

---

# Synthesizing Programs for Images using Reinforced Adversarial Learning

---

Yaroslav Ganin<sup>1</sup> Tejas Kulkarni<sup>2</sup> Igor Babuschkin<sup>2</sup> S. M. Ali Eslami<sup>2</sup> Oriol Vinyals<sup>2</sup>

## Abstract

Advances in deep generative networks have led to impressive results in recent years. Nevertheless, such models can often waste their capacity on the minutiae of datasets, presumably due to weak inductive biases in their decoders. This is where graphics engines may come in handy since they abstract away low-level details and represent images as high-level programs. Current methods that combine deep learning and renderers are limited by hand-crafted likelihood or distance functions, a need for large amounts of supervision, or difficulties in scaling their inference algorithms to richer datasets. To mitigate these issues, we present SPIRAL, an adversarially trained agent that generates a program which is executed by a graphics engine to interpret and sample images. The goal of this agent is to fool a discriminator network that distinguishes between real and rendered data, trained with a distributed reinforcement learning setup without any supervision. A surprising finding is that using the discriminator’s output as a reward signal is the key to allow the agent to make meaningful progress at matching the desired output rendering. To the best of our knowledge, this is the first demonstration of an end-to-end, unsupervised and adversarial inverse graphics agent on challenging real world (MNIST, OMNIGLOT, CELEBA) and synthetic 3D datasets. A video of the agent can be found at <https://youtu.be/iSyvwAwa7vk>.

## 1. Introduction

Recovering *structured* representations from raw sensations is an ability that humans readily possess and frequently use. Given a picture of a hand-written character, decomposing it into strokes can make it easier to classify or re-imagine

---

Work done while YG was an intern at DeepMind. <sup>1</sup>Montreal Institute for Learning Algorithms, Montréal, Canada <sup>2</sup>DeepMind, London, United Kingdom. Correspondence to: Yaroslav Ganin <yaroslav.ganin@gmail.com>.




Figure 1. SPIRAL takes as input either random noise or images and iteratively produces plausible samples or reconstructions via graphics program synthesis. The first row depicts an unconditional run given random noise. The second, third and fourth rows depict conditional execution given an image with a handwritten character, the Mona Lisa, and objects arranged in a 3D scene.

that character, and similarly, knowing the underlying layout of a room can aid with planning, navigation and interaction in that room. Furthermore, this structure can be exploited for generalization, rapid learning, and even communication with other agents. It is commonly believed that humans exploit *simulations* to learn this skill (Lake et al., 2017). By experimenting with a pen and a piece of paper we learn how our hand movements lead to written characters, and via imagination we learn how architectural layouts manifest themselves in reality.

In the visual domain, inversion of a renderer for the purposes of scene understanding is typically referred to as inverse graphics (Mansinghka et al., 2013; Kulkarni et al., 2015a). Training vision systems using the inverse graphics approach has remained a challenge. Renderers typically expect as input *programs* that have sequential semantics, are composed of discrete symbols (e.g., keystrokes in a CAD program), and are long (tens or hundreds of symbols). Additionally, matching rendered images to real data poses an optimization problem as black-box graphics simulators are not differentiable in general.

To address these problems, we present a new approach for interpreting and generating images using *Deep Reinforced Adversarial Learning*. In this approach, an adversarially trained agent generates visual programs which are in turn

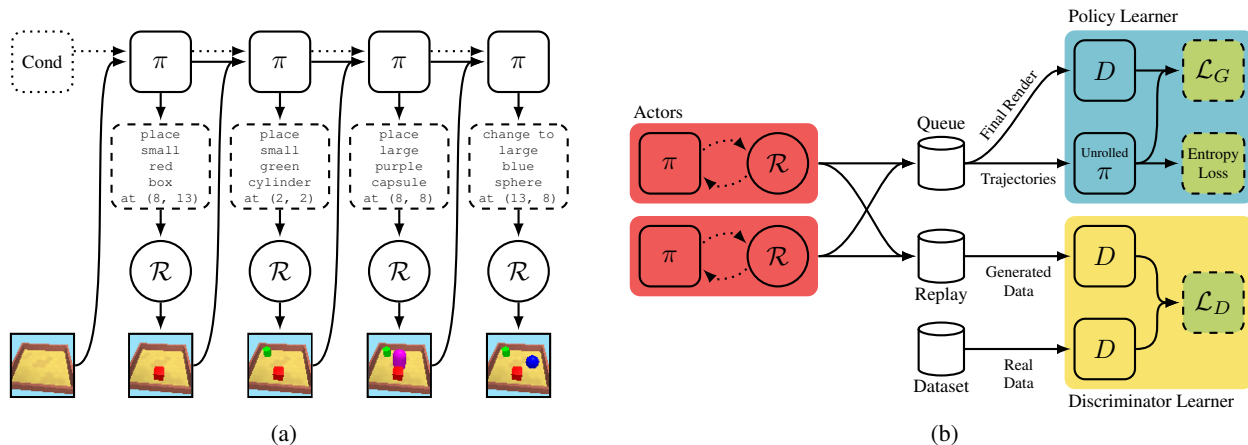


Figure 2. **The SPIRAL architecture.** (a) An execution trace of the SPIRAL agent. The policy outputs program fragments which are rendered into an image at each step via a graphics engine  $\mathcal{R}$ . The agent can make use of these intermediate renders to adjust its policy. The agent only receives a reward in the final step of execution. (b) Distributed training of SPIRAL. A collection of actors (in our experiments, up to 64), asynchronously and continuously produce execution traces. This data, along with a training dataset of ground-truth renderings, are passed to a Wasserstein discriminator on a separate GPU for adversarial training. The discriminator assesses the similarity of the final renderings of the traces to the ground-truth. A separate off-policy GPU learner receives batches of execution traces and trains the agent’s parameters via policy-gradients to maximize the reward assigned to them by the discriminator, *i.e.*, to match the distribution of the ground truth dataset.

executed by a graphics engine to generate images, either conditioned on data or unconditionally. The agent is rewarded by fooling a discriminator network, and is trained with distributed reinforcement learning without any extra supervision. The discriminator network itself is trained to distinguish between rendered and real images.

Our contributions are as follows:

- An adversarially trained reinforcement learning agent that interprets and generates images in the space of visual programs. Crucially, the architecture of our agent is agnostic both to the semantics of the visual program and to the domain.
- Scaling inverse graphics to real world and procedural datasets without the need for labels. In particular, our model discovers pen strokes that give rise to MNIST and OMNIGLOT characters, brush strokes that give rise to celebrity faces, and scene descriptions that, once rendered, reconstruct an image of a 3D scene (Figure 1).
- Evidence that utilizing a discriminator’s output as the reward signal for reinforcement learning is significantly better at optimizing the pixel error between renderings and data, compared to directly optimizing pixel error.
- A showcase of state-of-the-art deep reinforcement learning techniques, which can provide a scaling path for inverse graphics, and could lead to broader implications for program synthesis in future work.

