

Exploiting Spatial Invariance for Scalable Unsupervised Object Tracking

Eric Crawford

Mila/McGill University
Montreal, QC, Canada
eric.crawford@mail.mcgill.ca

Joelle Pineau

Mila/McGill University
Montreal, QC, Canada
jpineau@cs.mcgill.ca

Abstract

The ability to detect and track objects in the visual world is a crucial skill for any intelligent agent, as it is a necessary precursor to any object-level reasoning process. Moreover, it is important that agents learn to track objects without supervision (i.e. without access to annotated training videos) since this will allow agents to begin operating in new environments with minimal human assistance. The task of learning to discover and track objects in videos, which we call *unsupervised object tracking*, has grown in prominence in recent years; however, most architectures that address it still struggle to deal with large scenes containing many objects. In the current work, we propose an architecture that scales well to the large-scene, many-object setting by employing spatially invariant computations (convolutions and spatial attention) and representations (a spatially local object specification scheme). In a series of experiments, we demonstrate a number of attractive features of our architecture; most notably, that it outperforms competing methods at tracking objects in cluttered scenes with many objects, and that it can generalize well to videos that are larger and/or contain more objects than videos encountered during training.

1 Introduction

The ability to reason about objects is a vital skill for intelligent agents. Indeed, it underlies much of human intelligence, and is one of several “core” domains of human cognition, meaning that it is sufficiently important that evolution has endowed humans (and many animals) with special-purpose hardware for discovering, tracking and reasoning about objects (Carey 2009). Recently, significant progress has been made in the design of neural networks that can reason about objects, and these architectures have been shown to possess a number of significant advantages over systems that lack object-like representations (Diuk, Cohen, and Littman 2008; Chang et al. 2016; Kansky et al. 2017; Zambaldi et al. 2018). However, in order to reason in terms of objects, a system first needs a way of discovering and tracking objects in the world; moreover, the system should be able to do this without supervision, in order to be able to adapt to new environments with little human oversight. Much of the recent work on object-level reasoning assumes that object-like representations are

directly provided by the environment, thereby avoiding the problem of discovering and learning to track objects.

In the current work, we aim to make progress on this task of discovering and tracking objects, which we call *unsupervised object tracking*. Our general approach is to formulate a Variational Autoencoder (VAE) (Kingma and Welling 2013) for videos, endowed with a highly structured, object-like latent representation. This VAE is composed of a number of modules that are applied each timestep of the input video, notably a discovery module which detects new objects in the current frame, and a propagation module which updates the attributes of objects discovered in previous frames based on information from the current frame. For each input frame, a rendering module creates a corresponding output frame from the objects proposed by the discovery and propagation modules. The network is trained by maximizing the VAE evidence lower bound, which encourages the output frames to be accurate reconstructions of the input frames. At the end of training, it is expected that the discovery module will have become a competent object detector, while the propagation module will have learned object dynamics. This overall architecture was first proposed in Sequential Attend Infer Repeat (SQAIR) (Kosiorsek et al. 2018).

However, as we demonstrate empirically, SQAIR struggles at processing spatially large videos that contain many densely packed objects. We hypothesize that this is because SQAIR does not fully exploit the spatial statistics of objects. One example is SQAIR’s discovery module, which processes input frames holistically. Specifically, it computes an initial representation for an input frame by processing it with a Multi-layer Perceptron (MLP), effectively discarding all spatial structure, before recurrently emitting a sequence of detected objects. The discovery module is consequently unable to exploit the fact that objects are *spatially local*. A consequence of this spatial locality is that in order to detect an object, a local detector can be used which only has access to a sub-region of the frame. Doing so has the notable advantage that a sub-region will have significantly less complexity than the frame as a whole, making the detector’s job easier. In order to ensure the entire frame is covered, the local detector can be applied repeatedly to different (but possibly overlapping) sub-regions. This scheme can be implemented efficiently through the use of convolutional neural networks; indeed, this is the intuition behind the success of recent su-

ervised single-shot object detectors such as SSD (Liu et al. 2016) and YOLO (Redmon and Farhadi 2017).

In the current work, we propose Spatially Invariant Label-free Object Tracking (SILOT) (pronounced like “silo”), a differentiable architecture for unsupervised object tracking that is able to scale well to large scenes containing many objects. SILOT achieves this scalability by making extensive use of spatially invariant computations and representations, thereby fully exploiting the structure of objects in images. For example, in its discovery module SILOT employs a convolutional object detector with a spatially local object specification scheme, and spatial attention is used throughout to compute objects features in a spatially invariant manner.

Through a number of experiments, we demonstrate the concrete advantages that arise from this focus on spatial invariance. In particular, we show that SILOT has a greatly improved capacity for handling large, many-object videos, and that trained SILOT networks can generalize well to videos that are larger and/or contain different numbers of objects than videos encountered in training.

2 Related Work

Our work builds on a growing body of recent research on learning to detect and track objects without supervision. The pioneering work Attend, Infer, Repeat (AIR) (Eslami et al. 2016) formulated a VAE for images, with a highly structured, object-like latent representation. By training this VAE to reconstruct input images, the encoder module of the VAE is forced to learn to detect objects. AIR has since been extended in a number of directions. Sequential Attend, Infer, Repeat (SQAIR) (Kosiorek et al. 2018) extended AIR to handle videos rather than images. This was achieved by combining an object discovery module, which extracts objects from the scene and has a design similar to AIR, with an object propagation module which tracks changes in object attributes (e.g. position, size, appearance) over time. Similar extensions are also provided by (Hsieh et al. 2018; He et al. 2018). In an orthogonal direction, Spatially Invariant Attend, Infer, Repeat (SPAIR) (Crawford and Pineau 2019) improved on AIR’s ability to handle cluttered scenes by replacing AIR’s recurrent encoder network with a convolutional network and a spatially local object specification scheme. Finally, Discrete-AIR (Wang, Jamnik, and Lio 2019) showed how to adapt AIR to discover discrete object classes.

