

# Panoptic-based Image Synthesis

Aysegul Dundar

Karan Sapra

Guilin Liu

Andrew Tao

Bryan Catanzaro

NVIDIA Corporation

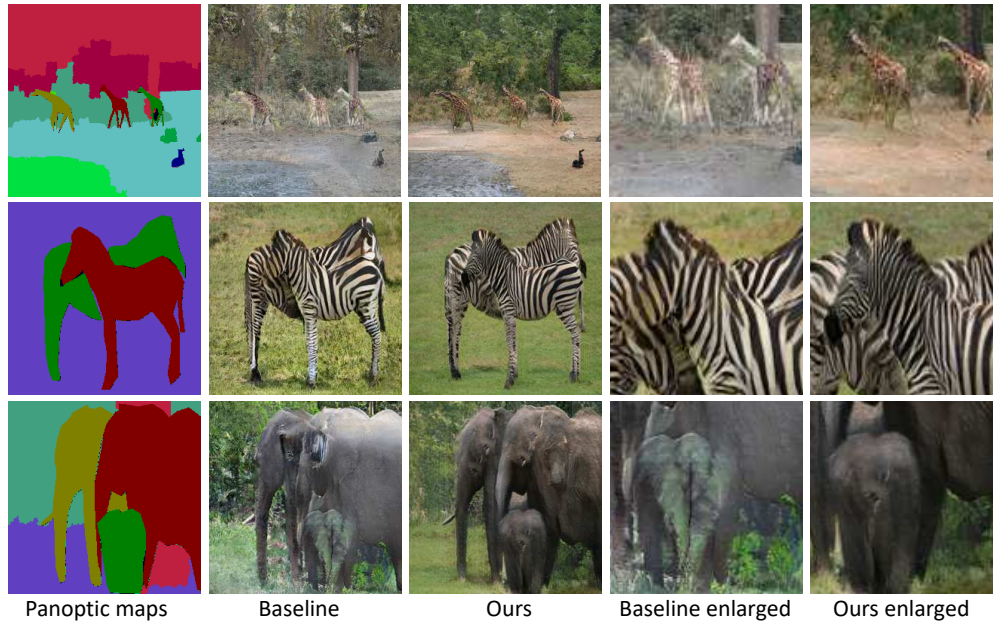


Figure 1. Unlike previous methods that rely on semantic and boundary label maps to synthesize images, our model uses panoptic maps. It generates instances with clear separation even in cluttered scenes where multiple instances occlude each other. The forth and fifth column show zoomed in patches from the second and third column images to highlight the boundaries between instances, where previous methods tend to blend instances together.

## Abstract

Conditional image synthesis for generating photorealistic images serves various applications for content editing to content generation. Previous conditional image synthesis algorithms mostly rely on semantic maps, and often fail in complex environments where multiple instances occlude each other. We propose a panoptic aware image synthesis network to generate high fidelity and photorealistic images conditioned on panoptic maps which unify semantic and instance information. To achieve this, we efficiently use panoptic maps in convolution and upsampling layers. We show that with the proposed changes to the generator, we can improve on the previous state-of-the-art methods by generating images in complex instance interaction environments in higher fidelity and tiny objects in more details. Furthermore, our proposed method also outperforms the

previous state-of-the-art methods in metrics of mean IoU (Intersection over Union), and detAP (Detection Average Precision).

## 1. Introduction

Image synthesis refers to the task of generating diverse and photo-realistic images, where a prevalent sub-category known as conditional image synthesis outputs images that are conditioned on some input data. Recently, deep neural networks have been successful at conditional image synthesis [11, 4, 32, 38, 39, 37, 1] where one of the conditional inputs is a semantic segmentation map. Extending this concept, in this paper, we are interested in the generation of photo-realistic images guided by panoptic maps. Panoptic maps unify semantic and instance maps. Specifically,

they provide information about object instances for countable classes which are called “things” such as people, animals, and cars. Additionally, they contain the semantic information about classes that are amorphous regions and repeat patterns or textures such as grass, sky, and wall. These classes are referred to as “stuff”.

We are interested in panoptic maps because semantic maps do not provide sufficient information to synthesize “things” (instances) especially in complex environments with multiple of them interacting with each other. Even the state-of-the-art baseline (SPADE [23]) which inputs boundary maps to the network fails to generate high fidelity images when objects are small and instances are partially occluding. This issue can be observed in Figure 1, in the continuous pattern extending from one zebra to the other. This is the result of conventional convolution and upsampling algorithms being independent of class and instance boundaries. To address this problem, we replace the convolution and upsampling layers in the generator with *Panoptic aware convolution* and *Panoptic aware upsampling* layers.