## 2. Related Work

The idea of inverting simulators to interpret images has been explored extensively in recent years (Nair et al., 2008; Paysan et al., 2009; Mansinghka et al., 2013; Loper & Black, 2014; Kulkarni et al., 2015a; Jampani et al., 2015). Structured object-attribute based ‘de-rendering’ models have been proposed for interpretation of images (Wu et al., 2017b) and videos (Wu et al., 2017a). Concurrent work has explored the use of *Constructive Solid Geometry* primitives for explaining binary images (Sharma et al., 2017). Loper & Black (2014) proposed the idea of differentiable inverse graphics, which is efficient for optimizing continuous variables but cannot handle discrete variables. Earlier work has also explored using reinforcement learning for automatic generation of single brush strokes (Xie et al., 2013). However, scaling these approaches to larger real-world datasets, particularly at test-time, has remained a challenge.

Inferring motor programs for the reconstruction of MNIST digits was first studied in (Nair & Hinton, 2006). The generative model is parametrized by two pairs of opposing springs whose stiffness is controlled by a motor program. The training procedure involved starting with a prototype program and its corresponding observation. Random noise was then added to this prototype in order to produce new training examples until the generated distribution stretched to cover the manifold of the training digits. In contrast, our model automatically learns the training curriculum via the discriminator and the same agent is suitable for a range of scene understanding problems, including those in 3D.

Visual program induction has recently been studied in the context of hand-written characters on the OMNIGLOT dataset (Lake et al., 2015). This model achieves impressive performance but requires parses from a hand-crafted algorithm to initialize training and was not demonstrated to generalize beyond hand-written characters. Ellis et al. (2017) proposed a visual program induction model to infer L<sup>A</sup>T<sub>E</sub>X programs for diagram understanding. More recently, the sketch-rnn model (Ha & Eck, 2017) used sequence-to-sequence learning (Sutskever et al., 2014) to produce impressive sketches both unconditionally and conditioned on data. However, similar to the aforementioned works and unlike SPIRAL, the model requires supervision in the form of sketches and corresponding sequences of strokes.

In the neural network community, there have been analogous attempts at inferring and learning feed-forward or recurrent procedures for image generation (LeCun et al., 2015; Hinton & Salakhutdinov, 2006; Goodfellow et al., 2014; Ackley et al., 1987; Kingma & Welling, 2013; Oord et al., 2016; Kulkarni et al., 2015b; Eslami et al., 2016; Reed et al., 2017; Gregor et al., 2015). These models demonstrate impressive image generation capabilities but generally lack the ability to infer structured representations of images.

Our approach employs adversarial training techniques, first used for generative modeling (Goodfellow et al., 2014) and domain adaptation (Ganin & Lempitsky, 2015). Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) were originally used for image generation but have now been successfully applied to model audio, text and motor behaviors (Ho & Ermon, 2016; Merel et al., 2017). Perhaps the most interesting extension in our context is their use in domain transfer, where images from one domain (e.g., segmentations) were mapped to another (e.g., pixels). Models such as pix2pix (Isola et al., 2017), CycleGAN (Zhu et al., 2017) and AIGN (Tung et al., 2017) fall in this category.

The SPIRAL agent builds upon this literature, has minimal hand-crafting in its design, requires no supervision in the form of pairs of programs and corresponding images and, as we demonstrate in the following sections, is applicable across a wide range of domains.

### 3. The SPIRAL Agent

#### 3.1. Overview

Our goal is to construct a generative model  $G$  capable of sampling from some target data distribution  $p_d$ . To that end, we propose using an external black-box rendering simulator  $\mathcal{R}$  that accepts a sequence of commands  $a = (a_1, a_2, \dots, a_N)$  and transforms them into the domain of interest, e.g., a bitmap. For example,  $\mathcal{R}$  could be a CAD program rendering descriptions of primitives into 3D scenes.

Thus, our task is equivalent to recovering a distribution  $p_a$  such that  $p_d \approx \mathcal{R}(p_a)$ . We model  $p_a$  with a recurrent neural network  $\pi$  which we call the *policy* network (or, somewhat sloppily, the *agent*). The generation process  $G$ , which consists of a policy  $\pi$  and a renderer  $\mathcal{R}$ , is illustrated in Figure 2a.

In order to optimize  $\pi$ , we employ the adversarial framework (Goodfellow et al., 2014). In this framework, the *generator* (denoted as  $G$ ) aims to maximally confuse a *discriminator* network  $D$  which is trained to distinguish between the samples drawn from  $p_d$  and the samples generated by the model. As a result, the distribution defined by  $G$  (denoted as  $p_g$ ) gradually becomes closer to  $p_d$ . Crucially, and unlike previous work, our training procedure does not require any strong supervision in the form of aligned examples from  $p_d$  and  $p_a$ .

We give the concrete training objectives for  $G$  and  $D$  below.

#### 3.2. Objectives

In our experiments, we found that the original minimax objective from Goodfellow et al. (2014) was hard to optimize in the context of our model, so we opted for using a variant that employs Wasserstein distance as a measure of divergence between distributions (Gulrajani et al., 2017), as it appears to handle vastly dissimilar  $p_g$  and  $p_d$  more gracefully. In this case, the discriminator is considered “confused” if it assigns the same scores (in expectation) to the inputs coming both from  $p_d$  and  $p_g$ . We note that our approach can be used in conjunction with other forms of GAN objectives as well.

**Discriminator.** Following (Gulrajani et al., 2017), we define the objective for  $D$  as:

$$\mathcal{L}_D = -\mathbb{E}_{\mathbf{x} \sim p_d} [D(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [D(\mathbf{x})] + R, \quad (1)$$

where  $R$  is a regularization term softly constraining  $D$  to stay in the set of Lipschitz continuous functions (for some fixed Lipschitz constant). It is worth noting that the solution to (1) is defined up to an additive constant. Due to the nature of how we use discriminator outputs during training of  $\pi$ , we found it beneficial to resolve this ambiguity by encouraging  $D(\mathbf{x})$  to be close to 0 on average for  $\mathbf{x} \sim \frac{1}{2}p_g + \frac{1}{2}p_d$ .

