

Full-Glow: Fully conditional Glow for more realistic image generation

Moein Sorkhei¹[0000–0001–6204–0778], Gustav Eje Henter¹[0000–0002–1643–1054], and Hedvig Kjellström^{1,2}[0000–0002–5750–9655]

¹ KTH Royal Institute of Technology, Sweden

² Silo AI, Sweden

{sorkhei, ghe, hedvig}@kth.se

Abstract. Autonomous agents, such as driverless cars, require large amounts of labeled visual data for their training. A viable approach for acquiring such data is training a generative model with collected real data, and then augmenting the collected real dataset with synthetic images from the model, generated with control of the scene layout and ground truth labeling. In this paper we propose *Full-Glow*, a fully conditional Glow-based architecture for generating plausible and realistic images of novel street scenes given a semantic segmentation map indicating the scene layout. Benchmark comparisons show our model to outperform recent works in terms of the semantic segmentation performance of a pretrained PSPNet. This indicates that images from our model are, to a higher degree than from other models, similar to real images of the same kinds of scenes and objects, making them suitable as training data for a visual semantic segmentation or object recognition system.

Keywords: Conditional image generation · Generative models · Normalizing flows.

1 Introduction

Autonomous mobile agents, such as driverless cars, will be a cornerstone of the smart society of the future. Currently available datasets of labeled street scene images, such as Cityscapes [6], are an important step in this direction, and could, e.g., be used for training models for semantic image segmentation. However, collecting such data poses challenges including privacy intrusions, the need for accurate crowd-sourced labels, and the requirement to cover a huge state-space of different situations and environments. Another approach – especially useful to gather data representing dangerous situations such as collisions with pedestrians – is to generate training images with known ground-truth labeling using game engines or other virtual worlds, but this approach requires object and state-space variability to be manually engineered into the system.

A viable alternative to both these approaches is to augment existing datasets with synthetically-generated novel datapoints, produced by generative image models trained on the existing data. This builds on recent applications of generative models for a variety of tasks such as image style transfer [15] and modality transfer in medical imaging [36].

Among currently-available deep generative approaches, GANs [10] are probably the most widely used in image generation, owing to their achievements in synthesizing realistic high-resolution output with novel and rich detail [3, 16]. Auto-regressive architectures [26, 27] are usually computationally demanding (not parallelizable) and not feasible for generating higher-resolution images. Image samples generated by early variants of VAEs [20, 33] tended to suffer from blurriness [44], although the realism of VAE output has improved in recent years [32, 40].

This article considers normalizing flows [7, 8], a different model class of growing interest. With recent improvements such as *Glow* [19], flows can generate images with a quality that approaches that produced by GANs. Flows have also achieved competitive results in other tasks such as audio and video generation [17, 21, 30]. Flow-based models exhibit several benefits compared to GANs: 1) stable, monotonic training, 2) learning an explicit representation useful for down-stream tasks such as style transfer, 3) efficient synthesis, and 4) exact likelihood evaluation that could be used for density estimation.

In this paper, we propose a new, fully conditional Glow-based architecture called *Full-Glow* for generating plausible street scene images conditioned on the structure of the image content (i.e., the segmentation mask). We show that, by using this model, we are able to synthesize moderately high-resolution images that are consistent with the given structure but differ substantially from the existing ground-truth images. A quantitative comparison against previously proposed Glow-based models [23, 36] and the popular GAN-based conditional image-generation model pix2pix [15], finds that our improved conditioning allows us to synthesize images that achieve better semantic classification scores under a pre-trained semantic classifier. We also provide visual comparisons of samples generated by different models.

The remainder of this article is laid out as follows: Section 2 presents prior work in street-scene generation and image-to-image translation, while Section 3 provides technical background on normalizing flows. Our proposed fully-conditional architecture is then introduced in Section 4 and validated experimentally in Section 5.

2 Related work

Synthetic data generation. Street-scene image datasets such as Cityscapes [6], CamVid [4], and the KITTI dataset [9] are useful for training vision systems for street-scene understanding. However, collecting and labeling such data is costly, resource demanding, and associated with privacy issues. An effective alternative that allows for ground-truth labels and scene layout control is synthetic data generation using game engines [34, 35, 39]. Despite these advantages, images generated by game engines tend to differ significantly from real-world images and may not always act as a replacement for real data. Moreover, game engines generally only synthesize objects from pre-generated assets or recipes, meaning that variation has to be hand-engineered in. It is therefore difficult and costly to obtain diverse data in this manner. Data generated from approaches such as ours address these shortcomings while maintaining the benefits of ground-truth labeling and scene layout control.

Image-to-image translation. In order to generate images for data-augmentation of supervised learning tasks, it is necessary to condition the image generation on an input,

such that the ground-truth labeling of the generated image is known. For street-scene understanding, this conditioning takes the form of per-pixel class labels (a *segmentation mask*), meaning that the augmentation task can be formulated as an image-to-image translation problem. GANs [10] have been employed for both paired and unpaired image-to-image translation problems [15, 45]. While GANs can generate convincing-looking images, they are known to suffer from mode collapse and low output diversity [11]. Consequentially, their value in augmenting dataset diversity may be limited.

