

Tracking Emerges by Looking Around Static Scenes, with Neural 3D Mapping

Adam W. Harley, Shrinidhi Kowshika Lakshmikanth,
Paul Schydlo, and Katerina Fragkiadaki

Carnegie Mellon University, Pittsburgh PA, 15213, USA
{aharley,kowshika,pschydlo,katef}@cs.cmu.edu

Abstract. We hypothesize that an agent that can look around in static scenes can learn rich visual representations applicable to 3D object tracking in complex dynamic scenes. We are motivated in this pursuit by the fact that the physical world itself is mostly static, and multiview correspondence labels are relatively cheap to collect in static scenes, e.g., by triangulation. We propose to leverage multiview data of *static points* in arbitrary scenes (static or dynamic), to learn a neural 3D mapping module which produces features that are correspondable across time. The neural 3D mapper consumes RGB-D data as input, and produces a 3D voxel grid of deep features as output. We train the voxel features to be correspondable across viewpoints, using a contrastive loss, and correspondability across time emerges automatically. At test time, given an RGB-D video with approximate camera poses, and given the 3D box of an object to track, we track the target object by generating a map of each timestep and locating the object’s features within each map. In contrast to models that represent video streams in 2D or 2.5D, our model’s 3D scene representation is disentangled from projection artifacts, is stable under camera motion, and is robust to partial occlusions. We test the proposed architectures in challenging simulated and real data, and show that our unsupervised 3D object trackers outperform prior unsupervised 2D and 2.5D trackers, and approach the accuracy of supervised trackers. This work demonstrates that 3D object trackers can emerge without tracking labels, through multiview self-supervision on static data.

1 Introduction

A large part of the real world almost never moves. This may be surprising, since moving entities easily attract our attention [13], and because we ourselves spend most of our waking hours continuously in motion. Objects like roads and buildings, however, stay put. Can we leverage this property of the world to learn visual features suitable for interpreting complex scenes with many moving objects?

In this paper, we hypothesize that a correspondence module learned in static scenes will also work well in dynamic scenes. This is motivated by the fact that the content of dynamic scenes is the same as the content of static scenes. We

would like our hypothesis to be true, because correspondences are far cheaper to obtain in static scenes than in dynamic ones. In static scenes, one can simply deploy a Simultaneous Localization and Mapping (SLAM) module to obtain a 3D reconstruction of the scene, and then project reconstructed points back into the input imagery to obtain multiview correspondence labels. Obtaining correspondences in dynamic scenes would require taking into account the motion of the objects (i.e., tracking).

We propose to leverage multiview data of static points in arbitrary scenes (static or dynamic), to learn a neural 3D mapping module which produces features that are correspondable across viewpoints *and* timesteps. The neural 3D mapper consumes RGB-D (color and depth) data as input, and produces a 3D voxel grid of deep features as output. We train the voxel features to be correspondable across viewpoints, using a contrastive loss. At test time, given an RGB-D video with approximate camera poses, and given the 3D box of an object to track, we track the target object by generating a map of each timestep and locating the object’s features within each map.

In contrast to models that represent video streams in 2D [48,49,46], our model’s 3D scene representation is disentangled from camera motion and projection artifacts. This provides an inductive bias that scene elements maintain their size and shape across changes in camera viewpoint, and reduces the need for scale invariance in the model’s parameters. Additionally, the stability of the 3D map under camera motion allows the model to constrain correspondence searches to the 3D area where the target was last seen, which is a far more reliable cue than 2D pixel coordinates. In contrast to models that use 2.5D representations (e.g., scene flow [27]), our neural 3D maps additionally provide features for partially occluded areas of the scene, since the model can infer their features from context. This provides an abundance of additional 3D correspondence candidates at test time, as opposed to being limited to the points observed by a depth sensor.

Our work builds on geometry-aware recurrent neural networks (GRNNs) [43]. GRNNs are modular differentiable neural architectures that take as input RGB-D streams of a static scene under a moving camera and infer 3D feature maps of the scene, estimating and stabilizing against camera motion. The work of Harley et al. [17] showed that training GRNNs for contrastive view prediction in static scenes helps semi-supervised 3D object detection, as well as moving object discovery. In this paper, we extend those works to learn from dynamic scenes with independently moving objects, and simplify the GRNN model by reducing the number of losses and modules. Our work also builds on the work of Vondrick et al. [46], which showed that 2D pixel trackers can emerge without any tracking labels, through self-supervision on a colorization task. In this work, we show that 3D voxel trackers can emerge without tracking labels, through contrastive self-supervision on static data. In the fact that we learn features from correspondences established through triangulation, our work is similar to Dense Object Nets [11], though that work used object-centric data, and used background subtraction to apply loss on static objects, whereas in our work we

do moving-object subtraction to apply loss on *anything* static. We do not assume a priori that we know which objects to focus on.