**Generator.** We formally define  $\pi$  as a network that at every time step  $t$  predicts a distribution over all possible commands  $\pi_t = \pi(a_t | s_t; \theta)$ , where  $s_t$  is the recurrent state of the network, and  $\theta$  is a set of learnable parameters. Given a sequence of samples  $(a_t | a_t \sim \pi_t, 1 \leq t \leq N)$ , a sample from  $p_g$  is then computed as  $\mathcal{R}(a_1, a_2, \dots, a_N)$ . Since the generator is an arbitrary non-differentiable function, we cannot optimize

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{x} \sim p_g} [D(\mathbf{x})] \quad (2)$$

with naive gradient descent. Therefore we pose this problem as maximization of the expected return which can be solved using standard techniques from reinforcement learning. Specifically, we employ a variant of the REINFORCE (Williams, 1992) algorithm, advantage actor-critic (A2C):

$$\mathcal{L}_G = - \sum_t \log \pi(a_t | s_t; \theta) [R_t - V^\pi(s_t)], \quad (3)$$

where  $V^\pi$  is an approximation to the value function which is considered to be independent of  $\theta$ , and  $R_t = \sum_t^N r_t$  is a 1-sample Monte-Carlo estimate of the return. Optimizing (3) recovers the solution to (2) if the rewards are set to:

$$r_t = \begin{cases} 0, & t < N, \\ D(\mathcal{R}(a_1, a_2, \dots, a_N)), & t = N. \end{cases} \quad (4)$$

One interesting aspect of this new formulation is that we can also bias the search by introducing intermediate rewards which may depend not only on the output of  $\mathcal{R}$  but also on commands used to generate that output. We present several examples of such rewards in Section 4.

### 3.3. Conditional Generation

So far, we have described the case of unconditional generation, but in many situations it is useful to condition the model on auxiliary input (Mirza & Osindero, 2014). For instance, one might be interested in finding a specific program that generates a given image  $\mathbf{x}_{\text{target}}$ . That could be achieved by supplying  $\mathbf{x}_{\text{target}}$  both to the policy and to the discriminator networks. In other words,

$$p_g = \mathcal{R}(p_a(a|\mathbf{x}_{\text{target}})), \quad (5)$$

while  $p_d$  becomes a Dirac  $\delta$ -function centered at  $\mathbf{x}_{\text{target}}$ . The first two terms in (1) thus reduce to

$$-D(\mathbf{x}_{\text{target}} | \mathbf{x}_{\text{target}}) + \mathbb{E}_{\mathbf{x} \sim p_g} [D(\mathbf{x} | \mathbf{x}_{\text{target}})]. \quad (6)$$

It can be shown that for this particular setting of  $p_g$  and  $p_d$ , the  $\ell^2$ -distance is an optimal discriminator. However, in general it is not a unique solution to (1) and may be a poor candidate to be used as the generator’s reward signal (see Section A in the appendix for details). In Section 4, we empirically evaluate both  $\ell^2$  and a dynamically learned  $D$  and conclude that those two options are not equivalent in practice (for example, see Figure 8a).

### 3.4. Distributed Learning

Our training pipeline is outlined in Figure 2b. It is an extension of the recently proposed *IMPALA* architecture (Espeholt et al., 2018). For training, we define three kinds of workers:

- *Actors* are responsible for generating the training trajectories through interaction between the policy network and the rendering simulator. Each trajectory contains a sequence  $((\pi_t, a_t) | 1 \leq t \leq N)$  as well as all intermediate renderings produced by  $\mathcal{R}$ .
- A *policy learner* receives trajectories from the actors, combines them into a batch and updates  $\pi$  by performing an SGD step on  $\mathcal{L}_G$  (2). Following common practice (Mnih et al., 2016), we augment  $\mathcal{L}_G$  with an entropy penalty encouraging exploration.
- In contrast to the base *IMPALA* setup, we define an additional *discriminator learner*. This worker consumes random examples from  $p_d$ , as well as generated data (final renders) coming from the actor workers, and optimizes  $\mathcal{L}_D$  (1).

In the original paper introducing WGAN with gradient penalty (Gulrajani et al., 2017), the authors note that in order to obtain better performance, the discriminator has to be updated more frequently than the generator. In our setting, generation of each model sample is expensive since it involves multiple invocations of an external simulator. We therefore do not omit any trajectories in the policy learner. Instead, we decouple the  $D$  updates from the  $\pi$  updates by introducing a *replay buffer* that serves as a communication layer between the actors and the discriminator learner. That allows the latter to optimize  $D$  at a higher rate than the training of the policy network due to the difference in network sizes ( $\pi$  is a multi-step RNN, while  $D$  is a plain CNN). We note that even though sampling from a replay buffer inevitably leads to smoothing of  $p_g$ , we found this setup to work well in practice.

## 4. Experiments

### 4.1. Datasets

We validate our approach on three real-world and one synthetic image dataset. The first, MNIST (LeCun et al., 1998), is regarded as a standard sanity check for newly proposed generative models. It contains 70,000 examples of handwritten digits, of which 10,000 constitute a test set. Each example is a  $28 \times 28$  grayscale image. Although the dataset is often considered “solved” by neural decoder-based approaches (including GANs and VAEs), these approaches do not focus on recovering interpretable structure from the data. We, therefore, choose not to discard MNIST from the empirical evaluation since it is likely that additional constraints increase the difficulty of the modeling task.

The second dataset, OMNIGLOT (Lake et al., 2015), comprises 1623 handwritten characters from 50 alphabets. Compared to MNIST, this dataset introduces three additional challenges: higher data variability, higher complexity of




Figure 3. Illustration of the **agent’s action space** in the `libmypaint` environment. We show three different strokes (red, green, blue) that can result from a single instruction from the agent to the renderer. Starting from a position on the canvas, the agent selects the coordinates of the next end point, the coordinates of the intermediate control point, as well as the brush size, pressure and color. See Section 4.2 for details.

symbols (e.g., disjoint subcurves) and fewer (only 20) data points per symbol class.

Since both MNIST and OMNIGLOT represent a restricted line drawing domain, we diversify our set of experiments by testing the proposed method on CELEBA (Liu et al., 2015). The dataset contains over 200,000 color headshots of celebrities with large variation in poses, backgrounds and lighting conditions.

Lastly, we are interested in evaluating our approach on the task of unsupervised 3D scene understanding which is a crucial precursor for manipulating and reasoning about objects in the real world. To that end, we created a procedural dataset called MUJOCO SCENES consisting of renders of simple 3D primitives (up to 5 objects) scattered around a square platform (see Figure 10). The training set is comprised of 50,000 RGB images generated by means of the MUJOCO environment discussed in the next section.

In each case, we rescale the images to  $64 \times 64$ , which allows us to reuse the same network architectures in all the experiments, demonstrating the generality of our method.

