



MIT Open Access Articles

Meta-Sim: Learning to Generate Synthetic Datasets

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	2020. "Meta-Sim: Learning to Generate Synthetic Datasets." Proceedings of the IEEE International Conference on Computer Vision, 2019-October.
As Published	10.1109/ICCV.2019.00465
Publisher	IEEE
Version	Original manuscript
Citable link	https://hdl.handle.net/1721.1/137178
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Meta-Sim: Learning to Generate Synthetic Datasets

Amlan Kar^{1,2,3} Aayush Prakash¹ Ming-Yu Liu¹ Eric Cameracci¹ Justin Yuan¹
Matt Rusiniak¹ David Acuna^{1,2,3} Antonio Torralba⁴ Sanja Fidler^{1,2,3*}
¹NVIDIA ²University of Toronto ³Vector Institute ⁴MIT

Abstract

Training models to high-end performance requires availability of large labeled datasets, which are expensive to get. The goal of our work is to automatically synthesize labeled datasets that are relevant for a downstream task. We propose *Meta-Sim*, which learns a generative model of synthetic scenes, and obtain images as well as its corresponding ground-truth via a graphics engine. We parametrize our dataset generator with a neural network, which learns to modify attributes of scene graphs obtained from probabilistic scene grammars, so as to minimize the distribution gap between its rendered outputs and target data. If the real dataset comes with a small labeled validation set, we additionally aim to optimize a meta-objective, i.e. downstream task performance. Experiments show that the proposed method can greatly improve content generation quality over a human-engineered probabilistic scene grammar, both qualitatively and quantitatively as measured by performance on a downstream task. Webpage: <https://nv-tlabs.github.io/meta-sim/>

1. Introduction

Data collection and labeling is a laborious, costly and time consuming venture, and represents a major bottleneck in most current machine learning pipelines. To this end, synthetic content generation [5, 35, 10, 33] has emerged as a promising solution since all ground-truth comes for free – via the graphics engine. It further enables us to train and test our models in virtual environments [37, 7, 46, 23, 40] before deploying to the real world, which is crucial for both scalability and safety. Unfortunately, an important performance issue arises due to the domain gap existing between the synthetic and real-world domains.

Addressing the domain gap issue has led to a plethora of work on synthetic-to-real domain adaptation [17, 27, 51, 9, 42, 33, 44]. These techniques aim to learn domain-invariant features and thus more transferrable models. One of the mainstream approaches is to learn to stylize syn-




Figure 1. Meta-Sim is a method to generate synthetic datasets that bridge the *distribution gap* between real and synthetic data and are optimized for *downstream task performance*

thetic images to look more like those captured in the real-world [17, 27, 48, 30, 18]. As such, these models address the *appearance gap* between the synthetic and real-world domains. They share the assumption that the domain gap is due to the differences that are fairly low level.

Here, we argue that domain gap is also due to a *content gap*, arising from the fact that the synthetic content (e.g. layout and types of objects) mimics a limited set of scenes, not necessarily reflecting the diversity and distribution of objects of those captured in the real world. For example, the Virtual KITTI [10] dataset was created by a group of engineers and artists, to match object locations and poses in KITTI [12] which was recorded in Karlsruhe, Germany. But what if the target city changes to Tokyo, Japan, which has much heavier traffic and many more high-rise buildings? Moreover, what if the downstream task that we want to solve changes from object detection to lane estimation or rain drop removal? Creating synthetic worlds that ensure realism and diversity for any desired task requires significant effort by highly-qualified experts and does not scale to the fast demand of various commercial applications.

In this paper, we aim to learn a generative model of synthetic scenes that, by exploiting a graphics engine, produces *labeled* datasets with a content distribution matching that of imagery captured in the desired real-world datasets. Our *Meta-Sim* builds on top of probabilistic scene grammars which are commonly used in gaming and graphics to create diverse and valid virtual environments. In particular, we assume that the structure of the scenes sampled from the grammar are correct (e.g. a driving scene has a road and

*Correspondence to amlan@cs.toronto.edu, sfidler@nvidia.com

cars), and learn to modify their attributes. By modifying locations, poses and other attributes of objects, Meta-Sim gains a powerful flexibility of adapting scene generation to better match real-world scene distributions. *Meta-Sim* also optimizes a meta objective of adapting the simulator to improve downstream real-world performance of a Task Network trained on the datasets synthesized by our model. Our learning framework optimizes several objectives using approximated gradients through a non-differentiable renderer.

We validate our approach on two toy simulators in controlled settings, where Meta-Sim is shown to excel at bridging the distribution gaps. We further showcase Meta-Sim on adapting a probabilistic grammar akin to SDR [33] to better match a real self-driving dataset, leading to improved content generation quality, as measured by sim-to-real performance. To the best of our knowledge, Meta-Sim is the first approach to enable dataset and task specific synthetic content generation, and we hope that our work opens the door to more adaptable simulation in the future.

2. Related Work

Synthetic Content Generation and Simulation. The community has been investing significant effort in creating high-quality synthetic content, ranging from driving scenes [37, 10, 35, 7, 33, 45], indoor navigation [46], household robotics [34, 23], robotic control [43], game playing [4], optical flow estimation [5], and quadcopter control and navigation [40]. While such environments are typically very realistic, they require qualified experts to spend a huge amount of time to create these virtual worlds. Domain Randomization (DR) is a cheaper alternative to such photo-realistic simulation environments [39, 42, 33]. The DR technique generates a large amount of diverse scenes by inserting objects in random locations and poses. As a result, the distribution of the synthetic scenes is very different to that of the real world scenes. We, on the other hand, aim to align the synthetic and real distributions through a direct optimization on the attributes and through a meta objective of optimizing for performance on a down-stream task.

Procedural modeling and probabilistic scene grammars are an alternative approach to content generation¹, which are able to produce worlds at the scale of full cities², and mimic diverse 3D scenes for self-driving³. However, the parameters for generating the distributions that control how a scene is generated need to be manually specified. This is not only tedious but also error-prone. There is no guarantee that the specified parameters can generate distributions that faithfully reflect real world distributions. [24, 32] use such probabilistic programs to invert the generative process and infer a program given an image, while we

aim to learn the generative process itself from real data.