We refer to this form of image synthesis as panoptic-based image synthesis. We evaluate our proposed image generator on two diverse and challenging datasets: Cityscapes [5], and COCO-Stuff [2]. We demonstrate that we are able to efficiently and accurately use panoptic maps to generate higher fidelity images and improve on evaluation metrics used by previous methods [4, 32, 23].

Our main contributions can be summarized as follows:

1. We propose to use *Panoptic aware convolution* that re-weights convolution based on the panoptic maps in conditional image generation setting. Similar mechanisms have been previously used for other tasks [15, 7] with binary masks and learned soft masks but not for image synthesis with multi-class panoptic masks.
2. We propose *Panoptic aware upsampling* that addresses the misalignment between the upsampled low resolution features and high resolution panoptic maps. This ensures that the semantic and instance details are not lost, and that we also maintain higher accuracy alignment between the generated images and the panoptic maps.
3. We demonstrate that using our proposed network architecture, not only do we see more photorealistic images, but we also observe significant improvements in object detection scores on both Cityscapes and COCO-Stuff datasets when evaluated with object detection model.

## 2. Related Work

**Generative Adversarial Networks (GANs)** [6] perform image synthesis by modelling the natural image dis-

tribution and synthesizing new samples that are indistinguishable from natural images. This is achieved by using a generator and a discriminator network that are both trying to optimize an opposing objective function, in a zero-sum game. Many conditional image synthesis works use GANs to generate realistic images, and so does ours.

**Conditional Image Synthesis** can vary based on different type of inputs to be conditioned upon. For example, inputs can be text [25, 38, 34, 9], natural and synthetic images [13, 41, 17, 42, 10, 40, 14], or unsupervised landmarks [18, 12, 28] to name a few. Recently, [24, 4, 11] use semantic maps and [32, 23] use both semantic maps and boundary maps as inputs to the generator, where the boundary maps are obtained from the instance maps. A pixel in the boundary map is set to 1 if its object identity is different from any of its 4 neighbors, and set to 0 otherwise. This approach does not preserve the whole information contained in an instance map especially when instances are occluding each other. The pixels that belong to the same instance may be separated by multiple boundaries.

**Content Aware Convolution.** There have been many works that learn to weight the convolution activations based on attention mechanisms [37, 33, 35, 7]. These mechanisms operate on feature maps to capture the spatial locations that are related to each other while making a decision. In another line of research, the spatial locations that should not contribute to an output may be given to us by binary masks such as in the case of image inpainting, the task of filling in holes in an image. In this task, [16, 30] use partial convolutions so that given a binary mask with holes and valid pixels, the convolutional results depend only on the valid pixels. Our convolution layer is similar to the one used in image inpainting, instead of masks with holes, we have panoptic maps given to us, and therefore we know that convolutional results of an instance should not depend on another instance or on pixels that belong to a different semantic class. We are not given binary masks, but we generate them efficiently on-the-fly based on panoptic maps.

**Content Aware Upsampling.** Nearest neighbor and bilinear interpolations are the most commonly used upsampling methods in deep learning applications. These methods use hand-crafted algorithms based on the relative positions of the pixel coordinates. There has been also great interest in learning the upsampling weights for the tasks of semantic segmentation [22, 29] and image and video super-resolution [27]. Recently, [19, 31] proposed feature guided upsampling algorithms. These methods operate on the feature maps to encode contents, and based on the contents they upsample the features. In our method, similar to the idea in the panoptic aware convolution layer, we take advantage of the high resolution panoptic maps to resolve the misalignments in upsampled feature maps and panoptic maps.




Figure 2. Panoptic aware partial convolution layer takes a panoptic map (colorized for visualization) and based on the center of each sliding window it generates a binary mask,  $\mathbf{M}$ . The pixels that share the same identity with the center of the window are assigned 1 and the others 0.

### 3. Method

In this section, we first detail the *Panoptic aware convolution* and *Panoptic aware upsampling* layers. We then describe the overall network architecture.

#### 3.1. Panoptic Aware Convolution Layer

We refer to the partial convolution operation using panoptic maps as a *Panoptic aware partial convolution* layer which shares the fundamentals with other works that use partial convolution for different tasks [7, 15]. Let  $\mathbf{W}$  be the convolution filter weights and  $b$  the corresponding bias.  $\mathbf{X}$  is the feature values,  $\mathbf{P}$  is the panoptic map values for the current convolution (sliding) window, and  $\mathbf{M}$  is the corresponding binary mask.

$\mathbf{M}$  defines which pixels will contribute to the output of the convolution operation based on the panoptic maps. The pixel coordinates which share the same identity with the center pixel in the panoptic map are assigned 1 in the mask, while the others are assigned 0. This is expressed as:

$$m_{(i,j)} = \begin{cases} 1, & \text{if } \mathbf{P}_{(i,j)} == \mathbf{P}_{(\text{center}, \text{center})} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

This can be implemented by first subtracting the center pixel from the patch and clipping the absolute value to (0, 1), then subtracting the clipped output from 1 to inverse the zeros and ones. Figure 2 depicts the construction of the mask  $\mathbf{M}$ .