We test the proposed architectures in simulated and real datasets of urban driving scenes (CARLA [9] and KITTI [15]). We evaluate the learned 3D visual feature representations on their ability to track objects over time in 3D. We show that the learned visual feature representations can accurately track objects in 3D, simply supervised by observing static data. Furthermore, our method outperforms 2D and 2.5D baselines, demonstrating the utility of learning a 3D representation for this task instead of a 2D one.

The main contribution of this paper is to show that learning feature correspondences from static 3D points causes 3D object tracking to emerge. We additionally introduce a neural 3D mapping module which simplifies prior works on 3D inverse graphics, and learns from a simpler objective than considered in prior works. Our code and data are publicly available¹.

2 Related Work

2.1 Learning to see by moving

Both cognitive psychology and computational vision have realised the importance of motion for the development of visual perception [16,50]. Predictive coding theories [36,14] suggest that the brain predicts observations at various levels of abstraction; temporal prediction is thus of central interest. These theories currently have extensive empirical support: stimuli are processed more quickly if they are predictable [26,33], prediction error is reflected in increased neural activity [36,4], and disproven expectations lead to learning [38]. Several computational models of frame predictions have been proposed [36,10,40,43,31]. Alongside future frame prediction, predicting some form of contextual or missing information has also been explored, such as predicting frame ordering [23], temporal instance-level associations [48] color from grayscale [46], egomotion [19,1] and motion trajectory forecasting [47]. Most of these unsupervised methods are evaluated as pre-training mechanisms for object detection or classification [28,47,48,17].

Video motion segmentation literature explores the use of videos in unsupervised moving object discovery [34]. Most motion segmentation methods operate in 2D image space, and cluster 2D optical flow vectors or 2D flow trajectories to segment moving objects [29,5,12], or use low-rank trajectory constraints [7,41,8]. Our work differs in that we address object detection and segmentation in 3D as opposed to 2D, by estimating 3D motion of the “imagined” (complete) scene, as opposed to 2D motion of the pixel observation stream.

2.2 Vision as inverse graphics

Earlier works in Computer Vision proposed casting visual recognition as inverse rendering [30,53], as opposed to feedforward labelling. The “blocks world” of

¹ https://github.com/aharley/neural_3d_tracking

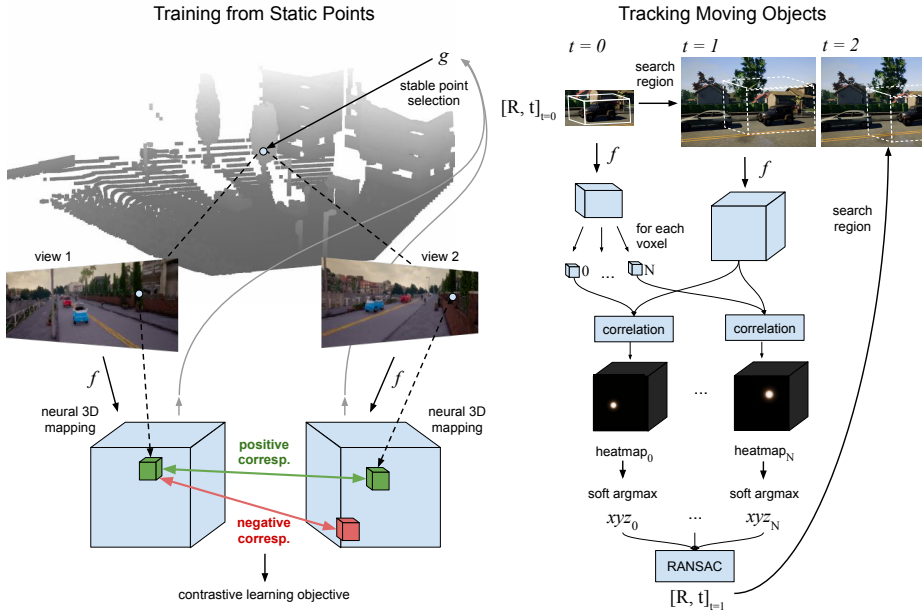


Fig. 1. Tracking emerges by looking around in static scenes, with neural 3D mapping. *Left: Training regime.* Our learned 3D neural mapper f maps RGB-D inputs into featurized 3D scene maps. Points that correspond across multiple views, provided by static data or selected by a reliability network g , provide labels for a contrastive learning objective. *Right: Testing regime.* Given an object box in the zeroth frame, our encoder creates features for the object, and searches for re-occurrences of those features in the 3D scene maps of future time steps. The model then explains the full-object motion via a robust estimate of the rigid transformation across frames. Each estimated transformation is used to initialize the search region in the next timestep. Our model does not use any human supervision, and never trains explicitly for tracking; instead, it is supervised only by observing static points in 3D scenes.