Domain Adaptation aims at addressing the domain gap, *i.e.* the mismatch between the distribution of data used to train a model and the distribution of data that the model is expected to work with. From synthetic to real, two kinds of domain gaps arise: the appearance (style) gap and the content (layout) gap. Most existing work [17, 27, 51, 9, 48, 30, 18] tackle the former by using generative adversarial networks (GANs) [13] to transform the appearance distribution of the synthetic images to look more like that of the real images. Others [17, 27] add additional task based constraints to ensure that the layout of the stylized images remain the same. Other techniques use pseudo label based self learning [51] and student-teacher networks [9] for domain adaptation. Our work is an early attempt to tackle the second kind of domain gap – the content gap. We note that the appearance gap is orthogonal to the content gap, and prior art could be directly plugged into our method.

Optimizing Simulators. Louppe *et al.* [31] attempt to optimize non-differentiable simulators with the key difference being in the method and the end goal. They optimize using a variational upperbound of a GAN-like objective to produce samples representative of a target distribution. We, on other hand, use the MMD [15] distance metric for comparing distributions and also optimize a meta objective to produce samples suitable for a downstream task. [6] learn to optimize simulator parameters for robotic control tasks, where trajectories between the real and simulated robot can be directly compared. [38] optimize high level exposed parameters by optimizing for downstream task performance using Reinforcement Learning. We, on the other hand, optimize low level scene parameters (at the level of every object) while also learning to match distributions along with optimizing downstream task performance. Ganin *et al.* [11] attempt to synthesize images by learning to generate even lower-level programs (at the level of brush strokes) that a graphics engine can interpret to generate realistic looking images, as measured by a trained discriminator.

3. Meta-Sim

In this section, we introduce *Meta-Sim*. Given a dataset of real imagery X_R and a task T (*e.g.* object detection), our goal is to *synthesize* a training dataset $D_T = (X_T, Y_T)$ with X_T synthesized imagery that resembles the given real imagery, and Y_T the corresponding ground-truth for task T . To simplify notation, we omit subscript T from here on.

We parametrize data synthesis with a neural network, *i.e.* $D(\theta) = (X(\theta), Y(\theta))$. Our goal in this paper is to learn the parameters θ such that the distribution of $X(\theta)$ matches that of X_R (real imagery). Optionally, if the real dataset comes with a small validation set V that is labeled for task T , we additionally aim to optimize a meta-objective, *i.e.* downstream task performance. The latter assumes we also have a

¹<https://www.sidefx.com/>

²<https://www.esri.com/en-us/arcgis/products/esri-cityengine/overview>

³<https://www.paralleldomain.com/>




Figure 2. Overview of Meta-Sim: The goal is to learn to transform samples coming from a probabilistic grammar with a distribution transformer, aiming to minimize the *distribution gap* between simulated and real data and maximize *sim-to-real performance*

trainable task solving module (*i.e.* another neural network), the performance of which we want to maximize by training it on our generated training data. We refer to this module as a Task Network, which will be treated as a black box in our work. Note that Meta-Sim has parallels to Neural Architecture Search [50], where our search is over the input datasets to a fixed neural network instead of a search over the neural network architecture given fixed data.

Image Synthesis vs Rendering. Generative models of pixels have only recently seen success in generating realistic high resolution images [3, 19]. Extracting task specific ground-truth (eg: segmentation) from them remains a challenge. Conditional generative models of pixels condition on input images and transform their appearance, producing compelling results. However, these methods assume ground truth labels remain unchanged, and thus are limited in their *content* (structural) variability. In Meta-Sim we aim to *learn a generative model of synthetic 3D content*, and obtain D via a graphics engine. Since the 3D assets come with semantic information (*i.e.*, we know an asset is a *car*), compositing or modifying the synthetic scenes will still render perfect ground-truth. The main challenge is to learn the 3D scene composition by optimizing solely the distribution mismatch of rendered with real imagery. The following subsections layout Meta-Sim in detail and are structured as follows: Sec. 3.1 introduces the representation of parametrized synthetic worlds, while Sec. 3.2 describes our learning framework.

3.1. Parametrizing Synthetic Scenes

Scene Graphs are a common way to represent 3D worlds in gaming/graphics. A scene graph represent elements of a scene in a concise hierarchical structure, with each element having a set of attributes (eg. class, location, or even the id of a 3D asset from a library) (see Fig. 3). The hierarchy defines parent-child dependencies, where the attributes of the child elements are typically defined relative to the parent’s, allowing for an efficient and natural way to create and modify scenes. The corresponding image and pixel-level annotations can be rendered easily by placing objects as described in the scene graph.

In order to generate *diverse* and *valid* 3D worlds, the typical approach is to specify the generative process of the graph by a *probabilistic scene grammar* [49]. For example, to generate a traffic scene, one might first lay out the centerline of the road, add parallel lanes, position aligned cars on each lane, etc. The structure of the scene is defined by the grammar, while the attributes are typically sampled from parametric distributions, which require careful tuning.

In our work, we assume access to a probabilistic grammar from which we can sample initial scene graphs. We assume the *structure* of each scene graph is correct, *i.e.* the driving scene has a road, sky, and a number of objects. This is a reasonable assumption, given that inferring structure (inverse graphics) is known to be a hard problem. Our goal is to modify the *attributes* of each scene graph, such that the transformed scenes, when rendered, will resemble the distribution of the real scenes. By modifying the attributes, we give the model a powerful flexibility to change objects’ locations, poses, colors, asset ids, etc. This amounts to learning a conditional generative model, which, by conditioning on an input scene graph transforms its node attributes. In essence, we keep the *structure* generated by the probabilistic grammar, but transform the distribution of the *attributes*. Thus, our model acts as a *Distribution Transformer*.

Notation. Let P denote the probabilistic grammar from which we can sample scene graphs $s \sim P$. We denote a single scene graph s as a set of vertices s_V , edges s_E and attributes s_A . We have access to a renderer R , that can take in a scene graph s and generate the corresponding image and ground truth, $R(s) = (x, y)$. Let G_θ refer to our Distribution Transformer, which takes an input scene graph s and outputs a scene graph $G_\theta(s)$, with transformed attributes but the same structure, *i.e.* $G_\theta(s = [s_V, s_E, s_A]) = [s_V, s_E, G_\theta(s_A)]$. Note that by sampling many scene graphs, transforming their attributes, and rendering, we obtain a synthetic dataset $D(\theta)$.