The partial convolution at every location is expressed as:

$$x' = \begin{cases} \mathbf{W}^T(\mathbf{X} \odot \mathbf{M}) \frac{\text{sum}(\mathbf{1})}{\text{sum}(\mathbf{M})} + b, & \text{if } \text{sum}(\mathbf{M}) > 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where  $\odot$  denotes element-wise multiplication and  $\mathbf{1}$  has same shape as  $\mathbf{M}$  but with all elements being 1. The scaling factor,  $\text{sum}(\mathbf{1})/\text{sum}(\mathbf{M})$ , applies normalization to account

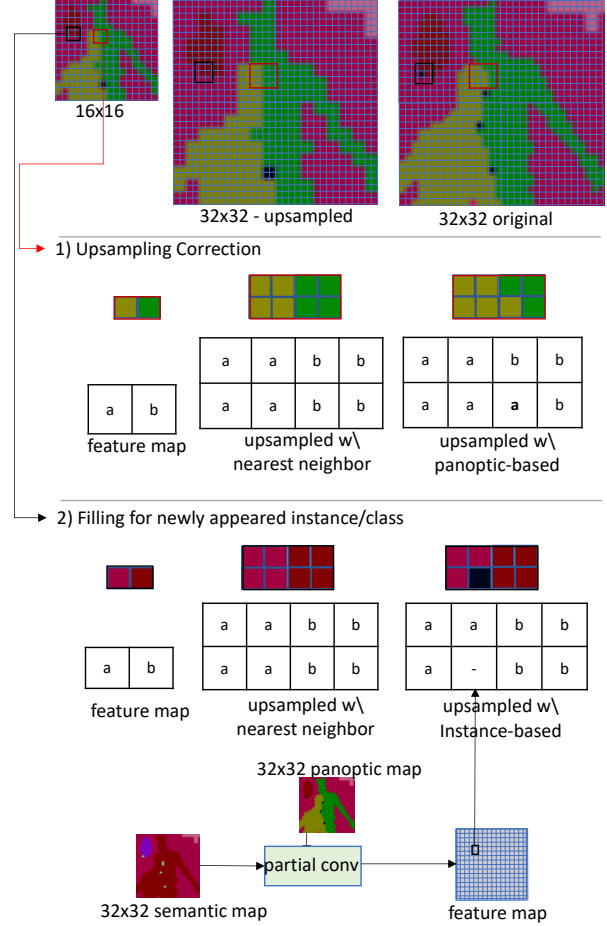


Figure 3. Overview of Panoptic aware upsampling module.  $16 \times 16$  and  $32 \times 32$  panoptic maps are nearest neighbor downsampled from original  $256 \times 256$  panoptic map.  $32 \times 32$  upsampled map is upsampled from  $16 \times 16$  panoptic map using nearest neighbor upsampling algorithm. Comparing the  $32 \times 32$  upsampled and  $32 \times 32$  original map, we can observe two issue: 1) Spatial misalignments and 2) Appearance of new classes or instances. As shown in Figure (top), first we correct misalignment by replicating a feature vector from a neighboring pixel that belongs to the same panoptic instance. This operation is different from nearest neighbor upsampling which would always replicate the top-left feature. Second, as shown in Figure (bottom), we resolve pixels where new semantic or instance classes have just appeared by encoding new features from semantic maps with *Panoptic aware convolution* layer.

for varying amount of valid inputs as in [15]. With Equation 2, convolution results of an instance or stuff depend only on the feature values that belong to the same instance or stuff.

#### 3.2. Panoptic Aware Upsampling Layer

We propose a *Panoptic aware upsampling* layer as an alternative to traditional upsampling layers when the higher resolution panoptic maps are available as in the case of im-




Figure 4. The percent of incorrectly mapped features using upsampling through different layers of network.

age synthesis for content generation task. Nearest neighbor upsampling is a popular conventional upsampling choice in conditional image synthesis tasks as used by [1, 23, 32, 23]. However, nearest neighbor upsampling algorithm is hand-crafted to do replication. For example, in a  $2 \times 2$  upsampling scenario, nearest neighbor algorithm will replicate the top-left corner to the neighboring pixels in a  $2 \times 2$  window. This creates two issues as shown in Figure 3.

First, it can create a spatial misalignment between the high resolution panoptic map and upsampled features. Figure 3-top illustrates this issue, where the features of instance  $id_B$  are replicated and incorrectly used for instance  $id_A$  following traditional upsampling approach. In Figure 3, we demonstrate the misalignments in the upsampled panoptic maps for clarity, but we are only interested in the alignment in feature maps. We refer to the operation for fixing this misalignment as “Upsampling alignment” correction. Secondly, as shown in Figure 3-bottom, the high resolution panoptic map may contain new classes and instances that might not exist in the lower resolution panoptic map. This implies that new features need to be generated and replaced in upsampled feature map. We refer to this operation as “Hole filling”.