Roberts [37] had the goal of reconstructing the 3D scene depicted in the image in terms of 3D solids found in a database. A key question to be addressed is: what representations should we use for the intermediate latent 3D structures? Most works seek to map images to explicit 3D representations, such as 3D pointclouds [51,44,54,45], 3D meshes [24,21], or binary 3D voxel occupancies [42,20,52]. The aforementioned manually designed 3D representations, e.g., 3D meshes, 3D key-points, 3D pointclouds, 3D voxel grids, may not be general enough to express the rich 3D world, which contains liquids, deformable objects, clutter, dirt, wind, etc., and at the same time may be over descriptive when detail is unnecessary. In this work, we opt for *learning-based 3D feature representations* extracted end-to-end from the RGB-D input stream as proposed by Tung et al. [43] and Harley et al. [17]. We extend the architectures of those works to handle and learn from videos of dynamic scenes, as opposed to only videos of static scenes.

3 Learning to Track with Neural 3D Mapping

We consider a mobile agent that can move around in a 3D scene, and observe it from multiple viewpoints. The scene can optionally contain dynamic (moving and/or deforming) objects. The agent has an RGB camera with known intrinsics, and a depth sensor registered to the camera’s coordinate frame. It is reasonable to assume that a mobile agent who moves at will has access to its approximate egomotion, since it chooses where to move and what to look at [32]. In simulation, we use ground truth camera poses; in real data, we use approximate camera poses provided by an inertial navigation system (GPS and IMU). In simulation, we use random viewpoints; in real data, we use just one forward-facing camera (which is all that is available). Note that a more sophisticated mobile agent might attempt to select viewpoints intelligently at training time.

Given the RGB-D and egomotion data, our goal is to learn 3D feature representations that can correspond entities across time, despite variations in pose and appearance. We achieve this by training inverse graphics neural architectures that consume RGB-D videos and infer 3D feature maps of the full scenes, as we describe in Section 3.1. To make use of data where some parts are static and other parts are moving, we learn to identify static 3D points by estimating a reliability mask over the 3D scene, as we describe in Section 3.2. Finally, we track in 3D, by re-locating the object within each timestep’s 3D map, as described in Section 3.3. Figure 1 shows an overview of the training and testing setups.

3.1 Neural 3D Mapping

Our model learns to map an RGB-D (RGB and depth) image to a 3D feature map of the scene in an end-to-end differentiable manner. The basic architecture is based on prior works [43,17], which proposed view prediction architectures with a 3D bottleneck. In our case, the 3D feature map is the *output* of the model, rather than an intermediate representation.

Let $\mathcal{M} \in \mathbb{R}^{W \times H \times D \times C}$ denote the 3D feature map representation, where W, H, D, C denote the width, height, depth and number of feature channels, respectively. The map corresponds to a large cuboid of world space, placed at some pose of interest (e.g., surrounding a target object). Every (x, y, z) location in the 3D feature map \mathcal{M} holds a C -length feature vector that describes the semantic and geometric content of the corresponding location of the world scene. To denote the feature map of timestep i , we write $\mathcal{M}^{(i)}$. We denote the function that maps RGB-D inputs to 3D feature maps as $f : (I, D) \mapsto \mathcal{M}$. To implement this function, we voxelize the inputs into a 3D grid, then pass this grid through a 3D convolutional network, and L_2 -normalize the outputs.

Tung et al. [43] learned the parameters of f by predicting RGB images of unseen viewpoints, and applying a regression loss; Harley et al. [17] demonstrated that this can be outperformed by contrastive prediction objectives, in 2D and 3D. Here, we drop the view prediction task altogether, and focus entirely on a 3D correspondence objective: if a static point (x, y, z) is observed in two views i and j , the corresponding features $m_i = \mathcal{M}_{x,y,z}^{(i)}, m_j = \mathcal{M}_{x,y,z}^{(j)}$ should be similar

to each other, and distinct from other features. We achieve this with a cross entropy loss [39,31,18,6]:

$$\mathcal{L}_{i,j} = -\log \frac{\exp(m_i^\top m_j / \tau)}{\sum_{k \neq i} \exp(m_i^\top m_k / \tau)}, \quad (1)$$

where τ is a temperature parameter, which we set to 0.07, and the sum over $k \neq i$ iterates over non-corresponding features. Note that indexing correctly into the 3D scene maps to obtain the correspondence pair m_i, m_j requires knowledge of the relative camera transformation across the input viewpoints; we encapsulate this registration and indexing in the notation $\mathcal{M}_{x,y,z}$. Following He et al. [18], we obtain a large pool of negative correspondences through the use of an offline dictionary, and stabilize training with a “slow” copy of the encoder, f_{slow} , which is learned via high-momentum updates from the main encoder parameters.

Since the neural mapper (a) does not know a priori which voxels will be indexed for a loss, and (b) is fully convolutional, it learns to generate view-invariant features densely in its output space, even though the supervision is sparse. Furthermore, since (a) the model is encouraged to generate correspondable features invariant to viewpoint, and (b) varying viewpoints provide varying contextual support for 3D locations, the model learns to *infer* correspondable features from limited context, which gives it robustness to partial occlusions.