Likelihood-based models, on the other hand, explicitly aim to learn the probability distribution of the data. These models generally prefer sample diversity, sometimes at the expense of sample quality [8], which has been linked to the mass-covering property of the likelihood objective [25, 37]. Like for GANs [3], perceived image quality can often be improved by reducing the entropy of the distribution at synthesis time, relative to the distribution learned during training, cf. [19, 40]. Flow-based models are a particular class of likelihood-based model that have gained recent attention after an architecture called Glow [19] demonstrated impressive performance in unconditional image generation. Previous works have applied flow-based models for image colorization [1, 2], image segmentation [23], modality transfer in medical imaging [36], and generating point clouds and image-to-image translation [31].

So far, Glow-based models proposed for image-to-image translation [23, 31, 36] have only considered low-resolution tasks. Although the results are promising, they do not assess the full capacity of Glow for generating realistic image detail, for example in street scenes. High-resolution street-scene synthesis has been performed by the GAN-based model pix2pixHD [41] on a GPU with very high memory capacity (24 GB). In the present work, we synthesize moderately high resolution street scene images using a GPU with lower memory capacity (11–12 GB). We extend previous works on Glow-based models by introducing a fully conditional architecture, and also by modeling high-resolution street-scene images, which is a more challenging task than the low-resolution output considered in prior work.

3 Flow-based generative models

Normalizing flows [28] are a class of probabilistic generative models, able to represent complex probability densities in a manner that allows both easy sampling and efficient training based on explicit likelihood maximization. The key idea is to use a sequence of *invertible* and *differentiable* functions/transformations which (nonlinearly) transform a random variable \mathbf{z} with a simple density function to another random variable \mathbf{x} with a more complex density function (and vice versa, thanks to invertibility):

$$\mathbf{x} \xleftrightarrow{\mathbf{f}_1} \mathbf{h}_1 \xleftrightarrow{\mathbf{f}_2} \mathbf{h}_2 \xleftrightarrow{\mathbf{f}_3} \dots \xleftrightarrow{\mathbf{f}_K} \mathbf{z}. \quad (1)$$

Each component transformation \mathbf{f}_i is called a *flow step*. The distribution of \mathbf{z} (termed the *latent*, *source*, or *base distribution*) is assumed to have a simple parametric form, such as an isotropic unit Gaussian. Similar to in GANs, the generative process can be formulated as:

$$\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z}), \quad (2)$$

$$\mathbf{x} = \mathbf{g}_{\theta}(\mathbf{z}) = \mathbf{f}_{\theta}^{-1}(\mathbf{z}) \quad (3)$$

where \mathbf{z} is sampled from the base distribution and \mathbf{g}_θ represents the cumulative effect of the parametric invertible transformations in Equation (1). The log-density function of \mathbf{x} under this transformation can be written as:

$$\log p_{\mathbf{x}}(\mathbf{x}) = \log p_{\mathbf{z}}(\mathbf{z}) + \sum_{i=1}^K \log \left| \det \frac{d\mathbf{h}_i}{d\mathbf{h}_{i-1}} \right| \quad (4)$$

using the *change-of-variables theorem*, where $\mathbf{h}_0 \triangleq \mathbf{x}$ and $\mathbf{h}_K \triangleq \mathbf{z}$. Equation (4) can be used to compute the exact dataset log-likelihood (not possible in GANs) and is the sole objective function for training flow-based models.

The central design challenge of normalizing flows is to create expressive invertible transformations (typically parameterized by deep neural networks) where the so-called *Jacobian log-determinant* in Equation (4) remains computationally feasible to evaluate. Often, this is achieved by designing transformations whose Jacobian matrix is triangular, making the determinant trivial to compute. An important example is NICE [7]. NICE introduced the *coupling layer*, which is a particular kind of flow nonlinearity that uses a neural network to invertibly transform half of the elements in \mathbf{h}_k with respect to the other half. RealNVP [8] improved on this architecture using more general invertible transformations in the coupling layer and by imposing a hierarchical structure where the flow is partitioned into *blocks* that operate at different resolutions. This hierarchy allows using smaller \mathbf{z} -vectors at the initial, smaller resolutions, speeding up computation, and has lately been used by other prominent image-generation systems [32, 40]. Glow [19] added *actnorm* as a replacement for batchnorm [14] and introduced invertible 1×1 convolutions to more efficiently mix variables in between the couplings.