AIR and its descendants generally describe each object using, at minimum, a bounding box and the appearance of the object within that box. A separate but related body of work uses an alternate scheme, specifying each object using two image-sized maps: a mask map, which determines how much the object “owns” each image pixel, and an appearance map, which specifies the object’s appearance at pixels that it owns. In these architectures, objects compete with one another for ownership of image pixels through their mask maps. Examples include (Greff, van Steenkiste, and Schmidhuber 2017; van Steenkiste et al. 2018; Goel, Weng, and Poupart 2018; Greff et al. 2019; Burgess et al. 2019). One weakness of this style of object specification is that the objects are relatively heavy-weight, which makes it computationally intensive to use large numbers of objects. Additionally, there is no way to

ensure that the objects are compact and object-like (in AIR-style architectures, a prior is used to encourage discovered objects to have small spatial extent). Together these weaknesses make these architectures poorly suited to handling large scenes with many objects. For instance, when the architecture proposed in (Goel, Weng, and Poupart 2018) was applied to scenes from the Space Invaders Atari game, it grouped the 36 aliens into a single object, rather than allocating a separate object per alien.

As a final note, an architecture similar to SILOT, with a similar emphasis on ability to scale to large numbers of objects, was concurrently developed in (Jiang et al. 2019).

3 Spatially Invariant, Label-free Object Tracking

SILOT is a Variational Autoencoder (VAE) which models a video as a collection of moving objects. The model is divided into modules: a discovery module, which detects objects from each frame; a propagation module, which updates the attributes of previously discovered objects; a selection module, which selects a small set of objects to keep from the union of the discovered and propagated objects; and a rendering module which renders selected objects into an output frame. The discovery, propagation and selection modules constitute the VAE encoder $h_\phi(z|x)$, with all encoder parameters gathered together in a vector ϕ . Meanwhile the rendering module constitutes the VAE decoder $g_\theta(x|z)$, with its parameters collected in a vector θ .

3.1 Overview

In this section we give a high-level overview of the architecture of SILOT. Throughout this work we generally surround temporal indices with brackets to differentiate them from other types of indices (e.g. indexing a specific object within a set of objects). To reduce clutter we often drop the temporal index when it can be inferred from the context.

Assume we are given a length- T input video $\{x_{(t)}\}_{t=0}^{T-1}$, with $x_{(t)}$ denoting an individual frame for $t \in \{0, \dots, T-1\}$. For each timestep, we consider a number of variable sets:

1. Discovered latents $\bar{z}_{(t)}$ and discovered objects $\bar{o}_{(t)}$
2. Propagated latents $\tilde{z}_{(t)}$ and propagated objects $\tilde{o}_{(t)}$
3. Selected objects $o_{(t)}$

The total set of latent variables for the VAE is:

$$z = \bigcup_{t=0}^{T-1} \bar{z}_{(t)} \cup \tilde{z}_{(t)}$$

For both discovery and propagation, the object sets are deterministic functions of the corresponding latent variable sets (and possibly other variables). The choice to employ a separate set of latents, apart from the objects, was made in part because it is generally inconvenient to define prior distributions (required for training the VAE) in the highly structured, interpretable space of objects. Past architectures have employed a similar approach (Eslami et al. 2016; Kosiorek et al. 2018).

The generative model assumed by SILOT, as well as the high-level structure of the SILOT neural network, is shown in Figure 1. There we see the relationships between the variables and modules discussed previously. Both the generative model and high-level network structure are similar to and inspired by SQAIR (Kosiorek et al. 2018). Within a timestep t , the flow of computation proceeds as follows:

1. Input frame $x_{(t)}$ and objects from the previous frame $o_{(t-1)}$ are passed into the propagation module, which predicts a set of updates $\tilde{z}_{(t)}$ and deterministically applies them, yielding propagated objects $\tilde{o}_{(t)}$.
2. $x_{(t)}$ and $\tilde{o}_{(t)}$ are passed into the discovery module, which discovers objects in the frame that are not accounted for by any propagated object, first yielding $\tilde{z}_{(t)}$ and then $\tilde{o}_{(t)}$ via a deterministic transformation.
3. $\tilde{o}_{(t)}$ and $\tilde{o}_{(t)}$ are passed into the selection module, which selects a subset of the objects to retain, yielding $o_{(t)}$.
4. $o_{(t)}$ is passed into the rendering module which yields an output frame $\hat{x}_{(t)}$.

On the initial timestep, the propagation module is not executed and a degenerate set of objects is used for $\tilde{o}_{(0)}$.

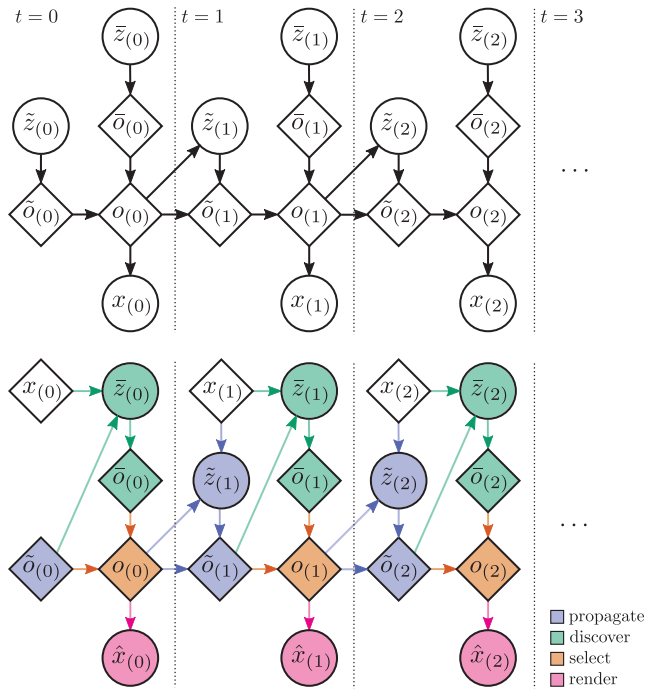


Figure 1: High-level overview of SILOT. Top: Generative model assumed by SILOT. Diamonds/circles are deterministic/stochastic functions of their inputs. Bottom: Structure of the SILOT neural network. Modules are indicated by color. Propagation, discovery and selection constitute the encoder/inference network $h_{\phi}(z|x)$, while rendering constitutes the decoder/likelihood network $g_{\theta}(x|z)$.