3.2 Inferring static points for self-supervision in dynamic scenes

The training objective in the previous subsection requires the location of a static point observed in two or more views. In a scenario where the data is made up entirely of static scenes, as can be achieved in simulation or in controlled environments, obtaining these static points is straightforward: any point on the surface of the scene will suffice, provided that it projects into at least two camera views.

To make use of data where some parts are static and other parts are moving, we propose to simply discard data that appears to be moving. We achieve this by training a neural module to take the difference of two scene features $\mathcal{M}^{(i)}, \mathcal{M}^{(j)}$ as input, and output a “reliability” mask indicating a per-voxel confidence that the scene cube within the voxel is static: $g : (\mathbf{sg}(\mathcal{M}^{(i)} - \mathcal{M}^{(j)})) \mapsto \mathcal{R}^{W,H,D}$, where \mathbf{sg} stops gradients from flowing from g into the function f which produces $\mathcal{M}^{(i)}, \mathcal{M}^{(j)}$. We implement g as a per-voxel classifier, with 2-layer fully-connected network applied fully convolutionally. We do not assume to have true labels of moving/static voxels, so we generate synthetic labels using static data: given two maps of the *same* scene $\mathcal{M}^{(i)}, \mathcal{M}^{(j)}$, we generate positive-label inputs with $\mathcal{M}^{(i)} - \mathcal{M}^{(j)}$ (as normal), and generate negative-label inputs with $\mathcal{M}^{(i)} - \text{shuffle}(\mathcal{M}^{(j)})$, where the shuffle operation ruins the correspondence between the two tensors. After this training, we deploy this network on pairs of frames from dynamic scenes, and use it to select high-confidence static data to further train the encoder f .

Our training procedure is then: (1) learn the encoder f on static data; (2) learn the reliability function g ; (3) in dynamic data, finetune f on data selected by g . Steps 2-3 can be repeated a number of times. In practice we find that results do not change substantially after the first pass.

3.3 Tracking via point correspondences

Using the learned neural 3D mapper, we track an object of interest over long temporal horizons by re-locating it in a map produced at each time step. Specifically, we re-locate each voxel of the object, and use these new locations to form a motion field. We convert these voxel motions into an estimate for the entire object by fitting a rigid transformation to the correspondences, via RANSAC.

We assume we are given the target object’s 3D box on the zeroth frame of a test video. Using the zeroth RGB-D input, we generate a 3D scene map centered on the object. We then convert the 3D box into a set of coordinates X_0, Y_0, Z_0 which index into the map. Let $m_i = \mathcal{M}_{x_i \in X_0, y_i \in Y_0, z_i \in Z_0}^{(0)}$ denote a voxel feature that belongs to the object. On any subsequent frame t , our goal is to locate the new coordinate of this feature, denoted x_i^t, y_i^t, z_i^t . We do so via a soft spatial argmax, using the learned feature space to provide correspondence confidences:

$$(x_i^t, y_i^t, z_i^t) = \sum_{x \in X, y \in Y, z \in Z} \left(\frac{\exp(m_i^\top \mathcal{M}_{x,y,z}^{(t)})}{\sum_{x_2 \in X, y_2 \in Y, z_2 \in Z} \exp(m_i^\top \mathcal{M}_{x_2, y_2, z_2}^{(t)})} (x, y, z) \right), \quad (2)$$

where X, Y, Z denote the set of coordinates in the search region. We then compute the motion of the voxel as $(\delta_{x_i}^t, \delta_{y_i}^t, \delta_{z_i}^t) = (x_i^t, y_i^t, z_i^t) - (x_i, y_i, z_i)$. After computing the motion of every object voxel in this way, we use RANSAC to find a rigid transformation that explains the majority of the correspondences. We apply this rigid transform to the input box, yielding the object’s location in the next frame.

Vondrick et al. [46] computed a similar attention map during tracking (in 2D), but did not compute its soft argmax, nor explain the object’s motion with a single transformation, but rather propagated a “soft mask” to the target frame, which is liable to grow or shrink. Our method takes advantage of the fact that all coordinates are 3D, and makes the assumption that the objects are rigid, and propagates a fixed-size box from frame to frame.