## 4.2. Environments

We introduce two new rendering environments. For MNIST, OMNIGLOT and CELEBA generation we use an open-source painting library `libmypaint` (libmypaint contributors, 2018). The agent controls a brush and produces a sequence of (possibly disjoint) strokes on a canvas  $C$ . The state of the environment is comprised of the contents of  $C$  as well as the current brush location  $l_t$ . Each action  $a_t$  is a tuple of 8 discrete decisions ( $a_t^1, a_t^2, \dots, a_t^8$ ) (see Figure 3). The first two components are the control point  $p_c$  and the end point  $l_{t+1}$  of the stroke, which is specified as a quadratic Bézier curve:

$$p(\tau) = (1 - \tau)^2 l_t + 2(1 - \tau)\tau p_c + \tau^2 l_{t+1}, \quad (7)$$

where  $\tau \in [0, 1]$ . In our experiments, we define the valid range of locations as a  $32 \times 32$  grid imposed on  $C$ . We set  $l_0$  to the upper left corner of the canvas. The next 5 components represent the appearance of the stroke: the pressure that the agent applies to the brush (10 levels), the brush size, and the stroke color characterized by mixture of red, green and blue (20 bins for each color component). The last element of  $a_t$  is a binary flag specifying the type of action: the agent can choose either to produce a stroke or to jump right to  $l_{t+1}$ . For grayscale datasets (MNIST and OMNIGLOT), we omit the color components.


In the MUJOCO SCENES experiment, we render images using a MUJOCO-based environment (Todorov et al., 2012). At each time step, the agent has to decide on the object type (4 options), its location on a  $16 \times 16$  grid, its size (3 options) and the color (3 color components with 4 bins each). The resulting tuple is sent to the environment, which adds an object to the scene according to the specification. Additionally, the agent can decide to skip a move or change the most recently emitted object. All three types of actions are illustrated in Figure 2a.

## 4.3. MNIST

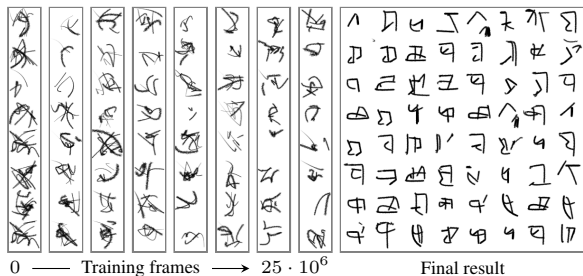
For the MNIST dataset, we conduct two sets of experiments. In the first set, we train an unconditional agent to model the data distribution. Along with the reward provided by the discriminator we also use auxiliary penalties expressing our inductive biases for the particular type of data. To encourage the agent to draw a digit in a single continuous motion of the brush, we provide a small negative reward for starting each continuous sequence of strokes. We also found it beneficial to penalize our model for not producing any visible strokes at all. The resulting agent manages to generate samples clearly recognizable as hand-written digits. Examples of such generations are shown in Figure 4a.

In the second set of experiments, we train an agent to generate the strokes for a given target digit, and we compare two kinds of rewards discussed in Section 3.3: fixed  $\ell^2$ -distance and the discriminator score. The results are summarized in Figure 8a (blue curves). We note that the discriminator-based approach significantly speeds up training of the model and achieves lower final  $\ell^2$  error. When no auxiliary rewards were employed,  $\ell^2$ -based runs failed to learn reasonable reconstructions. Figure 4b presents several conditional generations produced by our method.

Following (Sharma et al., 2017), we also train a “blind” version of the agent, i.e., we do not feed intermediate canvas states as an input to  $\pi$ . That means that the model cannot rely on reactive behaviour since it does not “see” the immediate consequences of its decisions. The training curve for this experiment is shown in Figure 8a (dotted blue line). Although the agent does not reach the level of performance



(a) MNIST unconditional generation



(a) Omniglot unconditional generation



(b) MNIST reconstruction



(b) Omniglot reconstruction

Figure 4. **MNIST.** (a) A SPIRAL agent is trained to draw MNIST digits via a sequence of strokes in the `libmypaint` environment. As training progresses, the quality of the generations increases. The final samples capture the multi-modality of the dataset, varying brush sizes and digit styles. (b) A conditional SPIRAL agent is trained to reconstruct using the same action space. Reconstructions (left) match ground-truth (right) accurately.

Figure 5. **Omniglot.** (a) A SPIRAL agent is trained to draw MNIST digits via a sequence of strokes in the `libmypaint` environment. As training progresses, the quality of the generations increase. The final samples capture the multi-modality of the dataset, varying brush sizes and character styles. (b) A conditional SPIRAL agent is trained to reconstruct using the same action space. Reconstructions (left) match ground-truth (right) accurately.

of the full model, it can still produce sensible reconstructions which suggests that our approach could be used in the more general setting of program synthesis, where access to intermediate states of the execution pipeline is not assumed.

#### 4.4. OMNIGLOT

In the previous section, we showed that our approach works reasonably well for handwritten digits. In this series of experiments, we test our agent in a similar but more challenging setting of handwritten characters. The difficulty of the dataset manifests itself in lower quality of unconditional generations (Figure 5a). Note that this task appears to be hard for other neural network based approaches as well: models that do produce good samples, such as (Rezende et al., 2016), do not do so in a manner that mimics actual strokes.

The conditional agent, on the other hand, managed to reach convincing quality of reconstructions (Figure 5b). Unfortunately, we could not make the  $\ell^2$ -based model work well in this setting (Figure 8a; dashed red line). This suggests not only that discriminator rewards speed up learning, but also that they allow successful training of agents in cases where

naïve rewards like  $\ell^2$  do not result in sufficient exploration.

Since OMNIGLOT contains a highly diverse set of symbols, over the course of training our model could learn a general notion of image reproduction rather than simply memorizing dataset-specific strokes. In order to test this, we feed a trained agent with previously unseen line drawings. The resulting reconstructions are shown in Figure 6. The agent handles out-of-domain images well, although it is slightly better at reconstructing the OMNIGLOT test set.

#### 4.5. CELEBA

Since the `libmypaint` environment is also capable of producing complex color paintings, we explore this direction by training a conditional agent on the CELEBA dataset. As in previous experiments, we use 20-step episodes, and as before, the agent does not receive any intermediate rewards. In addition to the reconstruction reward (either  $\ell^2$  or discriminator-based), we put a penalty on the earth mover’s distance between the color histograms of the model’s output and  $\mathbf{x}_{\text{target}}$ . We found this relatively task-agnostic penalty to slightly improve the performance of the method, but we would like to stress that it is by no means necessary.

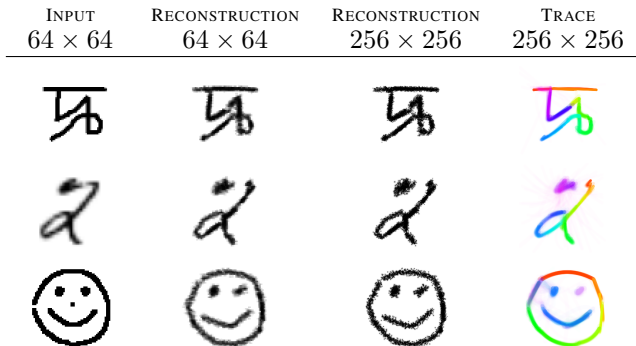


Figure 6. **Image parsing** using the SPIRAL agent trained on Omniglot. All images from test sets. Given a rastered input (a), the agent produces a reconstruction (b) that closely matches the input. (c) Having access to the underlying strokes, we can render the character at a higher resolution, or in a different stroke style. (d) The agent effectively parses the input image into strokes. Each stroke is depicted in a separate color (we show average across 100 samples).