3.2 Object Representation

In SILOT, an individual object is represented by a collection of variables, each called an *attribute*:

$$o^{\text{where}} \in \mathbb{R}^4 \quad o^{\text{what}} \in \mathbb{R}^A \quad o^{\text{depth}} \in [0, 1] \quad o^{\text{pres}} \in [0, 1]$$

o^{where} decomposes as $o^{\text{where}} = (o^y, o^x, o^h, o^w)$. o^y and o^x specify the location of the object’s center, while o^h and o^w specify the object’s size. o^{what} acts as a catch-all, storing information about the object that is not captured by other attributes (e.g. appearance, velocity). o^{depth} specifies the relative depth of the object; in the output image, objects with higher values for this attribute appear on top of objects with lower values. o^{pres} specifies the extent to which the object exists; objects with $o^{\text{pres}} = 0$ do not appear in the output image.

3.3 Discovery

The role of the object discovery module is to take in the current frame $x_{(t)}$ and the set of objects propagated from the previous frame $\tilde{o}_{(t-1)}$, and detect any objects in the frame that are not yet accounted for (either because those objects appeared for the first time in the current frame, or because they were not previously discovered due to error). Object discovery in SILOT is implemented as a convolutional neural network in order to achieve spatial invariance, and is heavily inspired by SPAIR, an unsupervised convolutional object detector (Crawford and Pineau 2019), as well as by single-shot supervised object detectors such as SSD and YOLO.

An initial convolutional network d_{ϕ}^{bu} extracts “bottom-up” information from the current input frame $x_{(t)}$, mapping to a feature volume $v_{(t)}^{\text{bu}}$:

$$v_{(t)}^{\text{bu}} = d_{\phi}^{\text{bu}}(x_{(t)})$$

The structure of network d_{ϕ}^{bu} can be taken to induce a spatial grid over the frame, as follows. Let c_h/c_w be the translation (in pixels) vertically/horizontally between receptive fields of adjacent spatial locations in $v_{(t)}^{\text{bu}}$. Then for an input frame with dimensions $(H_{\text{inp}}, W_{\text{inp}}, 3)$, we divide the frame up into an (H, W) grid of cells, each cell being c_h pixels high by c_w pixels wide, where $H = \lceil H_{\text{inp}}/c_h \rceil$, $W = \lceil W_{\text{inp}}/c_w \rceil$. The output volume $v_{(t)}^{\text{bu}}$ has spatial shape (H, W) , and we associate each of its spatial locations with the corresponding grid cell. Importantly, the input frame is padded on all sides to ensure that the receptive field for each spatial location in $v_{(t)}^{\text{bu}}$ is centered on the corresponding grid cell¹. This scheme is visualized in Figure 2.

The discovery module will ultimately yield a separate object for each grid cell. It is useful to think of the discovery module as consisting of a 2D array of identical object detectors, each operating on a different local region of the image. We call each of these local detectors a *discovery unit*; the structure of a unit is shown in Figure 3.

Variables for all units are grouped together into convolutional volumes, and computations are implemented by size-preserving convolutions on these volumes, essentially computing all discovery units in parallel. Thus, unless otherwise

¹(Dang Ha 2017) is a useful guide to receptive field arithmetic in convolutional nets.

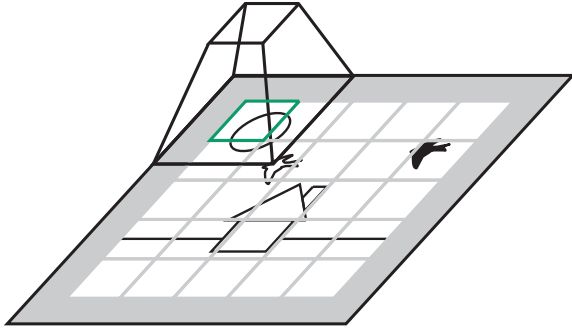


Figure 2: Schematic depicting grid cells and padding in the discovery module. The gray grid is the grid of cells described in Section 3.3. The top of the trapezoid is a spatial location in the output layer of d_ϕ^{bu} , associated with the grid cell highlighted in green. The bottom of the trapezoid is the receptive field of that spatial location; the frame is padded (solid gray border area) before being passed into d_ϕ^{bu} to ensure that all receptive fields are centered on their corresponding grid cells.




Figure 3: Schematic depicting the structure of a discovery unit with indices ij at time t , discovering the black bird which has just come into view from the right. Local bottom-up information from the current frame is processed by a convolutional filter (trapezoid), which has a receptive field (grey base of the trapezoid) centered on the discovery unit’s grid cell (green rectangle). Next, top-down information about nearby objects propagated from the previous frame (orange boxes) is summarized using spatial attention with a Gaussian kernel centered at the grid cell. Bottom-up and top-down information is then fused and used to autoregressively predict object attributes (here we have omitted the latent discovery variables \bar{z}_{ij}). $\bar{o}_{ij}^{\text{where}}$ is specified with respect to the grid cell, and the prediction for $\bar{o}_{ij}^{\text{what}}$ conditions on the output of a spatial transformer parameterized by $\bar{o}_{ij}^{\text{where}}$ (blue rectangle).

specified, all variables in this section are volumes with spatial dimensions (H, W) , and all networks are convolutional and preserve spatial dimensions. The convolutional layers use a kernel width and stride of 1 so that each discovery unit is independent (such a layer is equivalent to applying a fully

connected layer separately to each spatial location).