Figure 4 depicts how often the two issues mentioned above occur at different layers in the network for the Cityscapes dataset. As seen in the figure, especially in the early layers, over 30% of pixel features among the newly generated ones do not align with the panoptic maps, and many pixels that belong to a new instance or semantic map appear for the first time at the new scale.

To resolve these two issues, *Panoptic aware upsampling* layer performs a two-step process: Upsample alignment correction and Hole filling as shown in Figure 3. Let  $S$  be the semantic map, and  $F$  the feature to be upsampled. We are interested in  $2 \times 2$  upsampling as it is the most common upsampling scale used by image synthesis methods. Let  $P^d$  be the downsampled panoptic mask. We are interested in upsampling  $F^d$ , to generate the upsampled feature

#### Algorithm 1 Upsampling Alignment Correction.

---

**Initialize:**  $M^{correction} = 0, F'^u = 0,$

**for**  $i \in [0, 2W); j \in [0, 2H)$  **do**

**if**  $P_{i,j}^u == P_{i//2,j//2}^d$  **then**

$F'_{i,j}{}^u = F_{i//2,j//2}^d$

$M_{i,j}^{correction} = 1$

**end if**

**end for**

**for**  $i \in [0, 2W); j \in [0, 2H)$  **do**

**if**  $P_{i,j}^u == P_{i//2+1,j//2}^d$  **and**  $M_{i,j}^{correction} \neq 1$  **then**

$F'_{i,j}{}^u = F_{i//2+1,j//2}^d$

$M_{i,j}^{correction} = 1$

**end if**

**end for**

**for**  $i \in [0, 2W); j \in [0, 2H)$  **do**

**if**  $P_{i,j}^u == P_{i//2,j//2+1}^d$  **and**  $M_{i,j}^{correction} \neq 1$  **then**

$F'_{i,j}{}^u = F_{i//2,j//2+1}^d$

$M_{i,j}^{correction} = 1$

**end if**

**end for**

**for**  $i \in [0, 2W); j \in [0, 2H)$  **do**

**if**  $P_{i,j}^u == P_{i//2+1,j//2+1}^d$  **and**  $M_{i,j}^{correction} \neq 1$  **then**

$F'_{i,j}{}^u = F_{i//2+1,j//2+1}^d$

$M_{i,j}^{correction} = 1$

**end if**

**end for**

---

map,  $F'^u$  with the guidance from a higher scale panoptic and semantic maps,  $P^u$ , and  $S^u$ , and a mask,  $M^{correction}$ , that we will generate.

To correct the misalignment in the  $2 \times 2$  upsampling layer, we scan the four neighbors of each pixel. In the  $2 \times 2$  window, if we find a match between the corresponding pixels’ panoptic identity in higher resolution and a neighboring pixels panoptic identity in lower resolution, we copy over that neighboring feature to the corresponding indices in upsampled feature map. This method is depicted in Algorithm 1. Note that the first for loop would correspond to the nearest neighbor upsampling algorithm if there was no if statement in the loop. We also update the mask,  $M^{correction}$ , to keep track of which indices have been successfully aligned. In the subsequent for loops, for the indices that were not aligned yet, we check if any of the other neighbors are matching the panoptic identity with them.

After Algorithm 1, we end up with a partially filled upsampled feature map  $F'^u$  and a  $M^{correction}$  mask which defines which coordinates found a match. After that we calculate the final  $F'^u$  by:

$$F'^u_{(i,j)} = F'^u_{(i,j)} + \underbrace{(1 - M_{i,j}^{correction}) * f_{hole\ filling}(S^u_{(i,j)})}_{\text{Hole Filling}}$$



We generate  $f_{holefilling}$  with the *Panoptic aware convolution* layer by feeding semantic map ( $S^u$ ) as input and panoptic map ( $P^u$ ) for guidance. We use the semantic map to encode features from  $K \times 2W \times 2H$  semantic map where  $K$  is the number of classes to a higher dimension  $C \times 2W \times 2H$  with this layer.

$$f_{holefilling} = PanopticAwareConvolution(S^u) \quad (3)$$

With *Panoptic aware upsampling* layer, the features that have been tailored for a specific instance or a semantic map are not copied over to another one, which improves the accuracy of the generated images.

### 3.3. Network Architecture

Our final proposed architecture, motivated by SPADE [23] is described in Figure 5. Similar to SPADE, we feed a downsampled segmentation map to the first layer of the generator, but in our architecture, it is a Panoptic aware convolution layer that encodes features from  $\#Classes \times W \times H$  semantic map to a higher dimension  $1024 \times W \times H$ . In the rest of the network, we replace all convolution layers in the ResNet Blocks with the Panoptic aware convolution layers and all upsampling layers with Panoptic aware upsampling layers. Each block operates at a different scale and we downsample the semantic and panoptic maps to match the scale of the features.