A number of Glow-based architectures have been proposed for conditional image generation. In these models, the goal is to learn a distribution over the target image \mathbf{x}_b conditioned on the source image \mathbf{x}_a . C-Glow [23] is based on the standard Glow architecture from [19], but makes all sub-steps inside the Glow conditional on the raw conditioning image \mathbf{x}_a . The Dual-Glow [36] architecture instead builds a generative model of both source and target image together. It consists of two Glows where the base variables \mathbf{z}_a of the source-image Glow determine the Gaussian distribution of the corresponding base variables \mathbf{z}_b of the target-image Glow through a neural network. Because of the hierarchical structure of Glow, several different conditioning networks are used, one for each block of flow steps. C-Flow [31] described a similar structure of side-by-side Glows, but kept the Gaussian base distributions in the two flows independent. Instead, they used the latent variables $\mathbf{h}_{a,i}$ at every flow step i of the target-domain Glow to condition the transformation in the coupling layer at the corresponding level in the source-domain Glow. Compared to the raw image-data conditioning in C-Glow, Dual-Glow and C-Flow simplify the conditional mapping task at the different levels since the source and target information sit at comparable levels of abstraction.

4 Fully conditional Glow for scene generation

This section introduces our new, fully conditional Glow architecture for image-to-image translation, which combines key innovations from all three previous architectures, C-Glow, Dual-Glow, and C-Flow: Like Dual-Glow and C-Flow (but unlike C-Glow) we

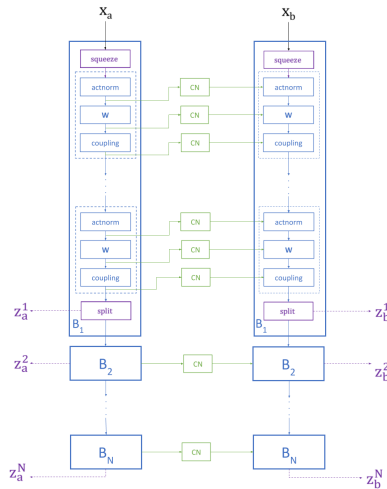


Fig. 1: The proposed architecture, where all substeps have been made conditional by inserting conditioning networks. \mathbf{x}_a and \mathbf{x}_b are paired images in the source and target domains, respectively.

use two parallel stacks of Glow, so that we can leverage conditioning information at the relevant level of hierarchy and not be restricted to always use the raw source image as input. In contrast to Dual-Glow and C-Flow (but reminiscent of C-Glow), we introduce conditioning networks that allow *all* operations in the target-domain Glow conditional on the source-domain information. The resulting architecture is illustrated in Figure 1. Because of its fully conditional nature, we dub this architecture *Full-Glow*.

In our proposed architecture, not only is the coupling layer conditioned on the output of the corresponding operation in the source Glow, but the actnorm and the 1×1 convolutions in the target Glow are also connected to the source Glow. In particular, the parameters of these two operations in each target-side step are generated by conditioning networks CN built from convolutional layers followed by fully connected layers. These networks also enable us to exploit other side information for conditioning, for instance by concatenating the side information with the other input features of each conditioning network. We experimentally show that making the model fully conditional indeed allows for learning a better conditional distribution (measured with lower conditional bits-per-dimension) and more semantically meaningful images (measured using a pre-trained semantic classifier).

We will now describe the architecture of the fully-conditional target-domain Glow in more detail. We describe the computational statistical inference (\mathbf{f}_θ), but every transformation is invertible for synthesis (\mathbf{g}_θ) given the conditioning image \mathbf{x}_a .

Conditional actnorm. The shift \mathbf{t} and scale \mathbf{s} parameters of the conditional actnorm are computed as follows:

$$\mathbf{s}, \mathbf{t} = \text{CN}(\mathbf{x}_{\text{act}}^{\text{source}}) \tag{5}$$

where $\mathbf{x}_{\text{act}}^{\text{source}}$ is the output of the corresponding actnorm in the source Glow. For initializing the actnorm conditioning network (CN), we set all parameters of the network except those of the output layer to small, random values. Similar to [19,42], the weights of the output layer are initialized to 0 and the biases are initialized such that the output activations of the target side after applying actnorm have mean 0 and std of 1 per channel for the first batch of data, similar to the scheme used for initialising actnorm in regular Glow.

Conditional 1×1 convolution. Like Glow, we represent the convolution kernel W using an LU decomposition for easy log-determinant computation, but we have conditioning networks generate the \mathbf{L} , \mathbf{U} matrices and the \mathbf{s} vector:

$$\mathbf{L}, \mathbf{U}, \mathbf{s} = \text{CN}(\mathbf{x}_{\mathbf{W}}^{\text{source}}) \quad (6)$$

where $\mathbf{x}_{\mathbf{W}}^{\text{source}}$ is the output of the corresponding 1×1 convolution in the source Glow, \mathbf{L} is a lower triangular matrix with ones on the diagonal, \mathbf{U} is an upper triangular matrix with zeros on the diagonal, and \mathbf{s} is a vector. Initialization again follows [19]: we first sample a random rotation matrix \mathbf{W}_0 per layer, which we factorise using the LU decomposition as $\mathbf{W}_0 = \mathbf{P}\mathbf{L}_0(\mathbf{U}_0 + \text{diag}(\mathbf{s}_0))$. The conditioning network is then set up similar as for the actnorm, with weights and biases set so that its outputs on the first batch are constant and equivalent to the sampled rotation matrix \mathbf{W}_0 . The permutation matrix \mathbf{P} remains fixed throughout the optimization.