In order to avoid rediscovering objects that are already accounted for, each discovery unit needs to be aware of “top-down” information about objects propagated from the previous frame that are near its grid cell. We thus employ a spatial attention step which, for each grid cell, extracts features of propagated objects, and then weights those features according to a Gaussian kernel centered at the cell (similar to the Neural Physics Engine (Chang et al. 2016)). The result is a feature volume $v_{(t)}^{\text{td}}$ with spatial shape (H, W) where each spatial location contains information about nearby propagated objects:

$$v_{(t)}^{\text{td}} = \text{SpatialAttention}_\phi^{\text{disc}}(\bar{o}_{(t)}, \sigma)$$

where σ is a hyperparameter giving the standard deviation of the Gaussian kernel. Details on the spatial attention step are provided in Section B of the Supplementary Material².

Next, a convolutional network d_ϕ^{fuse} combines bottom-up and top-down information (here we begin omitting temporal indices):

$$v = d_\phi^{\text{fuse}}(v^{\text{bu}}, v^{\text{td}})$$

We follow the convention that a neural network that takes multiple inputs first concatenates those inputs along their trailing dimension (i.e. the depth dimension in the case of convolutional volumes), and processes the result as usual.

The network then predicts parameters for distributions over the latent discovery variables \bar{z} , samples from the predicted distributions, and then maps the sampled latent to the more interpretable \bar{o} . This is done on an attribute-by-attribute basis (in order [where, what, depth, pres]), and is autoregressive, so that predictions for later attributes are conditioned on samples for earlier attributes. These steps are detailed below.

Predicting \bar{o}^{where} . We first use a network d_ϕ^{where} to predict parameters for a distribution over \bar{z}^{where} , and then sample:

$$\begin{aligned} \bar{\mu}^{\text{where}}, \bar{\sigma}^{\text{where}} &= d_\phi^{\text{where}}(v) \\ \bar{z}^{\text{where}} &\sim N(\bar{\mu}^{\text{where}}, \bar{\sigma}^{\text{where}}) \end{aligned}$$

Next we deterministically map to \bar{o}^{where} . We decompose \bar{z}^{where} as $\bar{z}^{\text{where}} = (\bar{z}^y, \bar{z}^x, \bar{z}^h, \bar{z}^w)$. Here it will be useful to narrow our focus to a single discovery unit with indices ij for $i \in \{0, \dots, H-1\}, j \in \{0, \dots, W-1\}$. \bar{z}_{ij}^y and \bar{z}_{ij}^x parameterize the position of the object according to:

$$\begin{aligned} b_{ij}^y &= b^{\min} + \text{sigmoid}(\bar{z}_{ij}^y) (b^{\max} - b^{\min}) \\ \bar{o}_{ij}^y &= (i + b_{ij}^y)c_h \\ b_{ij}^x &= b^{\min} + \text{sigmoid}(\bar{z}_{ij}^x) (b^{\max} - b^{\min}) \\ \bar{o}_{ij}^x &= (j + b_{ij}^x)c_w \end{aligned}$$

where b^{\min} and b^{\max} are fixed real numbers which, in effect, impose bounds on the distance between the object and the grid cell. \bar{z}_{ij}^h and \bar{z}_{ij}^w parameterize the size of the object as:

$$\bar{o}_{ij}^h = \text{sigmoid}(\bar{z}_{ij}^h)a_h \quad \bar{o}_{ij}^w = \text{sigmoid}(\bar{z}_{ij}^w)a_w$$

²Supplementary Material can be found at: e2crawfo.github.io/misc/silot_supp.pdf

for fixed real numbers a_h and a_w . (a_h, a_w) can be interpreted as the dimensions of an *anchor box* as used in supervised object detection (Ren et al. 2015). Specifying object size with respect to (a_h, a_w) , as opposed to the size of the input frame, ensures that $\bar{o}_{i,j}^h$ and $\bar{o}_{i,j}^w$ are meaningful regardless of the spatial dimensions of the input frame.

Predicting \bar{o}^{what} , \bar{o}^{depth} and \bar{o}^{pres} . In order to obtain highly location-specific information from the image, an array of glimpses \bar{g} (one glimpse per discovery unit) is extracted from the image using spatial transformers τ (Jaderberg et al. 2015) parameterized by \bar{o}^{where} . These glimpses are then mapped to a feature volume v^{obj} by a network d_ϕ^{obj} :

$$\begin{aligned}\bar{g} &= \tau(x, \bar{o}^{\text{where}}) \\ v^{\text{obj}} &= d_\phi^{\text{obj}}(\bar{g})\end{aligned}$$

We then autoregressively predict the remaining attributes:

$$\begin{aligned}\bar{\mu}^{\text{what}}, \bar{\sigma}^{\text{what}} &= d_\phi^{\text{what}}(v, v^{\text{obj}}, \bar{o}^{\text{where}}) \\ \bar{z}^{\text{what}} &\sim N(\bar{\mu}^{\text{what}}, \bar{\sigma}^{\text{what}}) \\ \bar{o}^{\text{what}} &= \bar{z}^{\text{what}} \\ \bar{\mu}^{\text{depth}}, \bar{\sigma}^{\text{depth}} &= d_\phi^{\text{depth}}(v, v^{\text{obj}}, \bar{o}^{\text{where}}, \bar{o}^{\text{what}}) \\ \bar{z}^{\text{depth}} &\sim N(\bar{\mu}^{\text{depth}}, \bar{\sigma}^{\text{depth}}) \\ \bar{o}^{\text{depth}} &= \text{sigmoid}(\bar{z}^{\text{depth}}) \\ \bar{\mu}^{\text{pres}} &= d_\phi^{\text{pres}}(v, v^{\text{obj}}, \bar{o}^{\text{where}}, \bar{o}^{\text{what}}, \bar{o}^{\text{depth}}) \\ \bar{z}^{\text{pres}} &\sim \text{Logistic}(\bar{\mu}^{\text{pres}}) \\ \bar{o}^{\text{pres}} &= \text{sigmoid}(\bar{z}^{\text{pres}})\end{aligned}$$

Note that \bar{o}^{pres} can be viewed as a set of BinConcrete random variables (Maddison, Mnih, and Teh 2016). The significance of this choice is discussed below in Section 3.5.