We empirically found that it is critical to constrain the search region of the tracker in 3D. In particular, on each time step we create a search region centered on the object’s last known position. The search region is $16m \times 2m \times 16m$ large, which is half of the typical full-scene resolution. This serves three purposes. The first is: it limits the number of spurious correspondences that the model can make, since it puts a cap on the metric range of the correspondence field, and thereby reduces errors. Second: it “re-centers” the model’s field of view onto the target, which allows the model to incorporate maximal contextual information surrounding the target. Even if the bounds were sufficiently narrow to reduce spurious correspondences, an object at the *edge* of the field of view will have less-informed features than an object at the middle, due to the model’s convolutional

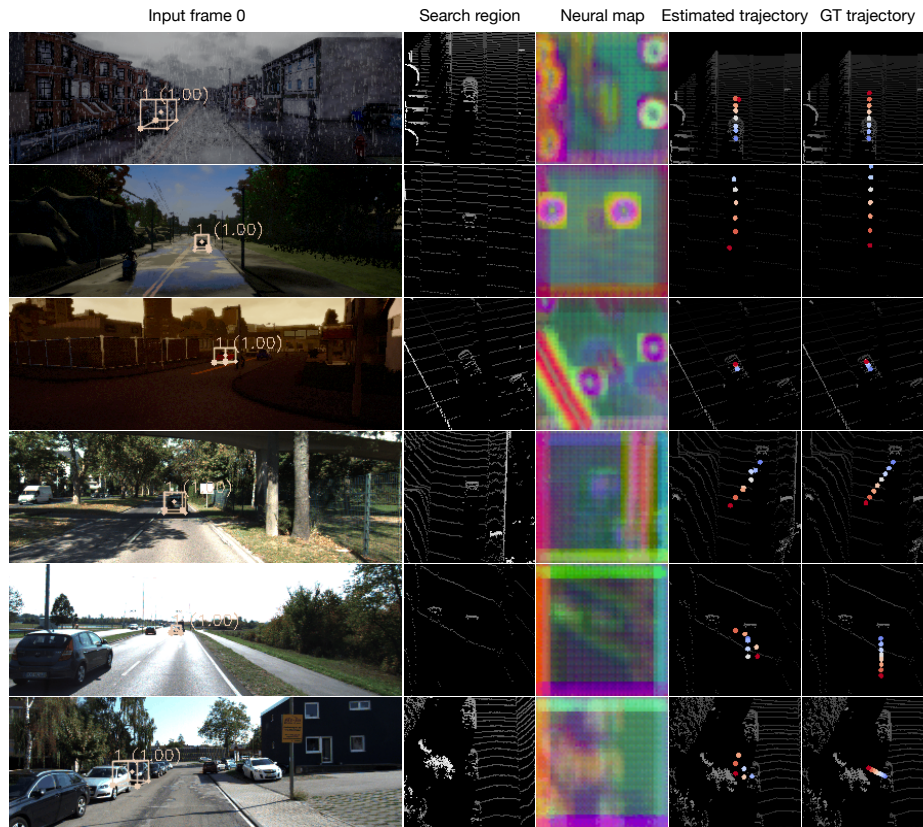


Fig. 2. Visualization of tracking inputs, inferred 3D scene features, and tracking outputs. Given the input frame on the left, along with a 3D box specifying the object to track, the model (1) creates a search region for the object centered on the object’s position (bird’s eye view of occupancy shown), (2) encodes this region into a 3D neural map (bird’s eye view of PCA compression shown), and (3) finds correspondences for the object’s features. The top three rows show CARLA results; the bottom three rows show KITTI results.

architecture. The third reason is computational: even 2D works [46] struggle with the computational expense of the large matrix multiplications involved in this type of soft attention, and in 3D the expense is higher. Searching locally instead of globally makes Eq. 2 tractable.

4 Experiments

We test our model in the following two datasets:

1. **Synthetic RGB-D videos of urban scenes rendered in the CARLA simulator** [9]. CARLA is an open-source photorealistic simulator of urban




Fig. 3. Visualization of tracking inputs, inferred 2D scene features, and tracking outputs, for our adaptation of Dense Object Nets [11]. This model operates in the same way as our full 3D one, but creates features in 2D, and unprojects them to a sparse 3D grid for tracking.

driving scenes. It permits moving the camera to any desired viewpoint in the scene. We obtain data from the simulator as follows. We begin by generating 10000 autopilot episodes of 16 frames each, at 10 FPS. We define 18 viewpoints along a 40m-radius hemisphere anchored to the ego-car (i.e., it moves with the car). In each episode, we sample 6 random viewpoints from the 18 and randomly perturb their pose, and then capture each timestep of the episode from these 6 viewpoints. We discard episodes that do not have an object in-bounds for the duration of the episode.

We treat the Town1 data as the “training” set, and the Town2 data as the “test” set, so there is no overlap between the train and test sets. This yields 4313 training videos, and 2124 test videos.

2. Real RGB-D videos of urban scenes, from the KITTI dataset [15].

This data was collected with a sensor platform mounted on a moving vehicle, with a human driver navigating through a variety of areas in Germany. We use the “left” color camera, and LiDAR sweeps synced to the RGB images. For training, we use the “odometry” subset of KITTI; it includes egomotion information accurate to within 10cm. The odometry data includes ten sequences, totalling 23201 frames.