The input to the SPADE module is kept as a semantic map which learns denormalization parameters. Panoptic maps are not suitable for this computation, since the convolution operation expects a fixed number of channels. Hence, we rely on SPADE to provide the network with the correct statistics of features based on the semantic classes.

We feed panoptic maps to the panoptic aware convolution layers in order to perform convolution operation based on the instances and classes. The original full resolution panoptic and semantic maps are also fed to the panoptic aware upsampling layers to perform upsampling alignment correction and hole filling.

The panoptic aware convolution layer that is in the first layer of the architecture which encodes features from  $\#Classes \times W \times H$  semantic map to a higher dimension encoded features are shared between panoptic aware upsampling layers in the rest of the network. When the number of channels the partial convolution layer from the first layer generates does not match the one expected at different blocks, we decrease the dimension with  $1 \times 1$  convolution layers. This layer is depicted by green boxes in Figure 5. Note that the green box is depicted multiple times in Figure but they are shared between stages. By sharing the weights, we do not introduce additional parameters to the baseline except the negligible cost of  $1 \times 1$  convolutions. Sharing these weights also makes sense since the task of this layer

is common at each stage which is to generate features for instances and semantic classes that appear for the first time at that stage.

## 4. Experiments

**Datasets.** We conduct our experiments on Cityscapes [5] and COCO-Stuff [2] datasets that have both instance and semantic segmentation labels available. The Cityscapes dataset contains 3,000 training images and 500 validation images of urban street scenes along with 35 semantic classes and 9 instance classes. All classes are used while synthesizing images but only 19 classes are used for semantic evaluation as defined by Cityscapes evaluation benchmark. COCO-Stuff dataset has 118,000 training images and 5,000 validation images from both indoor and outdoor scenes. This dataset has 182 semantic classes and 81 instance classes.

**Implementation Details.** We use the parameters provided by SPADE baseline [23]. Specifically, we use synchronized batch normalization to collect statistics across GPUs, and apply Spectral Norm [20] to all layers in the generator and discriminator. We train and generate images in  $256 \times 256$  resolution for COCO-Stuff, and  $256 \times 512$  for Cityscapes dataset. We train 200 epochs on Cityscapes dataset with batch size 16, and linearly decay the learning rate after 100 epochs as done by [23]. COCO-Stuff dataset is trained for 100 epochs with batch size of 48 with constant learning rate. Initial learning rates are set to 0.0001 and 0.0004 for the generator and discriminator respectively, and networks are trained with ADAM solver with  $\beta_1 = 0$  and  $\beta_2 = 0.999$ .

**Performance Metrics.** We adopt the evaluation metrics as previous conditional image synthesis work [23, 32] plus add another metric for detecting successfully generated object instances. The first two metrics, mean intersection over union (mIoU) and overall pixel accuracy (accuracy), are obtained by inferring a state-of-the-art semantic segmentation model on the synthesized images and comparing how well the predicted segmentation mask matches the ground truth semantic map. Additionally, we use Detection Average Precision (detAP) by using a trained object detection network to evaluate instance detection accuracy on synthesized images.

We use the same segmentation networks used in [23] for evaluation. Specifically, we use DeepLabV2 [3, 21] for COCO-Stuff, and DRN-D-105 [36] for Cityscapes datasets. For detection, we use Faster-RCNN [26] with ResNet-50 backbone. In addition to the mIoU, accuracy and detAP performance metrics, we use the Fréchet Inception Distance (FID) [8] to measure the distance between the distribution of synthesized results and the distribution of real images.

**Baselines.** We compare our method against three popular image synthesis frameworks, namely: cascaded refinement network (CRN) [4], semi-parametric image synthesis