Architecture of G_θ . Given the graphical structure of scene graphs, modeling G_θ via a Graph Neural Network is a natural choice. In particular, we use Graph Convolutional Networks (GCNs) [22]. We follow [47] and use a graph convolutional layer that utilizes two different weight matrices to capture top-down and bottom-up information flow separately. Our model makes per node predictions *i.e.* generates transformed attributes $G_\theta(s_A)$ for each node in s_V .

Mutable Attributes: We input to G_θ all attributes s_A , but we might want to only modify specific attributes and trust the probabilistic grammar P on the rest. For example, in Fig. 3 we may not want to change the heights of houses, or width of the sidewalks, if our final task is car detection. This reduces the number of exposed parameters our model is tasked to tune thus improving training time and complexity. Therefore, in the subsequent parts, we assume we have




Figure 3. Simple scene graph example for a driving scene.



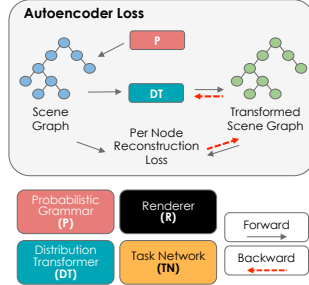


Figure 4. Illustration of different losses used in *Meta-Sim*, including forward and backward pass control flow for each step. We indicate transformed attributes of a scene graph by changing colors of the nodes.

a subset of attributes per node $v \in s_V$ which are mutable (modifiable), denoted by $s_{A,mult}(v)$. From here onwards, it is assumed that only the mutable attributes in $s_{A,mult}(v) \forall v$ are changed by G_θ ; others remain the same as in s .

3.2. Training Meta-Sim

We now introduce our learning framework. Since our learning problem is very hard and computationally intensive, we first pre-train our model using a simple autoencoder loss in Sec. 3.2.1. The distribution matching loss is presented in Sec 3.2.2, while meta-training is described in Sec 3.2.3. The overview of our model is given in Fig. 2, with the particular training objectives illustrated in Fig. 4.

3.2.1 Pre-training: Autoencoder Loss

A probabilistic scene grammar P represents a prior on how a scene should be generated. Learning this prior is a natural way to pre-train our *Distribution Transformer*. This amounts to training G_θ to perform the identity function *i.e.* $G_\theta(s) = s$. The input feature of each node is its attribute set (s_A), which is defined consistently across all nodes (see suppl.). Since s_A is composed of different categorical and continuous components, appropriate losses are used per feature component when training to reconstruct (*i.e.* cross-entropy loss for categorical attributes, and L1 loss for continuous attributes).

3.2.2 Distribution Matching

The first objective of training our model is to bring the distribution of the rendered images to be closer to the distribution of real imagery X_R . The Maximum Mean Discrepancy (MMD) [15] metric is a frequentist measure of the similarity of two distributions and has been used for training generative models [8, 29, 26] to match statistics of the generated distribution with the target distribution. An alternative, adversarial learning with discriminators, however, is known to suffer from mode collapse, and a general instability in training. Pixel-wise generative models with MMD have usually suffered from not being able to model high-frequency signals (resulting in blurry generations). Since our generative process goes through a renderer, we sidestep the issue altogether, and thus choose MMD for training stability.

We compute MMD in the feature space of an InceptionV3 [41] network (known as Kernel Inception Distance (KID) [2]). This feature extractor is denoted by the function ϕ . We use the kernel trick for the computation with a gaussian kernel $k(x_i, x_j)$. We refer the reader to [29] for more details. The *Distribution Matching* box in Fig. 4 shows the training procedure pictorially. Specifically, given scene graphs s_1, \dots, s_N sampled from P and target real images X_R , the squared MMD distance can be computed as,

$$\begin{aligned} \mathcal{L}_{MMD^2} = & \frac{1}{N^2} \sum_{i=1}^N \sum_{i'=1}^N k(\phi(X_\theta(s_i)), \phi(X_\theta(s_{i'}))) \\ & + \frac{1}{M^2} \sum_{j=1}^M \sum_{j'=1}^M k(\phi(X_R^j), \phi(X_R^{j'})) \\ & - \frac{1}{MN} \sum_{i=1}^N \sum_{j=1}^M k(\phi(X_\theta(s_i)), \phi(X_R^j)) \end{aligned} \quad (1)$$

where the image rendered from s is $X_\theta(s) = R(G_\theta(s))$.

Sampling from $G_\theta(s)$. For simplicity, we overloaded the notation $R(G_\theta(s))$, since R would actually require sampling from the prediction $G_\theta(s)$. In general, we assume independence across scenes, nodes and attributes, which lets each attribute of each node in the scene graph be sampled independently. While training with *MMD*, we mark categorical attributes in s_A as immutable. The predicted continuous attributes are directly passed as the sample.

Backprop through a Renderer. For optimizing the MMD loss, we need to backpropagate the gradient through the non-differentiable rendering function R . The gradient of $R(G_\theta(s))$ w.r.t. $G_\theta(s)$ can be approximated using the method of finite differences⁴. While this gives us noisy gradients, we found it sufficient to be able to train our models in practice, with the benefit of being able to use photorealistic rendering. We note that recent work on differentiable rendering [20, 28] could potentially benefit this work.