Conditional coupling layer. The conditional coupling layer resembles that of C-Flow (except that we use the whole source coupling output rather than half of it), where the network in the coupling layer takes input from both source and target sides:

$$\mathbf{x}_1^{\text{target}}, \mathbf{x}_2^{\text{target}} = \text{split}(\mathbf{x}^{\text{target}}) \quad (7)$$

$$\mathbf{o}_1, \mathbf{o}_2 = \text{CN}(\mathbf{x}_2^{\text{target}}, \mathbf{x}^{\text{source}}) \quad (8)$$

$$\mathbf{s} = \text{Sigmoid}(\mathbf{o}_1 + \mathbf{2}) \quad (9)$$

$$\mathbf{t} = \mathbf{o}_2 \quad (10)$$

where the `split` operation splits the input tensor along the channel dimension, $\mathbf{x}^{\text{source}}$ is the output of the corresponding coupling layer in the source Glow, and $\mathbf{x}^{\text{target}}$ is the output of the preceding 1×1 convolution in the target Glow. \mathbf{s} and \mathbf{t} are the affine coupling parameters. The conditioning network inputs are concatenated channel-wise.

The objective function for the model has the same form as that of Dual-Glow:

$$\frac{1}{N} \left[- \sum_{n=1}^N \lambda \log p_{\theta}(\mathbf{x}_a^{(n)}) - \sum_{n=1}^N \log p_{\phi}(\mathbf{x}_b^{(n)} | \mathbf{x}_a^{(n)}) \right] \quad (11)$$

where θ are the parameters of the source Glow and ϕ are the parameters of the target Glow. We note that there is one model (and term) for unconditional image generation in the source domain, coupled with a second model (and term) for conditional image generation in the target domain. With the tuning parameter λ set to unity, Equation 11 is the joint likelihood of the source-target image pair $(\mathbf{x}_a, \mathbf{x}_b)$, and puts equal emphasis on learning to generate (and to normalize/analyze) both images. In the limit $\lambda \rightarrow \infty$, we

will learn an unconditional model of source images only. Using a λ below 1, however, helps the optimization process instead put more importance on the conditional distribution, which is our main priority in image-to-image translation. This “exchange rate” between bits of information in different domains is reminiscent of the tuning parameter in the information bottleneck principle [38].

5 Experiments

This section reports our findings from applying the proposed model to the Cityscapes dataset from [6]. Each data instance is a photo of a street scene that has been segmented into objects of 30 different classes, such as road, sky, buildings, cars, and pedestrians. 5000 of these images come with fine per-pixel class annotations of the image, a so called *segmentation mask*. We used the data splits provided by the dataset (2975 training and 500 validation images), and trained a number of different models to generate street-scene images conditioned on their segmentation masks.

A common way to evaluate the quality of images generated based on the Cityscapes dataset is to apply well-known pre-trained classifiers such as FCN [22] and (here) PSP-Net [43] to synthesized images (as done by [15, 41]). The idea is that if a synthesized image is of high quality, a classifier trained on real data should be able to successfully classify different objects in the synthetic image, and thus produce an estimated segmentation mask that closely agrees with the ground-truth segmentation mask. For likelihood-based models we also consider the conditional bits per dimension (BPD), $-\log_2 p(\mathbf{x}_b|\mathbf{x}_a)$, as a measure of how well the conditional distribution learned by the model matches the real conditional distribution, when tested on held-out examples.

Implementation details. Our main experiments were performed on images from the Cityscapes data down-sampled to 256×256 pixels (higher than C-Flow [31] that uses 64×64 resolution). The Full-Glow model was implemented in PyTorch [29] and trained using the Adam optimizer [18] with a learning rate of 10^{-4} and a batch size of 1. The conditioning networks (CN) for the actnorm and 1×1 convolution in our model consisted of three convolutional layers followed by four fully connected layers. The CN for the coupling layer had two convolutional layers. Network weights were initialized as described in Section 4. We used $\lambda = 10^{-4}$ in the objective function Equation 11. Training was consistently stable and monotonic; see the loss curve in the supplement. Our implementation could be found at: <https://github.com/MoeinSorkhei/glow2>.

5.1 Quantitative comparison with other models

We compare the performance of our model against C-Glow [23] and Dual-Glow [36] (two previously proposed Glow-based models) and pix2pix [15] a widely used GAN-based model for image-to-image translation.