3.4 Propagation

Propagation at time t takes in the current frame $x_{(t)}$ and the objects from the previous timestep $o_{(t-1)}$, and propagates the objects forward in time, using information from the current frame to update the object attributes. Let K be the number of objects from the previous timestep. In this section, all variables are matrices with K as their leading dimension. All networks are fully-connected networks that are applied to each matrix row independently (i.e. ‘‘object-wise’’) and in parallel. This is similar to the convention used in the discovery module, but with the ‘‘object’’ dimension K taking the place of the two spatial dimensions (H, W) . The structure of propagation for an individual object is shown in Figure 4. Here we outline the major points of the propagation module; for additional details, see Section A.1 of the Supplementary Material. To reduce visual clutter, assume that all variables have temporal index t unless otherwise specified.

We begin by computing a feature vector for each object in $o_{(t-1)}$. In order to handle interactions between objects, we can have the feature vector for an object depend on other objects as well. Here we make the assumption that object interactions are spatially local (which is enough to handle collisions, for example). Thus we compute features using a

spatial attention step similar to the one used in the Discovery module, allowing the features for an object to depend on attributes of nearby objects:

$$u^{\text{td}} = \text{SpatialAttention}_\phi^{\text{prop}}(o_{(t-1)}, \sigma)$$

Next we need to condition attribute updates on the current frame. Rather than condition on the entire frame, we instead extract a set of glimpses in a region near each object’s location from the previous timestep:

$$\begin{aligned}u^{\text{where}} &= o_{(t-1)}^{\text{where}} + 0.1 \cdot p_\phi^{\text{glimpse}}(u^{\text{td}}) \\ \tilde{g} &= \tau(x, u^{\text{where}}) \\ u^{\text{bu}} &= p_\phi^{\text{bu}}(\tilde{g})\end{aligned}$$

A network p_ϕ^{fuse} then combines bottom-up and top-down information:

$$u = p_\phi^{\text{fuse}}(u^{\text{bu}}, u^{\text{td}})$$

From here we autoregressively predict new values for the object attributes; this is similar to attribute prediction in the discovery module, except that rather than directly predicting attribute values, we predict attribute *updates* and subsequently apply them:

$$\begin{aligned}\tilde{\mu}^{\text{where}}, \tilde{\sigma}^{\text{where}} &= p_\phi^{\text{where}}(u) \\ \tilde{z}^{\text{where}} &\sim N(\tilde{\mu}^{\text{where}}, \tilde{\sigma}^{\text{where}}) \\ \tilde{o}^{\text{where}} &= f^{\text{where}}(o_{(t-1)}^{\text{where}}, \tilde{z}^{\text{where}})\end{aligned}$$

Next we extract and process another set of glimpses at \tilde{o}^{where} :

$$\begin{aligned}\tilde{g}' &= \tau(x, \tilde{o}^{\text{where}}) \\ u^{\text{obj}} &= p_\phi^{\text{obj}}(\tilde{g}')\end{aligned}$$

Finally we predict updates to the remaining attributes:

$$\begin{aligned}\tilde{\mu}^{\text{what}}, \tilde{\sigma}^{\text{what}} &= p_\phi^{\text{what}}(u, u^{\text{obj}}, \tilde{o}^{\text{where}}) \\ \tilde{z}^{\text{what}} &\sim N(\tilde{\mu}^{\text{what}}, \tilde{\sigma}^{\text{what}}) \\ \tilde{o}^{\text{what}} &= f^{\text{what}}(o_{(t-1)}^{\text{what}}, \tilde{z}^{\text{what}}) \\ \tilde{\mu}^{\text{depth}}, \tilde{\sigma}^{\text{depth}} &= p_\phi^{\text{depth}}(u, u^{\text{obj}}, \tilde{o}^{\text{where}}, \tilde{o}^{\text{what}}) \\ \tilde{z}^{\text{depth}} &\sim N(\tilde{\mu}^{\text{depth}}, \tilde{\sigma}^{\text{depth}}) \\ \tilde{o}^{\text{depth}} &= f^{\text{depth}}(o_{(t-1)}^{\text{depth}}, \tilde{z}^{\text{depth}}) \\ \tilde{\mu}^{\text{pres}} &= p_\phi^{\text{pres}}(u, u^{\text{obj}}, \tilde{o}^{\text{where}}, \tilde{o}^{\text{what}}, \tilde{o}^{\text{depth}}) \\ \tilde{z}^{\text{pres}} &\sim \text{Logistic}(\tilde{\mu}^{\text{pres}}) \\ \tilde{o}^{\text{pres}} &= o_{(t-1)}^{\text{pres}} \cdot \text{sigmoid}(\tilde{z}^{\text{pres}})\end{aligned}$$

Notice that propagation cannot increase the value of the *pres* attribute, due to the form of the update (multiplication by a sigmoid). This ensures that objects are only ever discovered by the Discovery module, which is better equipped for it. f^{where} , f^{what} and f^{depth} determine the forms of the other updates; full details are left to the Supplementary Material.