3.2.3 Optimizing Task Performance

The second objective of training the model G_θ is to generate data $R(G_\theta(S))$ given samples $S = \{s_1, \dots, s_K\}$ from

⁴computed by perturbing each attribute in the scene graph $G_\theta(s)$

Algorithm 1 Pseudocode for Meta-Sim’s meta training phase

```

1: Given:  $P, R, G_\theta$  ▷ Probabilistic grammar, Renderer, GCN Model
2: Given:  $\text{TaskNet}, X_R, V$  ▷ Task Model, Real Images, Target Validation Data
3: Hyperparameters:  $E_m, I_m, B_m$  ▷ Epochs, Iters, Batch size
4: while  $e_m \leq E_m$  do ▷ Meta training
5:    $loss = 0$ ;
6:    $data = []$ ;  $samples = []$ ; ▷ Caching data & samples generated in epoch
7:   while  $i_m \leq I_m$  do
8:      $S = G_\theta(\text{sample}(P, B_m))$ ; ▷ Generate  $B_m$  samples from  $P$ 
9:     and transform them
10:     $D = R(S)$ ; ▷ Render images, labels from  $S$ 
11:     $data += D$ ;  $samples += S$ ;
12:     $loss += \mathcal{L}_{MMD^2}(D, X_R)$ ; ▷ MMD between generated and
13:    target real images
14:  end while
15:   $\text{TaskNet} = \text{train}(\text{TaskNet}, data)$ ; ▷ Train  $\text{TaskNet}$  on  $data$ 
16:   $score = \text{test}(\text{TaskNet}, V)$ ; ▷ Test  $\text{TaskNet}$  on target val
17:   $loss += -(score - \text{moving\_avg}(score)) \cdot \log p_{G_\theta}(samples)$ 
▷ Eq. 3
18:   $G_\theta = \text{optimize}(G_\theta, loss)$ ; ▷ SGD step
19: end while

```

the probabilistic grammar P , such that a model trained on this data achieves best performance when tested on target data V . This can be interpreted as a meta-objective, where the input data must be optimized to improve accuracy on a validation set. We introduce a task network TaskNet to train using our data and to measure validation performance on. We train G_θ under the following objective,

$$\max_{\theta} \mathbb{E}_{S' \sim G_\theta(S)} [\text{score}(S')] \quad (2)$$

where $\text{score}(S')$ is the performance metric achieved on validation data V after training TaskNet on data $R(G_\theta(S'))$. The task loss in Eq. 2 is not differentiable w.r.t the parameters θ , since the score is measured using validation data and not S' . We use the REINFORCE score function estimator (which is an unbiased estimator of the gradient) to compute the gradients of Eq. 2. Reformulating the objective as a loss and writing the gradient gives,

$$\begin{aligned} \mathcal{L}_{task} &= -\mathbb{E}_{S' \sim G_\theta(S)} [\text{score}(S')] \\ \nabla_{\theta} \mathcal{L}_{task} &= -\mathbb{E}_{S' \sim G_\theta(S)} [\text{score}(S') \times \nabla_{\theta} \log p_{G_\theta}(S')] \end{aligned} \quad (3)$$

To reduce the variance of the gradient from the estimator above, we keep track of an exponential moving average of previous scores and subtract it from the current score [14]. We approximate the expectation using one sample from $G_\theta(S)$. The *Task Optimization* box in Fig. 4 provides a pictorial overview of the task optimization.

Sampling from $G_\theta(s)$. Eq. 3 requires us to be able to sample (and measure its likelihood) from our model. For continuous attributes, we interpret our model to be predicting the mean of a normal distribution per attribute, with a pre-defined variance. We use the reparametrization trick to sample from this normal distribution. For categorical attributes, it is possible to sample from a multinomial distribution from the predicted log probabilities per category. In this paper, we keep categorical attributes immutable.

Calculating $\log p_{G_\theta}(S')$. Since we assume independence across scenes, attributes and objects in the scene, the likelihood in Eq 3 for the full scene is simply factorizable,

$$\log p_G(S') = \sum_{s' \in S'} \sum_{v \in s'_V} \sum_{a \in s'_{A,mut}(v)} \log p_{G_\theta}(s'(v, a)) \quad (4)$$

where $s'(v, a)$ represents the attribute a at node v in a single scene s' in batch S' . Note that the sum is only over mutable attributes per node $s_{A,mut}(v)$. The individual log probabilities come from the defined sampling procedure.

Training Algorithm. The algorithm for training with Distribution Matching and Task Optimization is presented in Algorithm 1.

4. Experiments

We evaluate Meta-Sim on three target datasets with three different tasks. The subsequent sections follow a general structure where we first outline the desired task, the target data and the task network⁵. Then, we describe the probabilistic grammar that the *Distribution Transformer* utilizes for its input, and the associated renderer that generates labeled synthetic data. We show quantitative and qualitative results after training the task network using synthetic data generated by Meta-Sim. We show strong boosts in quantitative performance and noticeable qualitative improvements in content-generation quality.

The first two experiments presented are in a controlled setting, each with increasing complexity. The aim here is to probe Meta-Sim’s capabilities when the shift between the target data distribution and the input distribution is known. The input distribution refers to the distribution of the scenes generated by samples from the probabilistic grammar that our *Distribution Transformer* takes as input. Target data for these tasks is created by carefully modifying the parameters of the probabilistic program, which represents a known distribution gap that the model must learn.

4.1. MNIST

We first evaluate our approach on digit **classification** on MNIST-like data. The probabilistic grammar samples a background texture, one digit texture (image) from the MNIST dataset [25] (which has an equal probability for any digit), and then samples a rotation and location for the digit. The renderer transforms the texture based on the sampled transformation and pastes it onto a canvas.

Task Network. Our task network is a small 2-layer CNN followed by 3 fully connected layers. We apply dropout in the fully connected layers (with 50, 100 and 10 features). We verify that this network can achieve greater than 99% accuracy on the regular MNIST classification task. We do not use data-augmentation while training (in all following

⁵Task Network training details in suppl. material



Figure 5. Examples from the rotated-MNIST dataset



Figure 6. Examples from the rotated and translated MNIST experiments as well), as it might interfere with our model’s training by changing the configuration of the generated data, making the task optimization signal unreliable.

Rotating MNIST. In our first experiment, the probabilistic grammar generates input samples that are upright and centered, like regular MNIST digits (Fig 7 bottom). The target data V and X_R are images (at 32×32 resolution) where digits centered and always rotated by 90 degrees (Fig 5). Ideally, the model will learn this exact transformation, and rotate the digits in the input scene graph while keeping them in the same centered position.

Rotating and Translating MNIST. For the second experiment, we additionally add translation to the distribution gap, making the task harder for Meta-Sim. We generate V and X_R as 1000 images (at 64×64 resolution) where in addition to being rotated by 90 degrees, the digits are moved to the bottom left corner of the canvas (Fig 6). The input probabilistic grammar remains the same, *i.e.* one that generates centered and upright digits (Fig. 8 bottom).