Since C-Glow was proposed to deal with low-resolution images, the authors exploited deep conditioning network in their model. We could not use equally deep conditioning networks in this task because the images we would like to generate are of higher resolution (256×256). To enable valid comparisons, we trained two versions of their model. In the first version, we allowed the conditioning networks to be deeper

while the Glow itself is shallower (3 Blocks each with 8 Flows). In the second version, the Glow model is deeper (4 Blocks and each with 16 Flows) but the conditioning networks are shallower. More details about the models and their hyper-parameters can be found in the supplementary material. Note that the Glow models in C-Glow version 2, Dual-Glow, and our model are all equally deep (4 Blocks and each having 16 Flows). All models, including Full-Glow, were trained for ~ 45 epochs using the same training procedure described earlier.³

We sampled from each trained model 3 times on the validation set, evaluated the synthesized images using PSPNet [43], and calculated the mean and standard deviation of the performance (denoted by \pm). The metrics used for evaluation are mean pixel accuracy, mean class accuracy, and mean intersection over union (IoU), as formulated in [22]. The mean pixel accuracy essentially computes mean accuracy over all the pixels of an image (which could easily be dominated by the sky, trees, and large objects that are mostly classified correctly.) Mean class accuracy, however, calculates the accuracy over the pixels of each class, and then takes average over different classes (where all classes are treated equally). Finally, mean class IoU calculates for each class the intersection over union for the objects of that class segmented in the synthesized image compared against the objects in the ground-truth segmentation. Optimally, this number should be 1, signifying complete overlap between segmented and ground-truth objects.

Quantitative results of applying each model to the Cityscapes dataset in the label \rightarrow photo direction could be seen in Table 1. The results show that street scene images generated by Full-Glow are of higher quality from the viewpoint of semantic segmentation. The noticeable difference in classification performance confirms that the objects in the images generated by our model are more easily *distinguishable* by the off-the-shelf semantic classifier. We attribute this to the fact that making the model fully conditional enables the target Glow to exploit the information available in the source image and to synthesize an image that follows the structure most.

³ We used these repositories 1, 2, 3 to obtain the official implementations of C-Glow, Dual-Glow, and pix2pix. We did not find any official implementation for C-Flow.

Table 1: Comparison of different models on the Cityscapes dataset for label \rightarrow photo image synthesis.

Model	Cond.	Mean	Mean	Mean
	BPD	pixel acc.	class acc.	class IoU
C-Glow v.1 [23]	2.568	35.02 ± 0.56	12.15 ± 0.05	7.33 ± 0.09
C-Glow v.2 [23]	2.363	52.33 ± 0.46	17.37 ± 0.21	12.31 ± 0.24
Dual-Glow [36]	2.585	71.44 ± 0.03	23.91 ± 0.19	18.96 ± 0.17
pix2pix [15]	—	60.56 ± 0.11	22.64 ± 0.21	16.42 ± 0.06
Our model	2.345	73.50 ± 0.13	29.13 ± 0.39	23.86 ± 0.30
<i>Ground-truth</i>	—	<i>95.97</i>	<i>84.31</i>	<i>77.30</i>

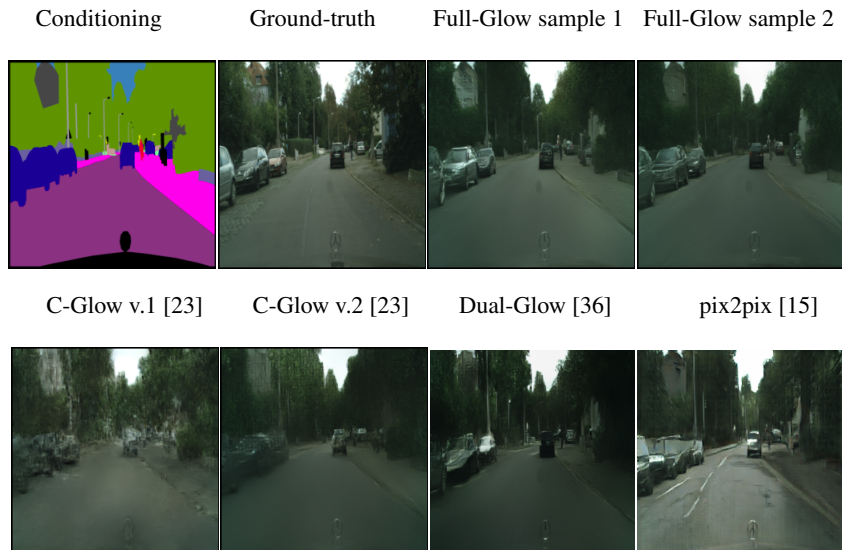


Fig. 2: Visual samples from different models. Samples from likelihood-based models are taken with temperature 0.7. Please zoom in to see more details.

5.2 Visual comparison with other models

It is interesting to see how samples generated by different models are different visually. Figure 2 illustrates samples from different models given the same condition. An immediate observation is that C-Glow v.1 [23] (which has deeper conditioning networks but shallower Glow) is essentially unable to generate any meaningful image. Dual-Glow [36], however, is able to generate plausible images. Samples generated by pix2pix [15] exhibit vibrant colors (especially for the buildings) but the important objects (such as cars) that constitute the general structure of the image are sometimes distorted. We believe this is the reason behind scoring low with the semantic classifier. Respecting the structure seems to be more important than having vibrant colors in order to get higher classification accuracy. Different samples taken from our model show the benefit of flow-based models in synthesizing different images every time we sample. Most of the difference can be seen in the colors of the objects such as cars.