Given that we made no effort whatsoever to adapt the action space for this domain, it is not surprising that it takes significantly more time to discover the policy that produces images resembling the target (Figure 7). As in the OMNIGLOT experiment, the  $\ell^2$ -based agent demonstrates some improvement over the random policy but gets stuck and, as a result, fails to learn sensible reconstructions (Figure 8b).

Although blurry, the model’s reconstruction closely matches the high-level structure of each image. For instance the background color, the position of the face and the color of the person’s hair. In some cases, shadows around eyes and the nose are visible. However, we observe that our model tends to generate strokes in the first half of the episode that are fully occluded by strokes in the second half. We hypothesize that this phenomenon is a consequence of credit assignment being quite challenging in this task. One possible remedy is to provide the agent with a mid-episode reward for reproducing a blurred version of the target image. We leave this prospect for future work.

#### 4.6. MUJOCO SCENES

For the MUJOCO SCENES dataset, we use our agent to construct simple CAD programs that best explain input images. Here we are only considering the case of conditional generation. Like before, the reward function for the generator can be either the  $\ell^2$  score or the discriminator output. We did not provide any auxiliary reward signals. The model is unrolled for 20 time steps, so it has the capacity to infer and represent up to 20 objects and their attributes. As we mentioned in Section 4.2, the training data consists of scenes with at most 5 objects. The agent does not have this knowledge a priori and needs to learn to place the right number of primitives.




Figure 7. **CELEBA reconstructions.** The SPIRAL agent reconstructs human faces in 20 strokes. Although blurry, the reconstructions closely match the high-level structure of each image, for instance the background color, the position of the face and the color of the person’s hair. In some cases, shadows around eyes and the nose are visible.

As shown in Figure 8b, the agent trained to directly minimize  $\ell^2$  is unable to solve the task and has significantly higher pixel-wise error. In comparison, the discriminator-based variant solves the task and produces near-perfect reconstructions on a holdout set (Figure 10).

We note that our agent has to deal with a high-cardinality action space intractable for a brute-force search. Indeed, the total number of possible execution traces is  $M^N$ , where  $M = 4 \cdot 16^2 \cdot 3 \cdot 4^3 \cdot 3$  is the total number of attribute settings for a single object (see Section 4.2 for details) and  $N = 20$  is the length of an episode.<sup>1</sup> In order to demonstrate the computational hardness of the task, we ran a general-purpose *Metropolis-Hastings* inference algorithm on a set of 100 images. The algorithm samples an execution trace defining attributes for a maximum of 20 primitives. These attributes are treated as latent variables. During each time step of inference, a block of attributes (including the presence/absence flag) corresponding to a single object is flipped uniformly within appropriate ranges. The resulting trace is rendered by the environment into an output sample which is then accepted or rejected using the Metropolis-Hastings update rule, with a Gaussian likelihood centered around the test image and a fixed diagonal covariance of 0.25. As shown in Figure 9, the MCMC search baseline was unable to solve the task even after a large number of evaluations.

## 5. Discussion

Scaling visual program synthesis to real world and combinatorial datasets has been a challenge. We have shown that it is

<sup>1</sup>The actual number of scene configurations is smaller but still intractable.

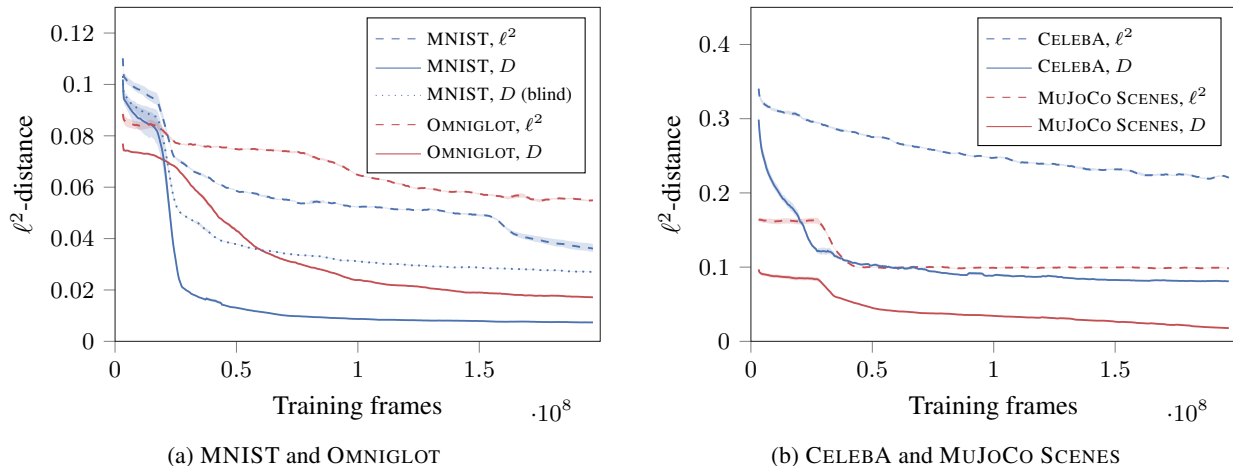


Figure 8.  $\ell^2$ -distance between reconstructions and ground truth images over the course of training. Across all datasets, we observe that training using a discriminator leads to significantly lower  $\ell^2$ -distances, than when directly minimizing  $\ell^2$ . We also show in (a) that the SPIRAL agent is capable of reconstructing even when it does not have access to the renderer in intermediate steps, however this does lead to a small degradation in performance.




Figure 9. **Blocked Metropolis-Hastings (MCMC) vs SPIRAL.** The MUJoCo SCENES dataset has a large combinatorial search space. We ran a general-purpose MCMC algorithm with object based blocked proposals and SPIRAL on 100 holdout images during inference time. SPIRAL reliably processes every image in a single pass. We ran the MCMC algorithm for thousands of evaluations but it was unable to solve the task.

possible to train an adversarial generative agent employing black-box rendering simulators. Our results indicate that using the Wasserstein discriminator’s output as a reward function with asynchronous reinforcement learning can provide a scaling path for visual program synthesis. The current exploration strategy used in the agent is entropy-based, but future work should address this limitation by employing sophisticated search algorithms for policy improvement. For instance, Monte Carlo Tree Search can be used, analogous to AlphaGo Zero (Silver et al., 2017). General-purpose inference algorithms could also be used for this purpose.