Quantitative Results. Table 1 shows classification on the target datasets with the two distribution gaps described above. The target datasets are fresh samples from the target distribution (separate from V). Training directly on the input scenes (coming from the input probabilistic grammar *i.e.* generating upright and centered digits in this case) results in just above random performance. Our model recovers the transformation causing the distribution gap, and achieves greater than 99% classification accuracy.

Data	Rotation	Rotation + Translation
Prob. Grammar	14.8	13.1
Meta-Sim	99.5	99.3

Table 1. Classification performance on our MNIST with different distribution gaps in the data

Qualitative Results. Fig. 7 and Fig. 8 show generations from our model at the end of training, and compares with the input scenes. Clearly, the model has learnt to perfectly transform the input distribution to replicate the target distribution, corroborating our quantitative results.

4.2. Aerial Views (2D)

Next, we evaluate our approach on **semantic segmentation** of simulated aerial views of simple roadways. In the probabilistic grammar, we sample a background grass texture, followed by a (straight) road at some location and



Figure 7. (bottom) Input scenes, (top) Meta-Sim’s generated examples for MNIST with rotation gap



Figure 8. (bottom) Input scenes, (top) Meta-Sim’s generated examples for MNIST with rotation and translation gap





Figure 9. Example label and image from Aerial2D validation



rotation on the background. Next, we sample two cars with independent locations (constrained to be in the road by parametrizing in the road’s coordinate system), and rotations. In addition, we also sample a tree and a house randomly in the scene. Each object in the scene gets a random texture from a set of textures we collected for each object. We ended up with nearly 600 car, 40 tree, 20 house, 7 grass and 4 road textures. Overall, this grammar has more complexity than MNIST, due to the scene graphs having higher depth, more objects, and variability in appearance.

V and X_R are created by tuning the grammar parameters to generate a realistic aerial view. (Fig. 9). The input probabilistic grammar uses random parameters (Fig. 10) bottom.

Task Network. We use a small U-Net architecture [36] with a total of 7 convolutional layers (with 16 to 64 filters in the convolution layers) as our task-network.

Quantitative Results. Table 2 shows semantic segmentation results on the target set. The results show that Meta-Sim effectively transforms the outputs of the probabilistic grammar, even in this relatively more complex setup, and improves the mean IoU. Specifically, it learns to drastically reduce the gap in performance for cars and also improves performance on roads.



Figure 11. (bottom) input scenes, (top) Meta-Sim’s generated examples for Aerial semantic segmentation

Data	Car	Road	House	Tree	Mean
Prob. Grammar	30.0	93.1	98.3	99.7	80.3
MetaSim	86.7	99.6	95.0	99.5	95.2

Table 2. Semantic segmentation results (IoU) on Aerial2D

Qualitative Results. Qualitative results in Fig. 11 show that the model indeed learns to exploit the convolutional structure of the task network, by only learning to orient. This is sufficient to achieve its job since convolutions are translation equivariant, but not rotation equivariant.

4.3. Driving Scenes (3D)

After validating our approach on controlled experiments in a simulated setting, we now evaluate our approach for **object detection** on the challenging KITTI [12] dataset. KITTI was captured with a camera mounted on top of a car driving around the city of Karlsruhe in Germany. It consists of challenging traffic scenarios and scenes ranging from highways to urban to more rural neighborhoods. Contrary to the previous experiments, the distribution gap which we wish to reduce arises naturally here.

Current open-source self driving simulators [7, 40] do not offer the amount of low level control on object attributes that we require in our model. We thus turn to probabilistic grammars for road scenarios [33, 45]. Specifically, SDR [33] is a road scene grammar that has been shown to outperform existing synthetic datasets as measured by sim-to-real performance. We adopt a simpler version of SDR and implement portions of their grammar as our probabilistic grammar. Specifically, we remove support for intersections and side-roads for computational reasons. The exact parameters of the grammar used can be found in the supplementary material. We use the Unreal Engine 4 (UE4) [1] game engine for the 3D rendering from scene graphs. Fig. 12(left column) shows example renderings of scenes generated using our version of the SDR grammar. The grammar parameters were mildly tuned, since we aim to have our model do the heavy lifting in subsequent parts.

Task Network. We use Mask-RCNN [16] with a Resnet-50-FPN backbone (ImageNet initialized) detection head as our task network for object detection.

Experimental Setup. Following SDR [33], we use car detection as our task. Validation data V is formed by taking 100 random images (and their labels) from the KITTI train set. The rest of the training data (images only) forms X_R .

We report results on the KITTI val set. Training and finer details can be found in the supplementary material.

Complexity. To reduce training complexity (coming from rendering and numerical gradients), we train Meta-Sim to optimize specific parts of the scene sequentially. We first train to optimize attributes of cars. Next, we optimize car and camera parameters, and finally add parameters of context elements (buildings, pedestrians, trees) together to the training. Similarly, we decouple distribution and task training. We first train the above with MMD, and finally optimize all parameters above with the meta task loss.

Quantitative Results. Table 3 reports the average precision at 0.5 IoU of the task network trained using data generated from different methods, when tested on the KITTI val set. We see that training with Meta-Sim beats just using the data from the probabilistic grammar.

Data	Easy	Moderate	Hard
Prob. Grammar	63.7	63.7	62.2
MetaSim (Cars)	66.4	66.5	65.6
+ Camera	65.9	66.3	65.9
+ Context	65.9	66.3	66.0
+ Task Loss	66.7	66.3	66.2

Table 3. AP @ 0.5 IOU for car detection on the KITTI val dataset

Training the task network online with meta-sim and offline on final generated data results in similar final detection performance. This ensures the quality of the final generated data, since training while the transformation of data is being learned could be seen as data augmentation.

Bridging the appearance gap. We additionally add a state-of-the-art image-to-image translation network, MUNIT [18] to attempt to bridge the appearance gap between the generated synthetic images and real images. Table 4 shows training with image-to-image translation still leaves a performance gap between MetaSim and the baseline, confirming our *content gap* hypothesis.

Data	Easy	Moderate	Hard
Prob. Grammar	71.1	75.5	65.3
Meta-Sim	77.5	75.1	68.2

Table 4. Effect of adding image-to-image translation to bridge the appearance gap in generated images

Training on V . Since we have access to some labelled training data, a valid baseline is to train the models on V (100 images from KITTI train split). In Table. 5 we show the effect of only training with V and finetuning using V .