Generally, the samples generated by likelihood-based models appear dimmed to some extent. This is in contrast with GAN-based samples, which often have realistic colors. This is probably related to the fundamental difference in the optimization process of the two categories. GAN-based models tend to collapse to regions of data where only plausible samples could have come from, and they might not have support over other data regions [11] – also seen as lack of diversity in their samples. In contrast, likelihood-based models try to learn a distribution that has support over wider data regions while maximizing the probability of the available datapoints. The latter approach seems to result in generating samples that are diverse but have somewhat muted colors (especially with lower temperatures).

Table 2: Effect of temperature T evaluated using a pre-trained PSPNet [43]. Each column lists the mean over repeated image samples.

T	Pixel acc.	Class acc.	Class IoU
0.0	27.83 \pm 0	12.48 \pm 0	7.69 \pm 0
0.1	29.47 \pm 0.04	13.15 \pm 0.05	8.33 \pm 0.01
0.2	34.97 \pm 0.13	15.83 \pm 0.13	10.65 \pm 0.01
0.3	43.36 \pm 0.08	20.22 \pm 0.05	14.69 \pm 0.02
0.4	55.26 \pm 0.12	25.90 \pm 0.15	20.01 \pm 0.11
0.5	69.83 \pm 0.01	31.66 \pm 0.16	25.82 \pm 0.12
0.6	79.38 \pm 0.05	34.90 \pm 0.28	29.37 \pm 0.23
0.7	82.14 \pm 0.09	35.51 \pm 0.23	29.98 \pm 0.21
0.8	83.52 \pm 0.04	35.94 \pm 0.19	30.67 \pm 0.17
0.9	82.27 \pm 0.05	35.24 \pm 0.37	29.85 \pm 0.36
1.0	73.50 \pm 0.13	29.13 \pm 0.39	23.86 \pm 0.30

5.3 Effect of temperature

As noted above, likelihood-based models such as Glow [19] generally tend to *overestimate* the variability of the data distribution [25, 37], hence occasionally generating implausible output samples. A common way to circumvent this issue is to reduce the diversity of the output at generation time. For flows, this can be done by reducing the standard deviation of the base distribution by a factor T (known as the temperature). While $T = 1$ corresponds to sampling from the estimated maximum-likelihood distribution, reducing T generally results in the output distribution becoming concentrated on a core region of especially-probable output samples. Similar ideas are widely used not only in flow-based models (cf. [19]) but also in other generative models such as GANs, VAEs, and Transformer-based language models [3, 5, 13, 40].

We investigated the effect of temperature by evaluating the performance of the model on samples generated at different temperatures (instead of $T = 1$ as in previous experiments). We sampled on the validation set 3 times with the trained model and evaluated using the PSPNet semantic classifier as before. The results are reported in Table 2 and suggest that the optimal temperature is around 0.8 for this task. That setting strikes a compromise where colors are vibrant while object structure is well-maintained, enabling the classifier to well understand the objects in the synthesized image. Also note the small standard deviation at lower temperatures, which agrees with our expectation that inter-sample variability would be small at low temperature. Example images generated at different temperatures are provided in the supplementary material.

5.4 Content transfer

Style transfer in general is an interesting application for synthesizing new images, which incorporates transferring the style of an image into a new structure (condition). In this experiment, we try to transfer the *content* of a real photo to a new structure (segmentation). Previous work [1, 31] has performed similar experiments with flow-based models but either on a different dataset or in very low resolution (64×64). We