We test our model in the validation set of the “tracking” subset of KITTI, which has twenty labelled sequences, totalling 8008 frames. For supervised baselines, we split this data into 12 training sequences and 8 test sequences. For evaluation, we create 8-frame subsequences of this data, in which a target object has a valid label for all eight frames. This subsequencing is necessary since objects are only labelled when they are within image bounds. The egomotion information in the “tracking” data is only approximate.

We evaluate our model on its ability to track objects in 3D. On the zeroth frame, we receive the 3D box of an object to track. On each subsequent frame,

we estimate the object’s new 3D box, and measure the intersection over union (IOU) of the estimated box with the ground truth box. We report IOUs as a function of timesteps.

4.1 Baselines

We evaluate the following baselines. We provide additional implementation details for each baseline (and for our own model) in the supplementary file.

- **Unsupervised 3D flow** [17]. This model uses an unsupervised architecture similar to ours, but with a 2-stage training procedure, in which features are learned first from static scenes and frozen, then a 3D flow module is learned over these features in dynamic scenes. We extend this into a 3D tracker by “chaining” the flows across time, and by converting the trajectories into rigid motions via RANSAC.
- **2.5D dense object nets** [11]. This model learns to map input images into dense 2D feature maps, and uses a contrastive objective at known correspondences across views. We train this model using static points for correspondence labels (like our own model). We extend this model into a 3D tracker by “unprojecting” the learned embeddings into sparse 3D scene maps, then applying the *same tracking pipeline* as our own model.
- **2.5D tracking by colorization** [46,22]. This model learns to map input images into dense 2D feature maps, using an an RGB reconstruction objective. The model trains as follows: given two RGB frames, the model computes a feature map for each frame; for each pixel of the first frame’s feature map, we compute that feature’s similarity with all features of the second frame, and then use that similarity matrix to take a weighted combination of the second frame’s colors; this color combination at every pixel is used as the reconstruction of the first frame, which yields an error signal for learning. We extend this model into a 3D tracker in the same way that we extended the “dense object nets” baseline.
- **3D neural mapping with random features**. This model is equivalent to our proposed model but with randomly-initialized network parameters. This model may be expected to perform at better-than-chance levels due to the power of random features [35] and due to the domain knowledge encoded in the architecture design.
- **3D fully convolutional siamese tracker (supervised)** [3]. This is a straightforward 3D upgrade of a fully convolutional 2D siamese tracker, which uses the object’s feature map as a cross correlation template, and tracks the object by taking an argmax of the correlation heatmap at each step. It is necessary to supervise this model with ground-truth box trajectories. We also evaluate a “cosine windowing” variant of this model, which suppresses correlations far from the search region’s centroid [3].
- **3D siamese tracker with random features**. This model is equivalent to the 3D siamese supervised tracker, but with randomly-initialized network parameters. Similar to the random version of 3D neural mapping, this model

measures whether random features and the implicit biases are sufficient to track in this domain.

- **Zero motion.** This baseline simply uses the input box as its final estimate for every time step. This baseline provides a measure for how quickly objects tend to move in the data, and serves as a lower bound for performance.

All of these are fairly simple trackers. A more sophisticated approach might incrementally update the object template [25], but we leave that for future work.

Ablations. We compare our own model against the following ablated versions. First, we consider a model without search regions, which attempts to find correspondences for each object point in the *entire* 3D map at test time. This model is far more computationally expensive, since it requires taking the dot product of each object feature with the entire scene. Second, we consider a similar “no search region” model but at half resolution, which brings the computation into the range of the proposed model. This ablation is intended to reveal the effect of resolution on accuracy. Third, we omit the “static point selection” (via the function g). This is intended to evaluate how correspondence errors caused by moving objects (violating the static scene assumption) can weaken the model.

4.2 Quantitative results

We evaluate the mean 3D IOU of our trackers over time. Figure 4-left shows the results of this evaluation in CARLA. As might be expected, the supervised 3D trackers perform best, and cosine windowing improves results.

Our model outperforms all other unsupervised models, and nearly matches the supervised performance. The 2.5D dense object net performs well also, but its accuracy is likely hurt by the fact that it is limited exclusively to points observed in the depth map. Our own model, in contrast, can match against both observed and unobserved (i.e., hallucinated or inpainted) 3D scene features. The colorization model performs under the 2.5D dense object net approach, likely because this model only indirectly encourages correspondence via the colorization task, and therefore is a weaker supervision than the multi-view correspondence objectives used in the other methods.

Random features perform worse than the zero-motion baseline, both with a neural mapping architecture and a siamese tracking architecture. Inspecting the results qualitatively, it appears that these models quickly propagate the 3D box off of the object and onto other scene elements. This suggests that random features and the domain knowledge encoded in these architectures are not enough to yield 3D trackers in this data.