TaskNet Initialization	Easy	Moderate	Hard
ImageNet	61.2	62.0	60.7
Prob. Grammar	71.3	72.7	72.7
Meta-Sim (Task Loss)	72.4	73.9	73.9

Table 5. Effect of finetuning on V




Figure 12. (left) samples from our prob. grammar, (middle) Meta-Sim’s corresponding samples, (right) random samples from KITTI

Figure 13. Animation showing evolution of a scene through Meta-Sim training (animation works on Adobe Reader)




Figure 14. Car detection results (top) of task network trained with Meta-Sim vs (bottom) trained with our prob. grammar

Qualitative Results. Fig. 12 shows a few outputs of Meta-Sim compared to the inputs sampled from the grammar, along with a few random samples from KITTI(train). There is a noticeable difference, as Meta-Sim’s cars are well aligned with the road, and the distances between cars are meaningful. Also notice the small changes in camera, and the differences in the context elements, including houses, trees and pedestrians. The last row in Fig. 12 represents a failure case where Meta-Sim is unable to clear up a dense initial scene, resulting in collided cars. Interestingly, Meta-Sim perfectly overlaps two cars in the same image such that a single car is visible from the camera (first car in front of camera). This behavior is seen multiple times, indicating that the model learns to cheat its way to good data. Fig. 13 shows an animation representing how a scene evolves through meta-sim training. Elements are moved to final configurations sequentially, following our training procedure. We remind the reader that these scene configura-

tions are learned with only image/task level supervision. In Fig. 18, we show results of training the task network on our grammar vs. training with Meta-Sim. We observe fewer false positives and negatives than the baseline. Meta-Sim shows better recall and GT overlap. Both models lose in precision, arguably because of not training for similar classes like Bus/Truck which would be negative examples.

5. Conclusion

We proposed Meta-Sim, an approach that generates synthetic data to match real content distributions while optimizing performance on downstream (real) tasks. Our model learns to transform sampled scenes from a probabilistic grammar so as to satisfy these objectives. Experiments on two toy and one real task showcased that Meta-Sim generates quantitatively better and noticeably higher quality samples than the baseline. We hope this opens a new exciting

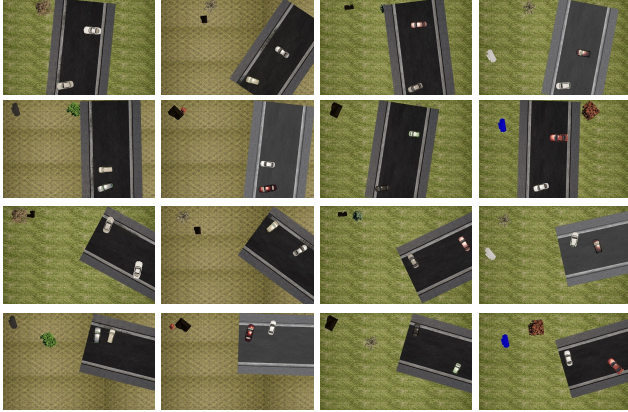


Figure 15. Meta-Sim generated samples (top two rows) and input scenes (bottom two rows)

direction for simulation in the computer vision community. Like any other method, it has its limitations. It relies on obtaining valid scene structures from a grammar, and hence is still limited in the kinds of scenes it can model. Inferring rules of the grammar from real images, learning to generate structure of scenes and introducing multimodality in the model are intriguing avenues for future work.

Acknowledgements: The authors would like to thank Shaad Boochoon, Felipe Alves, Gavriel State, Jean-Francois Lafleche, Kevin Newkirk, Lou Rohan, Johnny Costello, Dane Johnston and Rev Lebareddian for their help and support throughout this project.

6. Appendix

We provide additional details about Meta-Sim, as well as include more results.

6.1. Grammar and Attribute Sets

First, we describe the probabilistic grammars for each experiment in more detail, as well as attribute sets. We remind the reader that the probabilistic grammar defines the structure of the scenes, and is used to sample scene graphs that are input to Meta-Sim. The Distribution Transformer transforms the attributes on the nodes in order to optimize its objectives.

6.1.1 MNIST

Grammar. The sampling process for the probabilistic grammar used in the MNIST experiments is explained in the main paper, and we review it here. It samples a background texture and one digit texture (image) from the MNIST dataset [25] (which has an equal probability for any digit), and then samples a rotation and location for the digit, to place it on the scene.

Attribute Set. The attribute set s_A for each node consists of $[class, rotation, locationX, locationY, size]$. Class is a one-hot vector, spanning all possible node classes in the

graph. This includes, $[scene, background, 0, \dots, 9]$. $scene$ is the root node of the scene graph. The rotation, locations and size are floating point numbers between 0 to 1, and represent values in appropriate ranges *i.e.* rotation in 0 to 360 degrees, each location within the parent’s boundaries etc.

6.1.2 Aerial Views (2D)

Grammar. The probabilistic grammar is explained in the main paper, and we summarize it again here. First, a background grass texture is sampled. Next, a (straight) road is sampled at a location and rotation on the background. Two cars are then sampled with independent locations (constrained to be in the road by parametrizing in the roads coordinate system), and rotations. In addition, a tree and a house are sampled and placed randomly on the scene. Each object in the scene also gets assigned a random texture from a repository of collected textures.

Attribute Set. The attribute set s_A for each node in the Aerial 2D experiment consists of $[class, rotation, locationX, locationY, size]$. Our choice of class includes road, car, tree, house and background, with class as a one-hot vector. Subclasses, such as type of tree, car, etc, are not included in the attribute set. They are sampled randomly in the grammar and are left unchanged by the Distribution Transformer. The values are normalized similarly to that in the MNIST experiment. We learn to transform the $rotation$ and $location$ for every object in the scene, but leave $size$ unchanged, *i.e.* $size \notin s_{A,mut}$.

6.1.3 Driving Scenes (3D)

Grammar. The grammar used in the driving experiment is adapted from [33] (Section III). Our adaptation uses some fixed global parameters (*i.e.* weather type, sky type) as compared to SDR. Our camera optimization experiment learns the height of the camera from the ground, while scene contrast, saturation and light direction are randomized. We also do not implement side-streets and parking lanes.