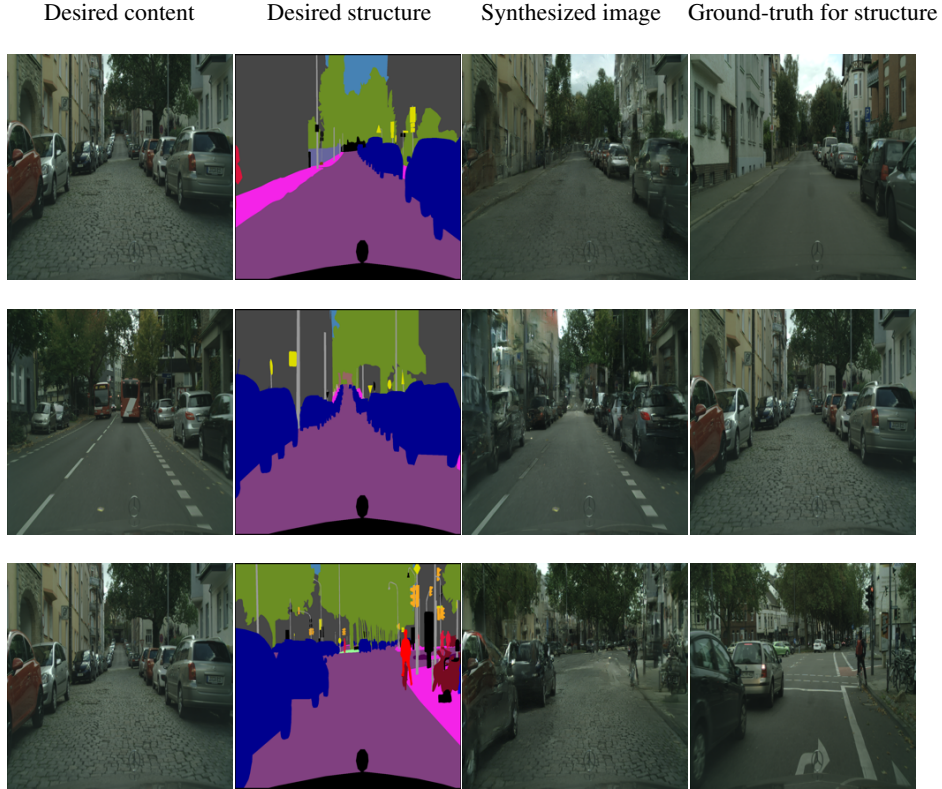


Fig. 3: Examples of applying a desired content to a desired structure. Please zoom in to see more details.

demonstrate how the learned representation enables us to synthesize an image given a desired content and a desired structure in relatively high resolution while maintaining the details of the content.

Suppose \mathbf{x}_b^1 is the image with the desired content (with \mathbf{x}_a^1 being its segmentation) and \mathbf{x}_a^2 is the new segmentation to which we are interested to transfer the content. We can take the following steps ($\mathbf{g}_\theta(\cdot)$ and $\mathbf{f}_\phi(\cdot)$ are the forward functions in the source and target Glows respectively):

1. Extract representation of the desired content: $\mathbf{z}_b^1 = \mathbf{f}_\phi(\mathbf{x}_b^1 | \mathbf{g}_\theta(\mathbf{x}_a^1))$.
2. Apply the content to the new segmentations: $\mathbf{x}_b^{\text{new}} = \mathbf{f}_\phi^{-1}(\mathbf{z}_b^1 | \mathbf{g}_\theta(\mathbf{x}_a^2))$.

Figure 3 shows examples of transferring the content of an image to another segmentation. We can see that the model is able to successfully apply the content of large objects such as buildings, trees, cars to the desired image structure. So often a given content and a given structure may not agree with each other. For instance, when there are cars in the content which are missing in the segmentation or vice versa. This kind of mismatch is quite common for content transfer on Cityscapes images as these im-



Fig. 4: Higher resolution samples of our model taken with temperature 0.9. Left: conditioning, middle-right: samples 1, 2, 3. Please zoom in to see more details.



Fig. 5: Higher resolution samples of our model taken with temperature 0.9. Left: conditioning, middle-right: samples 1, 2, 3. Please zoom in to see more details.

ages have a lot of objects placed in different positions. In such a case, the model tries to respect the structure as much as possible while filling it with the given content. The results show that content transfer is so useful in data augmentation since, given the desired content, the model can fill the structure with coherent information which makes the output image much realistic. This technique could practically be applied to any content and any structure (provided that they do not mismatch completely), hence enabling one to synthesize many more images.

5.5 Higher-resolution samples

In order to see how expressive the model is at even higher resolutions, we trained the proposed model on 512×1024 images. Example output images from the trained model are provided in Figures 4 and 5. It is known that higher temperatures show more diversity but the structures become somewhat distorted. We chose to sample with temperature 0.9 for these higher-resolution images. The diversity between multiple samples is especially obvious looking at the cars. Additional high-resolution samples are available in the supplementary material.

6 Conclusions

In this paper, we proposed a fully conditional Glow-based architecture for more realistic conditional street-scene image generation. We quantitatively compared our method against previous work and observed that our improved conditioning allows for generating images that are more interpretable by the semantic classifier. We also used the architecture to synthesize higher-resolution images in order to better show diversity of samples and the capabilities of Glow at higher resolutions. In addition, we demonstrated how new meaningful images could be synthesized based on a desired content and a desired structure, which is a compelling option for high-quality data augmentation.

Future work. While our results are promising, further work remains to be done in order to close the gap between synthetic images and real-world photographs, for instance by adding self-attention [12] and by leveraging approaches that combine the strong points of Full-Glow with the advantages of GANs, e.g., following [11, 24]. That said, we believe the quality of the synthetic images is already at a level where it also is worth exploring their utility in training systems for autonomous cars and other mobile agents, which remains to be observed in future works.