We perform the same evaluation in KITTI, and show the results in Figure 4-right. On this benchmark, accuracies are lower for all models, indicating that the task here is more challenging. This is likely related to the fact that (1) the egomotion is imperfect in this data, and (2) the tracking targets are frequently farther away than they are in CARLA. Nonetheless, the ranking of methods is the same, with our 3D neural mapping model performing best. One difference is




Fig. 4. Single-object tracking accuracy, in mean IOU across time, in CARLA (left) and KITTI (right). Methods are sorted in the legend in decreasing order of mean IOU.

that cosine windowing actually worsens siamese tracking results in KITTI. This is likely related to the fact that egomotion stabilization is imperfect in KITTI: a zero-motion prior is only helpful in frames where the target is stationary *and* camera motion is perfectly accounted for; otherwise it is detrimental.

We additionally split the evaluation on stationary vs. moving objects in CARLA. We use a threshold of $1m$ total distance (in world coordinates) across the 8-frame trajectories to split these categories. At the last timestep, the mean 3D IOU for all objects together it is 0.61 (as shown in Figure 4-left); for static objects only, the value is 0.64; for moving objects only, it is 0.56. This suggests that the model tracks stationary objects more accurately than moving objects, likely because their appearance changes less with respect to the camera and the background.

Finally, we evaluate the top two models in CARLA using the standard 2D tracking metrics, multi-object tracking accuracy (MOTA) and multi-object tracking precision (MOTP) [2], though we note that our task only has one target object per video. We find that the 3D siamese + cosine model (supervised) achieves a MOTA of 0.9578 and MOTP of 0.7407, while our model achieves MOTA 0.9125 and MOTP 0.8384. This suggests that the supervised tracker makes fewer misses, but our method delivers slightly better precision.

4.3 Qualitative results

We visualize our tracking results, along with inputs and colorized visualizations of our neural 3D scene maps, in Figure 2. To visualize the neural 3D maps, we take a mean along the vertical axis of the grid (yielding a “bird’s eye view” 2D grid), compress the deep features in each grid cell to 3 channels via principal component analysis, normalize, and treat these values as RGB intensities. Comparing the 3D features learned in CARLA vs those learned in KITTI reveals a very obvious difference: the KITTI features appear blurred and imprecise in comparison with the CARLA features. This is likely due to the imperfect ego-

motion information, which leads to slightly inaccurate correspondence data at training time (i.e., occasional failures by the static point selector g).

In Figure 3, we visualize the features learned by Dense Object Nets in this data. From the PCA colorization it appears that objects are typically colored differently from their surroundings, which is encouraging, but the embeddings are not as clear as those in the original work [11], likely because this domain does not have the benefit of background image subtraction and object-centric losses.

In the supplementary file, we include video visualizations of the learned features and 3D tracking results.

4.4 Ablations

We evaluate ablated versions of our model, to reveal the effect of (1) search regions, (2) resolution, and (3) static point selection in dynamic scenes. Results are summarized in Table 1.

Without search regions, the accuracy of our model drops by 20 IOU points, which is a strong impact. We believe this drop in performance comes from the fact that search regions take advantage of 3D scene constancy, by reducing spurious correspondences in far-away regions of the scene.

Resolution seems to have a strong effect as well: halving the resolution of the wide-search model reduces its performance by 15 IOU points. This may be related to the fact that fewer points are then available for RANSAC to find a robust estimate of the object’s rigid motion.

Since static point selection is only relevant in data with moving objects, we perform this experiment in KITTI (as opposed to CARLA, where the training domain is all static by design). The results show that performance degrades substantially without the static point selection. This result is to be expected, since this ablation causes erroneous correspondences to enter the training objective, and thereby weakens the utility of the self-supervision.

Table 1. Ablations of our model.

Method	Mean IOU
Ours in CARLA	0.61
... without search regions	0.40
... without search regions, at half resolution	0.25
Ours in KITTI	0.46
... without static point selection	0.39

4.5 Limitations

The proposed model has three important limitations. First, our work assumes access to RGB-D data with accurate egomotion data at training time, with a wide variety of viewpoints. This is easy to obtain in simulators, but real-world data of this sort typically lies along straight trajectories (as it does in KITTI), which limits the richness of the data. Second, our model architecture requires a lot of GPU memory, due to its third spatial dimension. This severely limits either the resolution or the metric span of the latent map \mathcal{M} . On 12G Titan X GPUs we encode a space sized $32m \times 4m \times 32m$ at a resolution of $128 \times 32 \times 128$, with a batch size of 4; iteration time is 0.2s/iter. Sparsifying our feature grid, or using points instead of voxels, are clear areas for future work. Third, our test-time tracking algorithm makes two strong assumptions: (1) a tight box is provided in the zeroth frame, and (2) the object is rigid. For non-rigid objects, merely propagating the box with the RANSAC solution would be insufficient, but the voxel-based correspondences might still be helpful.