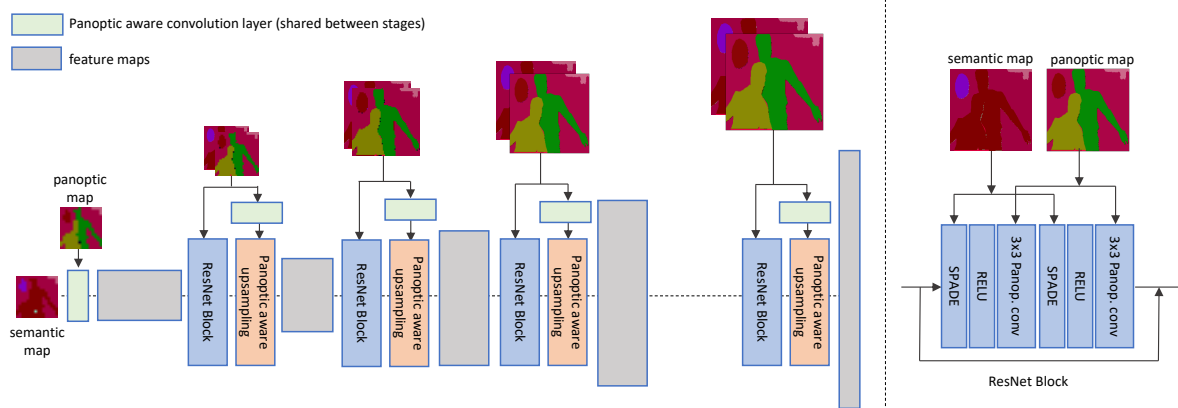


Figure 5. In our generator, each ResNet Block layer uses segmentation and panoptic masks to modulate the layer activations. (Left) The generator contains a series of the residual blocks with Panoptic aware convolution and upsampling layers. (Right) Structure of residual blocks.

Method	detAP	mIoU	accuracy	FID
CRN [4]	8.75	52.4	77.1	104.7
SIMS [24]	2.60	47.2	75.5	<b>49.7</b>
SPADE [23]	11.67	62.3	81.9	71.8
SPADE*	11.80	62.2	81.9	94.0
Ours	<b>13.43</b>	<b>64.8</b>	<b>82.4</b>	96.4

Table 1. Results on Cityscapes. Our method outperforms current leading methods in detAP, mIoU and overall pixel accuracy. SPADE\* is trained by us.

(SIMS) [24], and spatially-adaptive denormalization model (SPADE) [23]. CRN uses a deep network with given semantic label map, it repeatedly refines the output from low to high resolution without an adversarial training. SIMS uses a memory bank of image segments constructed from a training set of images and refines the boundaries via a deep network. Both SIMS and CRN operate only on a semantic map. SPADE is the current state-of-the-art conditional image synthesis method, and not only uses semantic map but also incorporates the instance information via a boundary map. The pixel in the boundary map is 1 if its object identity is different from any of its 4 neighbors, and 0 otherwise. This approach does not provide the full instance information especially in cluttered scenes with lots of objects occluding each other. We compare with SIMS on Cityscapes dataset but not on COCO-stuff dataset as SIMS requires queries to the training set images and it is computationally costly for a large dataset such as the COCO-stuff dataset.

**Quantitative Results.** In Tables 1 and 2, we provide the results for Cityscapes and COCO-Stuff datasets, respectively. We find that our method outperforms the current state-of-the-art by a large margin for object detection score, mIoU and pixel level accuracy in both datasets. Table 4 reports mIoU for each class in Cityscapes dataset. We improve al-

Method	detAP	mIoU	accuracy	FID
CRN [4]	22.7	23.7	40.4	70.4
SPADE [23]	28.5	37.4	67.9	<b>22.6</b>
SPADE*	29.0	38.2	68.6	25.3
Ours	<b>31.0</b>	<b>38.6</b>	<b>69.0</b>	28.8

Table 2. Results on COCO-Stuff. Our method outperforms current leading methods in detAP, mIoU and overall pixel accuracy. SPADE\* is trained by us.

most all of the classes significantly. Especially, our proposed method improves mIoU for traffic sign from 44.7 to 50.0, which is a challenging class because of the small size of the signs.

We observe a slight degradation in our FID score compared to the released SPADE models, and the SPADE models we trained with the parameters provided by [23]. FID score tries to match the variances/diversities between real and generated images, without caring about the correspondence with the conditioned semantic map and instance map. Our results have better correspondence with the underlying semantic and instance maps. Though this would be the desired behavior, the results may get affected by human annotation bias. We suspect, such annotation bias (e.g. straight line bias, over-simplified polygonal shape bias) in the inputs may deteriorate the matching of variances. Also note that SIMS produces images that have significantly lower FID score than the other methods even though it achieves worse detAP and mIoU scores. This is because SIMS copies image patches from training dataset, and sometimes the copied patch does not faithfully match the given segmentation mask. This issue becomes even more apparent in detAP score, as SIMS copies over patches without ensuring the number of cars being consistent with the panoptic map.