Propagation in SILOT is similar to propagation in SQAIR, with one significant exception. In SQAIR, objects within a timestep are updated sequentially; this allows objects within a

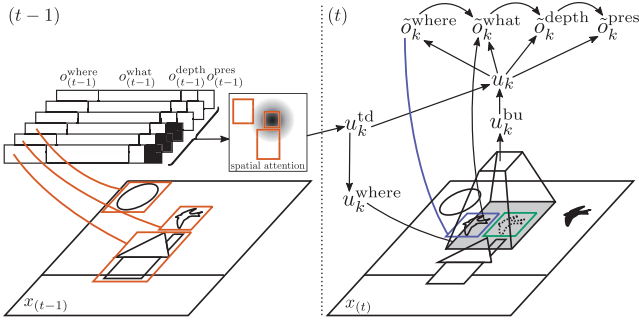


Figure 4: Schematic depicting the propagation module updating an object with index k , tracking the location of the white bird. A feature vector for the object, which also takes into account nearby objects, is first created using spatial attention. Next, an initial glimpse (grey region) is specified with respect to the object’s location from the previous time step (green box). This glimpse is then processed by a neural network (trapezoid) and used to predict and apply an update to the *where* attribute, resulting in $\tilde{o}_k^{\text{where}}$. The blue box corresponds to the location of the image referred to by $\tilde{o}_k^{\text{where}}$. Another glimpse is extracted at location $\tilde{o}_k^{\text{where}}$, and updates to the remaining attributes are predicted and applied autoregressively. Here we have omitted the latent propagation variables \tilde{z}_k .

timestep to condition on one another, facilitating coordination between objects and supporting behavior such as explaining away. However, this sequential processing can be computationally demanding when there are large numbers of objects. In contrast, SILOT updates all objects within a timestep in parallel; a degree of coordination between objects is achieved via the spatial attention step.

3.5 Selection

Past architectures in this domain, particularly AIR and SQAIR, use discrete Bernoulli random variables for their equivalents of the *pres* attribute, and are thus forced to employ reinforcement learning–inspired techniques for estimating gradients. In order to avoid this complication and make our architecture fully differentiable, we opted to model o^{pres} as real-valued BinConcrete random variables, which are differentiable relaxations of Bernoullis (Maddison, Mnih, and Teh 2016). One downside of this choice is that objects never “go away”, even if $o^{\text{pres}} = 0$; all posited objects are always present, but to varying degrees. If care is not taken this will result in scaling issues, since the discovery module yields HW new objects per frame; if all these objects are kept and propagated forward, we would end up with a collection of $HW(t + 1)$ objects at timestep t , which would quickly become intractable.

To fix this, we use a simple top- K selection strategy wherein we keep only the K objects from the union of $\tilde{o}_{(t)}$ and $\bar{o}_{(t)}$ with highest values for the *pres* attribute, for fixed integer K . While this hard selection step is not differentiable, the number of objects with non-negligible values for the *pres* attribute is intended to be small compared to $HW(t + 1)$ (see Section 3.8), and we have not found this non-differentiability

to cause problems in training as long as K is large enough. As a rule-of-thumb, we typically set K to be roughly 25% larger than the maximum number of objects that we expect to see in a single frame; this ensures there is always room for a reasonable number of objects with low values for *pres* to be propagated and receive gradient feedback. The output of this selection step is the final set of objects $o_{(t)}$.

3.6 Rendering

The differentiable rendering module is the sole constituent of the VAE decoder $g_\theta(x|z)$, taking in the current set of objects $o_{(t)}$ and yielding a reconstructed frame $\hat{x}_{(t)}$. The rendering process is highly structured, and this structure gives meaning to the object attributes. Object appearance and transparency are predicted from o^{what} , the object is placed at location o^{where} via spatial transformers, o^{depth} parameterizes a differentiable approximation of relative depth, and object transparency is multiplied by o^{pres} so that objects are only rendered to the extent that they exist. Full details can be found in Section A.2 of the Supplementary Material.

3.7 Training

Recall that propagation, discovery and selection form the VAE encoder $h_\phi(z|x)$, while rendering forms the VAE decoder $g_\theta(x|z)$. The network is trained by maximizing the evidence lower bound (Kingma and Welling 2013):

$$\mathcal{L}(\phi, \theta) := E_{x \sim h(x), z \sim h_\phi(z|x)} \left[\log \left(\frac{g_\theta(x|z)g(z)}{h_\phi(z|x)} \right) \right] \quad (1)$$

$$\phi^*, \theta^* = \arg \max_{\phi, \theta} \mathcal{L}(\phi, \theta) \quad (2)$$

Here $h(x)$ is the distribution over videos defined by the dataset, and $g(z)$ is the prior distribution over z . The optimization is performed by gradient ascent.

Curriculum Learning. Following SQAIR, we train the network with a form of curriculum learning. We begin by training on only the first 2 frames of each video. We then increase the number of training frames by 2 every N_{curric} update steps. After $(\lceil T/2 \rceil - 1)N_{\text{curric}}$ update steps the network will be training on complete videos. This was observed to help with training stability, possibly because it allows the network to learn object appearances before learning object dynamics.

Discovery Dropout. Early in development we found the network often tried to predict objects for new frames exclusively by way of the discovery module rather than propagating objects from the previous frame. This strategy can yield reasonable reconstruction performance, but will fail at object tracking since object identities are not maintained from one frame to the next. Moreover, this is a local minimum; once the network starts relying completely on discovery, propagation stops being trained and cannot improve. To discourage this behavior, we designed a technique that we call *discovery dropout*. Each timestep other than $t = 0$, the entire discovery module is turned off with probability p_{dd} . This forces the network do as much as possible through propagation rather than discovery, since the network is never sure whether the discovery module will be turned on for the next timestep. Throughout this work we use $p_{\text{dd}} = 0.5$.




Figure 5: Probing object tracking performance as number of digits per video varies in the Scattered MNIST task. All points are averages over 6 random seeds, filled regions are standard deviations (except for the fully deterministic ConnComp algorithm).

3.8 Prior Distribution

A crucial component of a VAE is the prior over latent variables, $g(z)$, which can be used to influence the statistics of latent variables yielded by the encoder. For the majority of the latents, we assume independent Normal distributions. However, for the Logistic random variables \tilde{z}^{pres} and \tilde{z}^{pres} we design a prior that puts pressure on the network to reconstruct the video using as few objects as possible (i.e. few objects with large values for o^{pres}). This pressure is necessary for the network to extract quality object-like representations; without it, the network is free to set all $o^{\text{pres}} = 1$, and the extracted objects become meaningless. The prior is identical to the one used in SPAIR (Crawford and Pineau 2019), applied to the union of the discovered and propagated objects for a given timestep.