Future work should explore different parameterizations of action spaces. For instance, the use of two arbitrary control points is perhaps not the best way to represent strokes, as it is hard to deal with straight lines. Actions could also directly



Figure 10. **3D scene reconstructions.** The SPIRAL agent is trained to reconstruct 3D scenes by emitting sequences of commands for the MUJoCo environment. In each pair, the left image corresponds to the model’s output while the right one is the target. Our method is capable of accurately inferring the number of objects, their locations, sizes and colors.

parametrize 3D surfaces, planes and learned texture models to invert richer visual scenes. On the reward side, using a joint image-action discriminator similar to BiGAN/ALI (Donahue et al., 2016; Dumoulin et al., 2016) (in this case, the policy can viewed as an encoder, while the renderer becomes a decoder) could result in a more meaningful learning signal, since  $D$  will be forced to focus on the semantics of the image.

We hope that this paper provides an avenue to further explore inverse simulation and program synthesis on applications ranging from vision, graphics, speech, music and scientific simulators.



## Acknowledgements

We would like to thank Mike Chrzanowski, Michael Figurnov, David Warde-Farley, Pushmeet Kohli, Sasha Vezhnevets and Suman Ravuri for helping with the manuscript preparation as well as Sergey Bartunov, Ian Goodfellow, Jacob Menick, Lasse Espeholt, Ivo Danihelka, Junyoung Chung, Çağlar Gülçehre, Dzmityry Bahdanau and Koray Kavukcuoglu for insightful discussions and support.

## References

- Ackley, David H, Hinton, Geoffrey E, and Sejnowski, Terrence J. A learning algorithm for boltzmann machines. In *Readings in Computer Vision*. Elsevier, 1987.
- Donahue, Jeff, Krähenbühl, Philipp, and Darrell, Trevor. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Dumoulin, Vincent, Belghazi, Ishmael, Poole, Ben, Lamb, Alex, Arjovsky, Martin, Mastropietro, Olivier, and Courville, Aaron. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Ellis, Kevin, Ritchie, Daniel, Solar-Lezama, Armando, and Tenenbaum, Joshua B. Learning to infer graphics programs from hand-drawn images. *arXiv preprint arXiv:1707.09627*, 2017.
- Eslami, SM Ali, Heess, Nicolas, Weber, Theophane, Tassa, Yuval, Szepesvari, David, Kavukcuoglu, Koray, and Hinton, Geoffrey E. Attend, infer, repeat: Fast scene understanding with generative models. In *NIPS*, 2016.
- Espeholt, Lasse, Soyer, Hubert, Munos, Remi, Simonyan, Karen, Mnih, Volodymir, Ward, Tom, Doron, Yotam, Firoiu, Vlad, Harley, Tim, Dunning, Iain, Legg, Shane, and Kavukcuoglu, Koray. IMPALA: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Ganin, Yaroslav and Lempitsky, Victor. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *NIPS*, 2014.
- Gregor, Karol, Danihelka, Ivo, Graves, Alex, Rezende, Danilo, and Wierstra, Daan. Draw: A recurrent neural network for image generation. In *ICML*, 2015.
- Gulrajani, Ishaan, Ahmed, Faruk, Arjovsky, Martin, Dumoulin, Vincent, and Courville, Aaron C. Improved training of wasserstein gans. In *NIPS*, 2017.
- Ha, David and Eck, Douglas. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*, 2017.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *CVPR*, 2016.
- Hinton, Geoffrey E and Salakhutdinov, Ruslan R. Reducing the dimensionality of data with neural networks. *Science*, 2006.
- Ho, Jonathan and Ermon, Stefano. Generative adversarial imitation learning. In *NIPS*, 2016.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 1997.
- Isola, Phillip, Zhu, Jun-Yan, Zhou, Tinghui, and Efros, Alexei A. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- Jaderberg, Max, Dalibard, Valentin, Osindero, Simon, Czarnecki, Wojciech M., Donahue, Jeff, Razavi, Ali, Vinyals, Oriol, Green, Tim, Dunning, Iain, Simonyan, Karen, Fernando, Chrisantha, and Kavukcuoglu, Koray. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.
- Jampani, Varun, Nowozin, Sebastian, Loper, Matthew, and Gehler, Peter V. The informed sampler: A discriminative approach to bayesian inference in generative computer vision models. *CVIU*, 2015.
- Kingma, Diederik P and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, Diederik P and Welling, Max. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kulkarni, Tejas D, Kohli, Pushmeet, Tenenbaum, Joshua B, and Mansinghka, Vikash. Picture: A probabilistic programming language for scene perception. In *CVPR*, 2015a.
- Kulkarni, Tejas D, Whitney, William F, Kohli, Pushmeet, and Tenenbaum, Josh. Deep convolutional inverse graphics network. In *NIPS*, 2015b.
- Lake, Brenden M, Salakhutdinov, Ruslan, and Tenenbaum, Joshua B. Human-level concept learning through probabilistic program induction. *Science*, 2015.
- Lake, Brenden M, Ullman, Tomer D, Tenenbaum, Joshua B, and Gershman, Samuel J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.

- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 1998.
- LeCun, Yann, Bengio, Yoshua, and Hinton, Geoffrey. Deep learning. *Nature*, 2015.
- libmypaint contributors. libmypaint. <https://github.com/mypaint/libmypaint>, 2018.
- Liu, Ziwei, Luo, Ping, Wang, Xiaogang, and Tang, Xiaoou. Deep learning face attributes in the wild. In *ICCV*, 2015.
- Loper, Matthew M and Black, Michael J. Opendr: An approximate differentiable renderer. In *ECCV*, 2014.
- Mansinghka, Vikash K, Kulkarni, Tejas D, Perov, Yura N, and Tenenbaum, Josh. Approximate bayesian image interpretation using generative probabilistic graphics programs. In *NIPS*, 2013.
- Merel, Josh, Tassa, Yuval, Srinivasan, Sriram, Lemmon, Jay, Wang, Ziyu, Wayne, Greg, and Heess, Nicolas. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.
- Mirza, Mehdi and Osindero, Simon. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Mnih, Volodymyr, Badia, Adria Puigdomenech, Mirza, Mehdi, Graves, Alex, Lillicrap, Timothy, Harley, Tim, Silver, David, and Kavukcuoglu, Koray. Asynchronous methods for deep reinforcement learning. In *ICML*, 2016.
- Nair, Vinod and Hinton, Geoffrey E. Inferring motor programs from images of handwritten digits. In *NIPS*, 2006.
- Nair, Vinod, Susskind, Josh, and Hinton, Geoffrey E. Analysis-by-synthesis by learning to invert generative black boxes. In *ICANN*, 2008.
- Oord, Aaron van den, Kalchbrenner, Nal, and Kavukcuoglu, Koray. Pixel recurrent neural networks. In *ICML*, 2016.
- Paysan, Pascal, Knothe, Reinhard, Amberg, Brian, Romdhani, Sami, and Vetter, Thomas. A 3d face model for pose and illumination invariant face recognition. In *AVSS*, 2009.
- Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Reed, Scott, Oord, Aaron, Kalchbrenner, Nal, Colmenarejo, Sergio Gómez, Wang, Ziyu, Chen, Yutian, Belov, Dan, and Freitas, Nando. Parallel multiscale autoregressive density estimation. In *ICML*, 2017.
- Rezende, Danilo, Danihelka, Ivo, Gregor, Karol, Wierstra, Daan, et al. One-shot generalization in deep generative models. In *ICML*, 2016.
- Sharma, Gopal, Goyal, Rishabh, Liu, Difan, Kalogerakis, Evangelos, and Maji, Subhansu. CSGNet: Neural shape parser for constructive solid geometry. *arXiv preprint arXiv:1712.08290*, 2017.
- Silver, David, Schrittwieser, Julian, Simonyan, Karen, Antonoglou, Ioannis, Huang, Aja, Guez, Arthur, Hubert, Thomas, Baker, Lucas, Lai, Matthew, Bolton, Adrian, et al. Mastering the game of go without human knowledge. *Nature*, 2017.
- Sutskever, Ilya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
- Todorov, Emanuel, Erez, Tom, and Tassa, Yuval. Mujoco: A physics engine for model-based control. In *IROS*, 2012.
- Tung, Hsiao-Yu Fish, Harley, Adam W, Seto, William, and Fragkiadaki, Katerina. Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation from unpaired supervision. In *ICCV*, 2017.
- Villani, Cédric. *Optimal transport: old and new*. Springer Science & Business Media, 2008.
- Williams, Ronald J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*. Springer, 1992.
- Wu, Jiajun, Lu, Erika, Kohli, Pushmeet, Freeman, Bill, and Tenenbaum, Josh. Learning to see physics via visual de-animation. In *NIPS*, 2017a.
- Wu, Jiajun, Tenenbaum, Joshua B, and Kohli, Pushmeet. Neural scene de-rendering. In *CVPR*, 2017b.
- Xie, Ning, Hachiya, Hirotaka, and Sugiyama, Masashi. Artist agent: A reinforcement learning approach to automatic stroke generation in oriental ink painting. *IEICE Transactions on Information and Systems*, 2013.
- Zhu, Jun-Yan, Park, Taesung, Isola, Phillip, and Efros, Alexei A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *CVPR*, 2017.

## A. Optimal $D$ for Conditional Generation

It turns out that in the case of conditional generation (*i.e.*,  $p_d$  is a Dirac  $\delta$ -function), we can derive an explicit form of the optimal (non-parametric) discriminator  $D$ . Indeed, (1) corresponds to the dual representation of the Wasserstein-1 metric (Villani, 2008). The primal form of that metric is defined as

$$W_1(p_g, p_d) = \inf_{\gamma \in \Gamma(p_g, p_d)} \int \|\mathbf{x} - \mathbf{y}\|_2 d\gamma(\mathbf{x}, \mathbf{y}), \quad (8)$$

where  $\Gamma(p_g, p_d)$  is a set of all couplings between  $p_g$  and  $p_d$ . Taking into account that the data distribution is a point mass, we can simplify (8):

$$W_1(p_g, p_d) = \mathbb{E}_{\mathbf{x} \sim p_g} \|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2. \quad (9)$$

The expression above gives the optimal value for  $\mathcal{L}_D$  in (1). Therefore  $D(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2$  is a solution of (1):

$$\begin{aligned} \mathcal{L}_D(\|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2) = & - \|\mathbf{x}_{\text{target}} - \mathbf{x}_{\text{target}}\|_2 \\ & + \mathbb{E}_{\mathbf{x} \sim p_g} \|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2 \\ & + 0, \end{aligned} \quad (10)$$

where the last term ( $R$ ) is zero since the Euclidean distance belongs to the set of 1-Lipschitz functions.

This result suggests that for inverse graphics, in (4), one may use a fixed image distance (like the Euclidean distance  $\ell^2$ ) instead of a parametric function optimized via the WGAN objective. Note, however, that  $\ell^2$  is not a unique solution to (1). Consider, for example, the case where the model distribution is also a Dirac delta centered at  $\mathbf{x}_g$ . The Wasserstein distance is equal to  $d = \|\mathbf{x}_g - \mathbf{x}_{\text{target}}\|_2$ . In order to achieve that value in (1), we could take any  $D$  such that  $\forall \alpha \in [0, 1]$

$$D(\alpha \mathbf{x}_g + (1 - \alpha) \mathbf{x}_{\text{target}}) = \alpha \cdot d, \quad (11)$$

and  $\text{Lip}(D) \leq 1$ . One example of such function would be a hyperplane  $H$  containing the segment (11). Let us now consider a set of points

$$V = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_g\|_2 < \varepsilon, \|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2 = d\} \quad (12)$$

in an  $\varepsilon$ -vicinity of  $\mathbf{x}_g$ . By definition,  $\|\mathbf{x} - \mathbf{x}_{\text{target}}\|_2$  is constant for any  $\mathbf{x} \in V$ . That means that  $\ell^2$  expresses no preference over points that are equidistant from  $\mathbf{x}_{\text{target}}$  even though some of them may be semantically closer to  $\mathbf{x}_{\text{target}}$ . This property may significantly slow down learning if we are relying on  $\ell^2$  (or similar distance) as our training signal. Functions like  $H$ , on the other hand, have non-zero slope in  $V$  and therefore can potentially shift the search towards more promising subspaces.

One other reason why discriminator training is different from using a fixed image distance is that in practice, we do

not optimize the exact dual formulation of the Wasserstein distance and, on top of that, use stochastic gradient descent methods which we do not run until convergence. A toy example illustrating that difference is presented in Figure 11.

## B. Network Architectures

The policy network (shown in Figure 12) takes the observation (*i.e.*, the current state of the canvas  $C_t$ ) and conditions it on a tuple corresponding to the last performed action  $a_t$ . The resulting features are then downsampled to a lower-dimensional spatial resolution by means of strided convolutions and passed through a stack of ResNet blocks (He et al., 2016) followed by a fully-connected layer. This yields an embedding which we feed into an LSTM (Hochreiter & Schmidhuber, 1997). The LSTM produces a hidden vector  $z_0$  serving as a seed for the action sampling procedure described below.