**Qualitative Comparison.** In Figures 6 and 7, we provide image synthesis results of our method and other competing

Method	road	swalk	build.	wall	fence	pole	tlght	tsign	veg.	terr.	sky	person	rider	car	truck	bus	train	mbike	bike
CRN [4]	96.9	79.5	76.7	29.0	10.6	34.8	39.8	44.3	68.4	54.4	91.9	63.0	39.7	87.8	25.0	56.2	31.8	14.5	52.2
SIMS [24]	93.3	66.1	73.6	33.1	34.5	30.3	27.2	39.5	73.4	46.2	56.6	42.9	31.0	70.3	35.8	42.5	37.3	20.3	43.1
SPADE [23]	97.4	80.0	87.9	50.6	47.2	<b>35.9</b>	39.0	44.7	88.2	66.1	91.6	62.3	<b>38.7</b>	88.7	65.0	70.2	41.4	28.6	58.8
Ours	<b>97.7</b>	<b>82.5</b>	<b>89.2</b>	<b>60.6</b>	<b>54.2</b>	35.3	<b>39.8</b>	<b>50.0</b>	<b>89.5</b>	<b>69.0</b>	<b>92.4</b>	<b>63.2</b>	38.2	<b>90.6</b>	<b>66.7</b>	<b>72.2</b>	<b>48.8</b>	<b>31.2</b>	<b>59.1</b>

Table 3. Per-class mIoU results on Cityscapes.

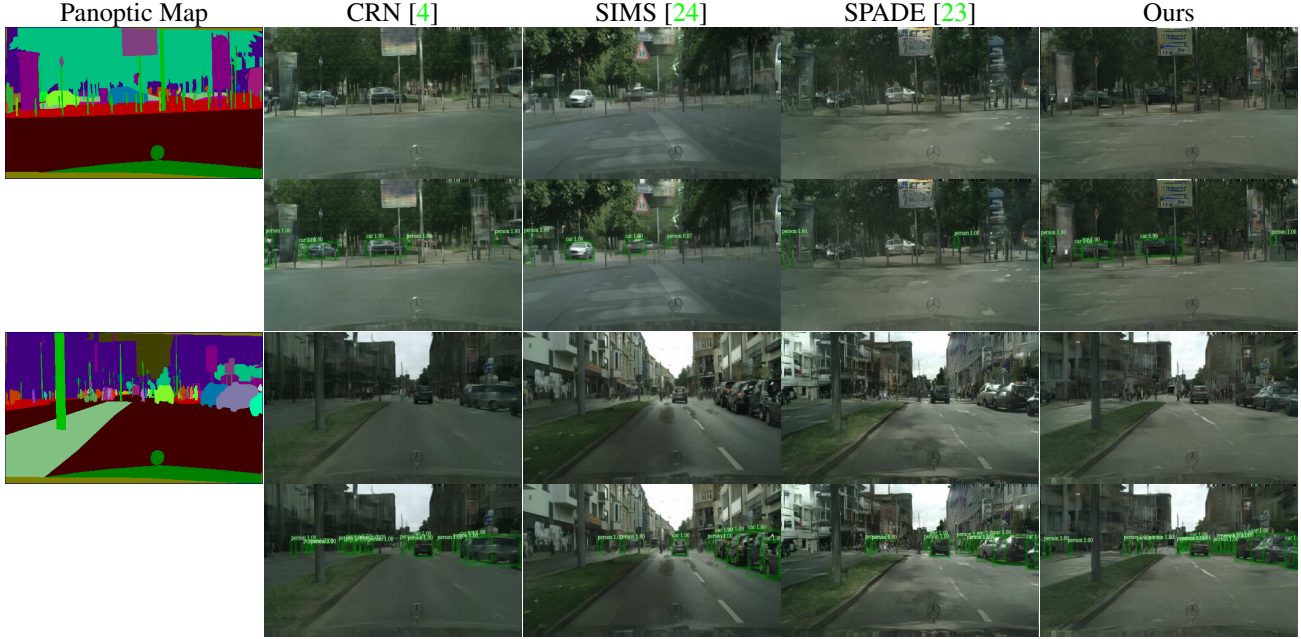


Figure 6. Visual comparison of image synthesis results on Cityscapes dataset. We also provide the bounding box detection predictions from Faster-RCNN. The cars in the first row images are occluded by poles which create a challenge for the image synthesis methods. CRN generates cars that can be detected by Faster-RCNN but visually look less pleasing. SIMS loosely follows the provided semantic map, and the cars SPADE generates are not distinctive enough to be detected by Faster-RCNN. In the third row, cars generated on the right side of the images present a challenge for the algorithms that only use semantic maps as seen in the images from CRN and SIMS where CRN generates two cars and SIMS generates four cars while three cars should be present. Thanks to the boundary maps used in SPADE, it can generate the correct number of cars. However, our proposed method along with generating the correct number of car instances, also generates more instances of persons that can be detected with higher accuracy.

methods. We also provide the bounding box detection predictions from Faster-RCNN. We especially provide examples where multiple instances occlude each other. We find that our method produces instances with better visual quality in challenging scenarios. Specifically, we find that our method generates distinct cars even when they are behind poles, and can generate detectable people even when they are far away as shown in Figure 6. As can be seen in Figure 7, we find that other methods may blend the pattern and texture of objects among neighboring instances, whereas our method clearly separates them.