As done in SQAIR, we can also use a learned prior $g_\theta(z)$, whose parameters can be trained alongside the other parameters. In the current work, we do this only for the propagation module. In effect, this amounts to learning a separate prior propagation module which does not take into account information from the current frame. This prior propagation module employs an architecture that is similar to the main propagation module, except that it omits the glimpse components which extract information from the current frame. The prior propagation module thus is forced to truly learn the dynamics of the objects, whereas the main propagation module can rely on information from the current frame. In the training objective (Equation (1)), we use the static prior $g(z)$ for the discovery latents, and an even mixture of the static and learned priors for the propagation latents.

4 Experiments

We tested SILOT in a number of challenging object discovery and tracking tasks, emphasizing large videos containing many objects. We use 3 metrics to assess model performance: Multi-Object Tracking Accuracy (MOTA) (with IoU threshold 0.5), a standard measure of object tracking performance (Milan et al. 2016); Average Precision (AP) (with IoU=0.1:0.1:0.9), a standard measure of object detection performance (Everingham et al. 2010); and Count Abs. Error, or the absolute difference between the number of objects predicted by the

model and the ground truth number of objects (on a frame-by-frame basis). Intersection over Union (IoU) is a measure of the overlap between two objects.

We compare against a simple baseline algorithm called ConnComp (Connected Components), which works by treating each connected region of similarly-colored pixels in a frame as a separate object, and computing object identities over time using the Hungarian algorithm (Kuhn 1955). The performance of ConnComp can be interpreted as a measure of the difficulty of the dataset; it will be successful only to the extent that objects can be tracked by color alone.

Code for running these experiments is available online at github.com/e2crawfo/silot. To supplement the quantitative results provided in this section, we also provide qualitative results: videos depicting the performance of trained SILOT networks can be found at sites.google.com/view/silot, and static visualizations are shown in Supplementary Material.

4.1 Scattered MNIST

In this experiment, each 8-frame video is generated by first selecting a number of MNIST digits to include and sampling that number of digits. Videos have spatial size 48×48 pixels, digits are scaled down to 14×14 pixels, and initial digit positions are randomly chosen so that the maximum overlap allowed is small but non-trivial. For each digit we also uniformly sample an initial velocity vector with a fixed magnitude of 2 pixels per frame. Digits pass through one another without event, and bounce off the edges of the frames.

In order to test generalization ability, we used two training conditions: training on videos containing 1–6 digits, and training on videos containing 1–12 digits. In both cases, we test performance on videos containing up to 12 digits; networks trained in the 1–6 condition will thus be required to generalize beyond their training experience.

We compare SILOT against SQAIR. We tested SQAIR with two different discovery backbone networks (the network which first processes each input frame): an MLP, and a convolutional network with the same structure as SILOT’s convolutional backbone d_ϕ^{bu} . Note that using a convolutional network as a backbone, rather than an MLP, should improve SQAIR’s spatial invariance to a degree; however it will still




Figure 6: Probing SILOT’s object tracking performance as the number of shapes per video varies in the Scattered Shapes task. All points are averages over 4 random seeds, and filled regions are standard deviations (except for the fully deterministic ConnComp algorithm). We manipulated 2 different aspects of the training setup: training on images that contain 1–10 shapes vs 21–30 shapes, and training on random 60×60 crops vs. full 96×96 images (full videos are always used at test time).

lack, among other features, SILOT’s local object specification scheme, which is part of what allows SILOT’s discovery module to be interpreted as an array of local discovery units.

We also experimented with Tracking by Animation (TbA) (He et al. 2018), but were unable to obtain good tracking performance on these densely cluttered videos. One relevant point is that TbA lacks a means of encouraging the network to explain scenes using few objects, and we found TbA often using several internal objects to explain a single object in the video; in contrast, both SILOT and SQAIR use priors on o^{pres} which encourage o^{pres} to be near 0, forcing the networks to use internal objects efficiently.

Results are shown in Figure 5. We have omitted both SQAIR networks trained on the 1–12 digit condition; SQAIR’s simpler discovery module was unable to handle the more densely packed scenes, resulting in highly varied performance that would make the plot unreadable. From the plot we can see that SILOT outperforms SQAIR by a large margin (especially when extrapolating beyond training experience), and that SILOT trained on 1–6 digits suffers only a minor performance hit compared to SILOT trained on 1–12 digits, a result of SILOT’s spatially invariant architecture. In the Supplementary Material, we also show results obtained when using the learned prior propagation module (see Section 3.8) in place of the regular propagation module.

Note the slight improvement in performance of the AP of the SQAIR networks when going from 1 to 2 digits. This reflects an interesting strategy developed by the SQAIR networks. On the first timestep, they often turn “on” several more objects than are in the video; this pushes down the AP score, with a greater negative effect when there are fewer objects. We hypothesize that it settles on this strategy because the backbone networks are unable to get an accurate read on the number of objects present, and so the network settles for a strategy of casting a “wide net” on the first timestep.

Additionally, when testing on videos containing 1–6 digits, the SQAIR networks were used in the same fashion as they were during training. However, when testing on videos containing 7 or more digits, we found that the network’s performance at guessing the number of objects to use became

very poor; thus, for such test cases we artificially told the network the correct number of objects to use by running SQAIR’s recurrent network for the correct number of steps.

4.2 Scattered Shapes

In this experiment we tackle larger images with significantly more objects to track. Here each video has spatial size 96×96 , and contains a number of moving monochrome shapes. We use 6 different colors and 5 different shapes. Initial shape velocities are uniformly sampled with a fixed magnitude of 5 pixels per frame.

Here we test ability to generalize to scenes that are larger and/or contain different numbers of objects than scenes encountered during training. To assess the former, networks were trained on either random 60×60 crops of input videos, or full 96×96 videos (note that SILOT’s fully convolutional discovery module allows it to process videos of any size). To assess the latter, networks were trained on videos containing either 1–10 shapes or 21–30 shapes. This yields 4 different training conditions. At test time, we use full 96×96 videos containing up to 35 shapes.