References

1. Ardizzone, L., Kruse, J., Lüth, C., Bracher, N., Rother, C., Köthe, U.: Conditional invertible neural networks for diverse image-to-image translation. In: DAGM German Conference on Pattern Recognition. pp. 373–387 (2020)
2. Ardizzone, L., Lüth, C., Kruse, J., Rother, C., Köthe, U.: Guided image generation with conditional invertible neural networks. arXiv preprint arXiv:1907.02392 (2019)
3. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2019)
4. Brostow, G.J., Fauqueur, J., Cipolla, R.: Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters* **30**(2), 88–97 (2009)
5. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: *Neural Information Processing Systems*. pp. 1877–1901 (2020)
6. 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: *IEEE Conference on Computer Vision and Pattern Recognition* (2016)
7. Dinh, L., Krueger, D., Bengio, Y.: NICE: Non-linear independent components estimation. arXiv preprint arXiv:1410.8516 (2014)
8. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using Real NVP. In: *International Conference on Learning Representations* (2017)
9. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research* **32**(11), 1231–1237 (2013)
10. 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* (2014)
11. Grover, A., Dhar, M., Ermon, S.: Flow-GAN: Combining maximum likelihood and adversarial learning in generative models. In: *AAAI Conference on Artificial Intelligence* (2018)
12. Ho, J., Chen, X., Srinivas, A., Duan, Y., Abbeel, P.: Flow++: Improving flow-based generative models with variational dequantization and architecture design. arXiv preprint arXiv:1902.00275 (2019)
13. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: *International Conference on Learning Representations* (2020)
14. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)
15. 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* (2017)

16. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
17. Kim, S., Lee, S.g., Song, J., Kim, J., Yoon, S.: FloWaveNet: A generative flow for raw audio. In: International Conference on Machine Learning (2019)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
19. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. In: Neural Information Processing Systems (2018)
20. Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. In: International Conference on Learning Representations (2014)
21. Kumar, M., Babaeizadeh, M., Erhan, D., Finn, C., Levine, S., Dinh, L., Kingma, D.: VideoFlow: A flow-based generative model for video. arXiv preprint arXiv:1903.01434 (2019)
22. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (2015)
23. Lu, Y., Huang, B.: Structured output learning with conditional generative flows. In: AAAI Conference on Artificial Intelligence (2020)
24. Lucas, T., Shmelkov, K., Alahari, K., Schmid, C., Verbeek, J.: Adaptive density estimation for generative models. In: Neural Information Processing Systems. pp. 11993–12003 (2019)
25. Minka, T.: Divergence measures and message passing. Tech. Rep. MSR-TR-2005-173, Microsoft Research, Cambridge, UK (2005)
26. van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., Kavukcuoglu, K.: Conditional image generation with pixelCNN decoders. In: Neural Information Processing Systems (2016)
27. Oord, A.v.d., Kalchbrenner, N., Kavukcuoglu, K.: Pixel recurrent neural networks. arXiv preprint arXiv:1601.06759 (2016)
28. Papamakarios, G., Nalisnick, E., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing flows for probabilistic modeling and inference. arXiv preprint arXiv:1912.02762 (2019)
29. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch. In: NIPS 2017 Workshop Autodiff (2017)
30. Prenger, R., Valle, R., Catanzaro, B.: WaveGlow: A flow-based generative network for speech synthesis. In: IEEE International Conference on Acoustics, Speech and Signal Processing (2019)
31. Pumarola, A., Popov, S., Moreno-Noguer, F., Ferrari, V.: C-Flow: Conditional generative flow models for images and 3D point clouds. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
32. Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with VQ-VAE-2. In: Neural Information Processing Systems (2019)
33. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: International Conference on Machine Learning (2014)
34. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: European Conference on Computer Vision (2016)
35. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
36. Sun, H., Mehta, R., Zhou, H.H., Huang, Z., Johnson, S.C., Prabhakaran, V., Singh, V.: DUAL-GLOW: Conditional flow-based generative model for modality transfer. In: IEEE International Conference on Computer Vision (2019)
37. Theis, L., van den Oord, A., Bethge, M.: A note on the evaluation of generative models. In: International Conference on Learning Representations (2016)

38. Tishby, N., Pereira, F.C., Bialek, W.: The information bottleneck method. In: Proceedings of the Allerton Conference on Communication, Control and Computing. vol. 37, pp. 368–377 (2000)
39. Tsirikoglou, A., Kronander, J., Wrenninge, M., Unger, J.: Procedural modeling and physically based rendering for synthetic data generation in automotive applications. arXiv preprint arXiv:1710.06270 (2017)
40. Vahdat, A., Kautz, J.: NVAE: A deep hierarchical variational autoencoder. arXiv preprint arXiv:2007.03898 (2020)
41. 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: IEEE Conference on Computer Vision and Pattern Recognition (2018)
42. Zhang, H., Dauphin, Y.N., Ma, T.: Fixup initialization: Residual learning without normalization. In: International Conference on Learning Representations (2019)
43. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: IEEE Conference on Computer Vision and Pattern Recognition (2017)
44. Zhao, S., Song, J., Ermon, S.: Towards deeper understanding of variational autoencoding models. arXiv preprint arXiv:1702.08658 (2017)
45. 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 (2017)