Attribute Set. The attribute set s_A for each node in the Driving Scene 3D experiment consists of $[class, rotation, distance, offset]$. The classes, subclasses and ranges are treated exactly like in the Aerial 2D experiment. $Distance$ indicates how far along the parent spline the object is being placed and $offset$ indicates how much across the parent spline the object is placed. For each experiment, a subset of nodes in the graph is kept mutable. We optimized the attributes of cars, context elements such as buildings, foliage and people. Camera height from the global parameters was optimized, which was injected into the sky node encoded in the distance attribute.

6.2. Training Details

Let a_{in} be the number of attributes in s_A for each dataset. Our Distribution Transformer use Graph Convolutional lay-



Figure 16. **Successful cases:** (left) input scenes from the probabilistic grammar, (right) Meta-Sim’s generated examples for the task of car detection on KITTI. Notice how meta-sim learns to align objects in the scene, slightly change the camera position and move context elements such as buildings and trees, usually densifying the scene.

ers with different weight matrices for each direction of the edge following [47]. We use 2 layers for the encoder and 2 layers for the decoder. In the MNIST and Aerial2D experiments, it has the following structure: Encoder ($a_{in} - > 16 - > 10$), and Decoder ($10 - > 16 - > a_{in}$). For 3D driving scenes, we use 28 and 18 feature size in the intermediate layers.

After training the autoencoder step, the encoder is kept frozen and only the decoder is trained (*i.e.* for the distribution matching and task optimization steps). For the MMD computation, we always use the pool1 and pool2 feature layers of the InceptionV3 network [41]. For the MNIST and Aerial2D experiments, the images are not resized before passing them through the Inception network. For the



Figure 17. **Failure cases:** (left) input scenes from the probabilistic grammar, (right) Meta-Sim’s generated examples for the task of car detection on KITTI. Initially dense scenes are sometimes unresolved, leading to collisions in the final scenes. There are unrealistic colours on cars since they are sampled from the prior and not optimized in this work. In the last row, meta-sim moves a car very close to the ego-car (camera).

Driving Scenes (3D) experiment, they are resized to 299 x 299 (standard inception input size) before computing the features. The task network training is iterative, our Distribution Transformer is trained for one epoch, followed by training of the task network for a few epochs.

6.2.1 MNIST

For the MNIST experiment, 500 samples coming from the probabilistic grammar are said to constitute 1 epoch.

Distribution Transformer. We first train the autoencoder step of the Distribution Transformer for 8 epochs, with batch size of 8 and learning rate 0.001 using the adam optimizer [21]. Next, training is done with the joint loss, i.e. Distribution Loss and Task Loss, for 100 epochs. Note that when training with the task loss, there is only get 1 gradient step per epoch. We note that the MNIST tasks could be solved with the Distribution Loss alone, in around 3 epochs (since we can backpropagate the gradient for the distribution loss per minibatch when training only with the Distribution Loss).

Task Network. After getting each epoch of data from the Distribution Transformer, the Task Network is trained with this data for 2 epochs. The optimization was done using SGD with learning rate 0.01 and momentum 0.9. Note that the task network is not trained from scratch every time. Rather, training is resumed from the checkpoint obtained in the previous epoch.

6.2.2 Aerial Views (2D)

For the Aerial Views (2D) experiment, 100 samples coming from the probabilistic grammar are considered to be 1 epoch.

Distribution Transformer. Here, the autoencoder step is trained for 300 epochs using the adam optimizer with batch size 8 and learning rate 0.01, and then for an additional 100 epochs with learning rate 0.001. Finally, the model is trained jointly with the distribution and task loss for 100 epochs with the same batch size and a learning rate of 0.001.

Task Network. After getting each epoch of data from the Distribution Transformer, the Task Network is trained with this data for 8 epochs using SGD with a learning rate of 1e-8. Similar to the previous experiment, the task-network training is resumed instead of training from scratch.

6.2.3 Driving Scenes (3D)

For Driving Scenes (3D), 256 samples coming from the prob. grammar are considered to be 1 epoch.

Distribution Transformer. In this case, the autoencoder step is trained for 100 epochs using the adam optimizer with a batch size of 16 and learning rate 0.005. Next, the model is trained with the distribution loss for 40 epochs with the same batch size and learning rate 0.001, finally stopping at the best validation (in our case, performance on the 100 images taken from the training set) performance. Finally, the model is trained with the task loss for 8 epochs with the same learning rate.




Figure 18. Car detection results with Mask-RCNN trained on (right) the dataset generated by Meta-Sim, and (left) the original dataset obtained by the probabilistic grammar. Notice how the models trained with meta-sim have lesser false positives and negatives.

Task Network. For task network, we use Mask-RCNN [16] with a FPN and Resnet-101 backbone. We pre-train the task network (with freely available data sampled from the probabilistic grammar), saving training time by not having to train the detector from scratch. We also do not train the task network from scratch every epoch. Rather, we resume training from the checkpoint from the previous epoch. We also heavily parallelize the rendering computation by working on multiple UE4 workers simultaneously and utilizing their batched computation features.

6.3. Additional Results

6.3.1 Aerial Views (2D)

Fig. 15 shows additional qualitative results for the Aerial 2D experiment, comparing the original samples to those

obtained via Meta-Sim. The network learns to utilize the translation equivariance of convolutions and makes sufficient transformations to the objects for semantic segmentation to work on the validation data. Translating objects (such as cars) in the scene is still important, since our task-network for Aerial-2D has a maximum receptive field which is a quarter of the longest side of the image.

6.3.2 Driving Scenes (3D)

Fig 16 shows a few positive results for Meta-Sim. Our approach learns to rotate cars to look more like those in KITTI scenes and learns to place objects to avoid collisions between objects. The examples show that it can handle these even with a lot of objects present in the scene, while failing in some cases. It has also learnt to push cars out of the

view of the camera when required. Sometimes, it positions two cars perfectly on top of each other, so that only one car is visible from the camera. This is in fact a viable solution since we solve for occlusion and truncation before generating ground truth bounding boxes. We also notice that the model modifies the camera height slightly, and moves context elements, usually densifying scenes.