Results are shown in Figure 6. There we see that all SILOT networks drastically outperform the baseline. Additionally, there is significant evidence that the networks are able to generalize well to both different numbers of shapes and different sized videos than what was encountered during training. In particular, the network trained on 60×60 crops of videos containing 1–10 shapes achieves reasonable performance when asked to process 96×96 images containing up to 35 shapes.

4.3 Atari

We also applied SILOT to videos obtained from select Atari games (Bellemare et al. 2013) (the possibility of applying a SQAIR-style model to reinforcement learning was first suggested in (Kosiorek et al. 2018)). Results are shown in Table 1. Ground truth object labels for evaluation were obtained using the ConnComp algorithm, which is reasonably effective for the chosen games. Our goal here was to push scalability, rather than ability to deal with e.g. overlapping objects.

Game	MOTA	AP	Count Abs. Error
Space Invaders	.89	.73	2.94
Asteroids	.67	.67	1.81

Table 1: SILOT performance on Atari videos.

5 Conclusion

In the current work we proposed an architecture for scalable unsupervised object tracking, and empirically demonstrated a number of its benefits. Many interesting directions are left for future research. For one, it should be possible to extend the discovery module to discover objects at different scales, similar to the supervised object detection architecture SSD (Liu et al. 2016). Additionally, in order to have SILOT and related approaches work on real videos, it will be necessary to include a means of dealing with complex, dynamic backgrounds; past work in this area has generally preprocessed videos with an off-the-shelf background subtractor, largely avoiding the problem of modeling the background (though see concurrent work in (Jiang et al. 2019) which makes significant progress on this issue). Finally, as suggested in (Kosiorok et al. 2018), these architectures are quite complex, and it would be worthwhile to perform ablation studies to determine whether all components are necessary for good performance.

Acknowledgements

Funding for this work was provided by Samsung.

References

Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* 47:253–279.

Burgess, C. P.; Matthey, L.; Watters, N.; Kabra, R.; Higgins, I.; Botvinick, M.; and Lerchner, A. 2019. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*.

Carey, S. 2009. *The origin of concepts*. Oxford University Press.

Chang, M. B.; Ullman, T.; Torralba, A.; and Tenenbaum, J. B. 2016. A compositional object-based approach to learning physical dynamics. *arXiv preprint arXiv:1612.00341*.

Crawford, E., and Pineau, J. 2019. Spatially invariant, unsupervised object detection with convolutional neural networks. In *Thirty-Third AAAI Conference on Artificial Intelligence*.

Dang Ha, T. H. 2017. A guide to receptive field arithmetic for convolutional neural networks. <https://medium.com/mlreview/a-guide-to-receptive-field-arithmetic-for-convolutional-neural-networks-e0f514068807>. Accessed: Aug 7, 2019.

Diuk, C.; Cohen, A.; and Littman, M. L. 2008. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*, 240–247. ACM.

Eslami, A.; Heess, N.; Weber, T.; Tassa, Y.; Szepesvari, D.; Hinton, G. E.; et al. 2016. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, 3225–3233.

Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision* 88(2):303–338.

Goel, V.; Weng, J.; and Poupart, P. 2018. Unsupervised video object segmentation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, 5683–5694.

Greff, K.; Kaufmann, R. L.; Kabra, R.; Watters, N.; Burgess, C.; Zoran, D.; Matthey, L.; Botvinick, M.; and Lerchner, A. 2019. Multi-object representation learning with iterative variational inference. *arXiv preprint arXiv:1903.00450*.

Greff, K.; van Steenkiste, S.; and Schmidhuber, J. 2017. Neural expectation maximization. In *Advances in Neural Information Processing Systems*, 6691–6701.

He, Z.; Li, J.; Liu, D.; He, H.; and Barber, D. 2018. Tracking by animation: Unsupervised learning of multi-object attentive trackers. *arXiv preprint arXiv:1809.03137*.

Hsieh, J.-T.; Liu, B.; Huang, D.-A.; Fei-Fei, L. F.; and Niebles, J. C. 2018. Learning to decompose and disentangle representations for video prediction. In *Advances in Neural Information Processing Systems*, 517–526.

Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in neural information processing systems*, 2017–2025.

Jiang, J.; Janghorbani, S.; de Melo, G.; and Ahn, S. 2019. Scalable object-oriented sequential generative models. *arXiv preprint arXiv:1910.02384*.

Kansky, K.; Silver, T.; Mély, D. A.; Eldawy, M.; Lázaro-Gredilla, M.; Lou, X.; Dorfman, N.; Sidor, S.; Phoenix, S.; and George, D. 2017. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. In *International Conference on Machine Learning*, 1809–1818.

Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.

Kosiorok, A.; Kim, H.; Teh, Y. W.; and Posner, I. 2018. Sequential attend, infer, repeat: Generative modelling of moving objects. In *Advances in Neural Information Processing Systems*, 8606–8616.

Kuhn, H. W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2(1-2):83–97.

Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*, 21–37. Springer.

Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2016. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*.

Milan, A.; Leal-Taixé, L.; Reid, I.; Roth, S.; and Schindler, K. 2016. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*.

Redmon, J., and Farhadi, A. 2017. Yolo9000: better, faster, stronger. *arXiv preprint*.

Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.

van Steenkiste, S.; Chang, M.; Greff, K.; and Schmidhuber, J. 2018. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. *arXiv preprint arXiv:1802.10353*.

Wang, D.; Jamnik, M.; and Lio, P. 2019. Unsupervised and interpretable scene discovery with discrete-attend-infer-repeat. *arXiv preprint arXiv:1903.06581*.

Zambaldi, V.; Raposo, D.; Santoro, A.; Bapst, V.; Li, Y.; Babuschkin, I.; Tuyls, K.; Reichert, D.; Lillicrap, T.; Lockhart, E.; et al. 2018. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*.