**Ablation Studies.** We conduct controlled experiments and gradually add our proposed components. We start with a baseline SPADE model [23]. We train the model three times, and report the average results. First, we replace the convolutions in ResNet blocks and the first layer with panoptic aware convolution layers. Second, we additionally replace nearest neighbor upsampling layers with panoptic

Method	mIoU	detAP
Baseline (SPADE)	60.00	10.97
+Panoptic-aware Partial Conv	61.24	11.50
+Panoptic-aware Upsampling	64.55	13.04

Table 4. Ablation studies on Cityscapes dataset. Results are averaged over 3 runs, and they are slightly different than the results in Table 1.

aware upsampling layers. The segmentation mIoU scores and detAP scores of the generated images by each setup are shown in Table 4 where each added module increases the performance.

## 5. Conclusion

In conclusion, we propose a panoptic-based image synthesis network that generates images with higher fidelity to the underlying segmentation and instance information.






Figure 7. Visual comparison of image synthesis results on the COCO-stuff dataset. We also display the bounding box detection predictions from Faster-RCNN. Other methods generate patterns that are continues throughout instances which makes the instances indistinguishable. Also note that in the last row our method is able to produce detectable car instances in a cluttered scene.

We show that our method is better at generating distinct instances in challenging scenarios and outperforms previous state-of-the-art significantly in detAP metric, a metric which has not been used to evaluate conditional image synthesis results before.

**Future Work.** Multi-modal image synthesis and controllability of styles are very important for content generation applications. The architecture in our experiments does not support style-guided image synthesis. However, our work can be extended to output multiple styles via an encoder-

decoder architecture as proposed in pix2pixHD [32]. Furthermore, the proposed panoptic aware convolution and up-sampling layers can be used for feature maps that decode styles, and can provide further improvements. We leave this as future work.

## References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthe-



- sis. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 4
- [2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocosuff: Thing and stuff classes in context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1209–1218, 2018. 2, 5
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 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. 5
- [4] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1511–1520, 2017. 1, 2, 5, 6, 7, 8
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 2, 5
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2
- [7] Adam W Harley, Konstantinos G Derpanis, and Iasonas Kokkinos. Segmentation-aware convolutional networks using local attention masks. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, page 7, 2017. 2, 3
- [8] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. 5
- [9] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [10] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. *European Conference on Computer Vision (ECCV)*, 2018. 2
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [12] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *Advances in Neural Information Processing Systems*, 2018. 2
- [13] Levent Karacan, Zeynep Akata, Aykut Erdem, and Erkut Erdem. Learning to generate images of outdoor scenes from attributes and semantic layouts. *arXiv preprint arXiv:1612.00215*, 2016. 2
- [14] Levent Karacan, Zeynep Akata, Aykut Erdem, and Erkut Erdem. Manipulating attributes of natural scenes via hallucination. *arXiv preprint arXiv:1808.07413*, 2018. 2
- [15] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018. 2, 3
- [16] Guilin Liu, Kevin J. Shih, Ting-Chun Wang, Fitsum A. Reda, Karan Sapra, Zhiding Yu, Andrew Tao, and Bryan Catanzaro. Partial convolution based padding, 2018. 2
- [17] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, 2017. 2
- [18] Dominik Lorenz, Leonard Bereska, Timo Milbich, and Björn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *CVPR*, 2019. 2
- [19] Davide Mazzini. Guided upsampling network for real-time semantic segmentation. *arXiv preprint arXiv:1807.07466*, 2018. 2
- [20] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018. 5
- [21] Kazuto Nakashima. Deeplab-pytorch. <https://github.com/kazuto1011/deeplab-pytorch>, 2018. 5
- [22] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. *IEEE International Conference on Computer Vision (ICCV)*, Dec 2015. 2
- [23] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. 2, 4, 5, 6, 7, 8
- [24] Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun. Semi-parametric image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8808–8816, 2018. 2, 6, 7
- [25] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning (ICML)*, 2016. 2
- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 5
- [27] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [28] Kevin J Shih, Aysegul Dundar, Animesh Garg, Robert Pottorf, Andrew Tao, and Bryan Catanzaro. Video interpolation and prediction with unsupervised landmarks. *arXiv preprint arXiv:1909.02749*, 2019. 2
- [29] Hang Su, Varun Jampani, Deqing Sun, Orazio Gallo, Erik Learned-Miller, and Jan Kautz. Pixel-adaptive convolutional

- neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11166–11175, 2019. 2
- [30] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. *arXiv preprint arXiv:1708.06500*, 2017. 2
- [31] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. Carafe: Content-aware reassembly of features, 2019. 2
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 1, 2, 4, 5, 8
- [33] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. 2
- [34] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [35] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016. 2
- [36] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [37] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018. 1, 2
- [38] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaolei Huang, Xiaogang Wang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2
- [39] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018. 1
- [40] Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [42] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, 2017. 2