8901–8908 (2019)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Neural Information Processing Systems. pp. 5998–6008 (2017)
- 44. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: Highresolution image synthesis and semantic manipulation with conditional GANs. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8798–8807 (2018)
- 45. 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. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1316–1324 (2018)
- 46. Yan, X., Yang, J., Sohn, K., Lee, H.: Attribute2image: Conditional image generation from visual attributes. In: European Conference on Computer Vision (ECCV). pp. 776–791. Springer (2016)
- Yang, J., Kannan, A., Batra, D., Parikh, D.: LR-GAN: layered recursive generative adversarial networks for image generation. In: International Conference on Learning Representations (ICLR) (2017)
- Yin, G., Liu, B., Sheng, L., Yu, N., Wang, X., Shao, J.: Semantics disentangling for text-to-image generation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2327–2336 (2019)
- 49. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: International Conference on Machine Learning (ICML) (2019)
- 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: IEEE International Conference on Computer Vision (ICCV). pp. 5907–5915 (2017)
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stack-GAN++: Realistic image synthesis with stacked generative adversarial networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 41(8), 1947–1962 (2018)
- 52. Zhao, B., Meng, L., Yin, W., Sigal, L.: Image generation from layout. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: IEEE International Conference on Computer Vision (ICCV). pp. 2223–2232 (2017)