Fig. 17 shows failure cases for Meta-Sim. Sometimes, the model does not resolve occlusions/intersections correctly, resulting in final scenes with cars intersection each other. Sometimes it places an extra car between two cars (top row), attempts a partial merge of two cars (second row), causes collision (third row) and fails to rotate cars in the driving lane (final row). In this experiment, we did not model the size of objects in the features used in the Distribution Transformer. This could be a major reason behind this, since different assets (cars, houses) have different sizes and therefore placing them at the same location is not enough to ensure perfect occlusion between them. Meta-Sim sometimes fails to deal with situations where the probabilistic grammar samples objects densely around a single location (Fig. 17). Another failure mode is when the model moves cars unrealistically close to the ego-car (camera) and when it overlaps buildings to create unrealistic composite buildings. The right column of fig. 18 show additional detection results on KITTI when the task-network is trained with images generated by Meta-Sim. The left column shows results when trained with samples from the probabilistic grammar. In general, we see fewer false positives and negatives with Meta-sim. The results are also significant given that the task network has not seen any real KITTI images during training.

References

- [1] <https://www.unrealengine.com/>. 7
- [2] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying mmd gans. *ICLR*, 2018. 4
- [3] A. Brock, J. Donahue, and K. Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 3
- [4] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. In *arXiv:1606.01540*, 2016. 2
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *ECCV*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, 2012. 1, 2
- [6] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Isaac, N. Ratliff, and D. Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. *arXiv preprint arXiv:1810.05687*, 2018. 2
- [7] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator. In *CORL*, pages 1–16, 2017. 1, 2, 7
- [8] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *UAI*, 2015. 4
- [9] G. French, M. Mackiewicz, and M. Fisher. Self-ensembling for visual domain adaptation. In *ICLR*, 2018. 1, 2
- [10] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016. 1, 2
- [11] Y. Ganin, T. Kulkarni, I. Babuschkin, S. Eslami, and O. Vinyals. Synthesizing programs for images using reinforced adversarial learning. *arXiv preprint arXiv:1804.01118*, 2018. 2
- [12] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *CVPR*, 2012. 1, 7
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [14] E. Greensmith, P. L. Bartlett, and J. Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *JMLR*, 5(Nov):1471–1530, 2004. 5
- [15] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *JMLR*, 2012. 2, 4
- [16] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 7, 11
- [17] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. S. A. A. Efros, and T. Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 1, 2
- [18] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. 1, 2, 7
- [19] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018. 3
- [20] H. Kato, Y. Ushiku, and T. Harada. Neural 3d mesh renderer. In *CVPR*, 2018. 4
- [21] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 11
- [22] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 3
- [23] E. Kolve, R. Mottaghi, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi. Ai2-thor: An interactive 3d environment for visual ai. In *arXiv:1712.05474*, 2017. 1, 2
- [24] T. D. Kulkarni, P. Kohli, J. B. Tenenbaum, and V. Mansinghka. Picture: A probabilistic programming language for scene perception. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4390–4399, 2015. 2
- [25] Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>. 5, 9
- [26] C.-L. Li, W.-C. Chang, Y. Cheng, Y. Yang, and B. Póczos. Mmd gan: Towards deeper understanding of moment matching network. In *NIPS*, 2017. 4
- [27] P. Li, X. Liang, D. Jia, and E. P. Xing. Semantic-aware grad-gan for virtual-to-real urban scene adaption. In *BMVC*, 2018. 1, 2

- [28] T.-M. Li, M. Aittala, F. Durand, and J. Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 2018. 4
- [29] Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. In *ICML*, 2015. 4
- [30] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *NIPS*, 2017. 1, 2
- [31] G. Louppe and K. Cranmer. Adversarial variational optimization of non-differentiable simulators. *arXiv preprint arXiv:1707.07113*, 2017. 2
- [32] V. K. Mansinghka, T. D. Kulkarni, Y. N. Perov, and J. Tenenbaum. Approximate bayesian image interpretation using generative probabilistic graphics programs. In *Advances in Neural Information Processing Systems*, pages 1520–1528, 2013. 2
- [33] A. Prakash, S. Boochoon, M. Brophy, D. Acuna, E. Cameracci, G. State, O. Shapira, and S. Birchfield. Structured domain randomization: Bridging the reality gap by context-aware synthetic data. In *arXiv:1810.10093*, 2018. 1, 2, 7, 9
- [34] X. Puig, K. Ra, M. Boben, J. Li, T. Wang, S. Fidler, and A. Torralba. Virtualhome: Simulating household activities via programs. In *CVPR*, 2018. 2
- [35] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. 1, 2
- [36] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241. Springer, 2015. 6
- [37] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. Lopez. The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, 2016. 1, 2
- [38] N. Ruiz, S. Schuler, and M. Chandraker. Learning to simulate. *arXiv preprint arXiv:1810.02513*, 2018. 2
- [39] F. Sadeghi and S. Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016. 2
- [40] S. Shah, D. Dey, C. Lovett, and A. Kapoor. Aerial Informatics and Robotics platform. Technical Report MSR-TR-2017-9, Microsoft Research, 2017. 1, 2, 7
- [41] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 4, 10
- [42] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS*, 2017. 1, 2
- [43] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *Intl. Conf. on Intelligent Robots and Systems*, 2012. 2
- [44] Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018. 1
- [45] M. Wrenninge and J. Unger. Synscapes: A photo-realistic synthetic dataset for street scene parsing. In *arXiv:1810.08705*, 2018. 2, 7
- [46] Y. Wu, Y. Wu, G. Gkioxari, and Y. Tian. Building generalizable agents with a realistic and rich 3d environment. In *arXiv:1801.02209*, 2018. 1, 2
- [47] T. Yao, Y. Pan, Y. Li, and T. Mei. Exploring visual relationship for image captioning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 684–699, 2018. 3, 9
- [48] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 1, 2
- [49] S.-C. Zhu, D. Mumford, et al. A stochastic grammar of images. *Foundations and Trends® in Computer Graphics and Vision*, 2(4):259–362, 2007. 3
- [50] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017. 3
- [51] Y. Zou, Z. Yu, B. V. K. V. Kumar, and J. Wang. Domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018. 1, 2