In order to obtain  $a_{t+1}$ , we employ an *autoregressive decoder* depicted in Figure 13. Each component  $a_{t+1}^i$  is sampled from a categorical distribution whose parameters are computed as a function of  $z_i$ . We use two kinds of functions depending on whether  $a_{t+1}^i$  corresponds to a scalar (*e.g.*, brush size) or to a spatial location (*e.g.*, a control point of a Bézier curve). In the scalar case,  $z_i$  is transformed by a fully-connected layer, otherwise we process it using several ResNet blocks followed by a series of transpose convolutions and a final convolution. After  $a_{t+1}^i$  is sampled, we obtain an updated hidden vector  $z_{i+1}$  by embedding  $a_{t+1}^i$  into a 16-dimensional code and combining it with  $z_i$ . The procedure is repeated until the entire action tuple has been generated.

For the discriminator network, we use a conventional architecture similar to DCGAN (Radford et al., 2015).

## C. Training Details

Following standard practice in the GAN literature, we optimize the discriminator objective using Adam (Kingma & Ba, 2014) with a learning rate of  $10^{-4}$  and  $\beta_1$  set to 0.5. For generator training, we employ population-based exploration of hyperparameters (PBT) (Jaderberg et al., 2017) to find values for the entropy loss coefficient and learning rate of the policy learner. A population contains 12 training instances with each instance running 64 CPU actor jobs and 2 GPU jobs (1 for the policy learner and 1 for the discriminator learner). We assume that discriminator scores are compatible across different instances and use them as a measure of fitness in the exploitation phase of PBT.

The batch size is set to 64 on both the policy learner and discriminator learner. The generated data is sampled uniformly from a replay buffer with a capacity of 20 batches.

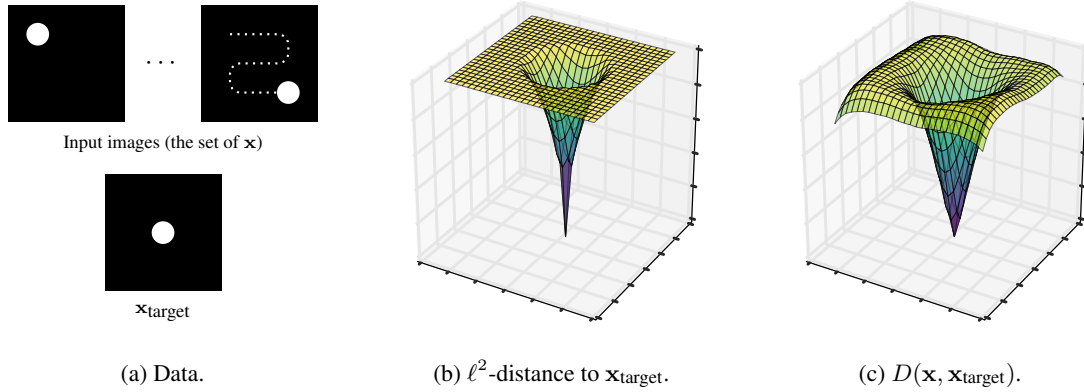


Figure 11. A toy experiment illustrating the **difference between  $\ell^2$  and discriminator training in practice**. (a) We collect a dataset of images with a single solid circle in all possible locations (top) and pick one of them as a target image (bottom). (b) The  $\ell_2$ -distance (in the pixel space) from the input images to the target as a function of the circle location; the surface is flat around the borders since the circles do not overlap. (c) We train a discriminative model  $D$  that takes a pair of images and tells whether they match or not; just like  $\ell^2$ , the resulting function has a pit centered at the target location, but unlike (b), the surface now has better behaved gradients.

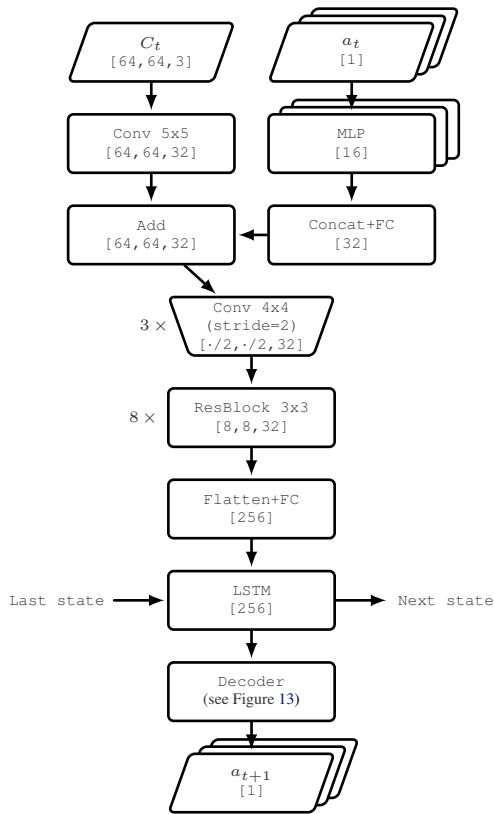


Figure 12. The architecture of the **policy network** for a single step. FC refers to a fully-connected layer, MLP is a multilayer perceptron, Conv is a convolutional layer and ResBlock is a residual block. We give the dimensions of the output tensors in the square brackets. ReLU activations between the layers have been omitted for brevity.

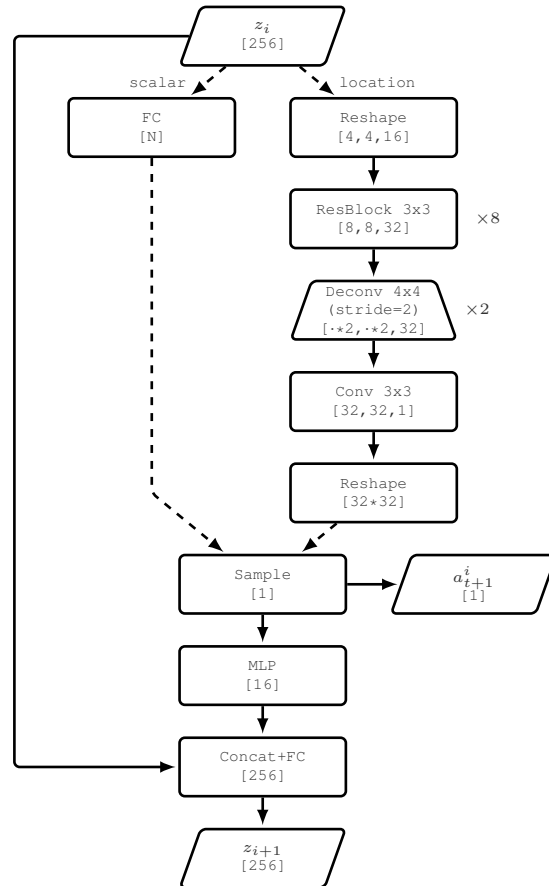


Figure 13. The architecture of the **autoregressive decoder** for sampling an element  $a_{t+1}^i$  of the action tuple. The initial hidden vector  $z_0$  is provided by an upstream LSTM. Depending on the type of the subaction to be sampled, we use either the `scalar` or the `location` branch of the diagram.