5 Conclusion

We propose a model which learns to track objects in dynamic scenes just from observing static scenes. We show that a multi-view contrastive loss allows us to learn rich visual representations that are correspondable not only across views, but across time. We demonstrate the robustness of the learned representation by benchmarking the learned features on a tracking task in real and simulated data. Our approach outperforms prior unsupervised 2D and 2.5D trackers, and approaches the accuracy of supervised trackers. Our 3D representation benefits from denser correspondence fields than 2.5D methods, and is invariant to the artifacts of camera projection, such as apparent scale changes of objects. Our approach opens new avenues for learning trackers in arbitrary environments, without requiring explicit tracking supervision: if we can obtain an accurate pointcloud reconstruction of an environment, then we can learn a tracker for that environment too.

Acknowledgements

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. We also acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), AiDTR, the DARPA Machine Common Sense program, and the AWS Cloud Credits for Research program.

References

1. Agrawal, P., Carreira, J., Malik, J.: Learning to see by moving. In: ICCV (2015)
2. Bernardin, K., Elbs, A., Stiefelhagen, R.: Multiple object tracking performance metrics and evaluation in a smart room environment. In: Sixth IEEE International Workshop on Visual Surveillance, in conjunction with ECCV. vol. 90, p. 91. Cite-seer (2006)
3. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: European conference on computer vision. pp. 850–865. Springer (2016)
4. Brodski, A., Paasch, G.F., Helbling, S., Wibral, M.: The faces of predictive coding. *Journal of Neuroscience* **35**(24), 8997–9006 (2015)
5. Brox, T., Malik, J.: Object segmentation by long term analysis of point trajectories. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV. pp. 282–295 (2010)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint arXiv:2002.05709 (2020)
7. Cheriyyadat, A., Radke, R.J.: Non-negative matrix factorization of partial track data for motion segmentation. In: ICCV (2009)
8. Costeira, J., Kanade, T.: A multi-body factorization method for motion analysis. ICCV (1995)
9. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: CORL. pp. 1–16 (2017)
10. Eslami, S.M.A., Jimenez Rezende, D., Besse, F., Viola, F., Morcos, A.S., Garnelo, M., Ruderman, A., Rusu, A.A., Danihelka, I., Gregor, K., Reichert, D.P., Buesing, L., Weber, T., Vinyals, O., Rosenbaum, D., Rabinowitz, N., King, H., Hillier, C., Botvinick, M., Wierstra, D., Kavukcuoglu, K., Hassabis, D.: Neural scene representation and rendering. *Science* **360**(6394), 1204–1210 (2018). <https://doi.org/10.1126/science.aar6170>
11. Florence, P.R., Manuelli, L., Tedrake, R.: Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. In: CoRL (2018)
12. Fragkiadaki, K., Shi, J.: Exploiting motion and topology for segmenting and tracking under entanglement. In: CVPR (2011)
13. Franconeri, S.L., Simons, D.J.: Moving and looming stimuli capture attention. *Perception & psychophysics* **65**(7), 999–1010 (2003)
14. Friston, K.: Learning and inference in the brain. *Neural Networks* **16**(9), 1325–1352 (2003)
15. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)* (2013)
16. Gibson, J.J.: *The Ecological Approach to Visual Perception*. Houghton Mifflin (1979)
17. Harley, A.W., Lakshmikanth, S.K., Li, F., Zhou, X., Tung, H.Y.F., Fragkiadaki, K.: Learning from unlabelled videos using contrastive predictive neural 3d mapping. In: ICLR (2020)
18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
19. Jayaraman, D., Grauman, K.: Learning image representations tied to ego-motion. In: ICCV (2015)
20. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. In: NIPS (2017)
21. Kato, H., Ushiku, Y., Harada, T.: Neural 3d mesh renderer. In: CVPR (2018)

22. Lai, Z., Lu, E., Xie, W.: MAST: A memory-augmented self-supervised tracker. In: CVPR (2020)
23. Lee, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Unsupervised representation learning by sorting sequences. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 667–676 (2017)
24. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM Trans. Graph.* **34**(6), 248:1–248:16 (Oct 2015). <https://doi.org/10.1145/2816795.2818013>, <http://doi.acm.org/10.1145/2816795.2818013>
25. Matthews, L., Ishikawa, T., Baker, S.: The template update problem. *IEEE transactions on pattern analysis and machine intelligence* **26**(6), 810–815 (2004)
26. McClelland, J.L., Rumelhart, D.E.: An interactive activation model of context effects in letter perception: I. an account of basic findings. *Psychological review* **88**(5), 375 (1981)
27. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: CVPR (2015)
28. Misra, I., Zitnick, C.L., Hebert, M.: Unsupervised learning using sequential verification for action recognition. In: ECCV (2016)
29. Ochs, P., Brox, T.: Object segmentation in video: A hierarchical variational approach for turning point trajectories into dense regions. In: ICCV (2011)
30. Olshausen, B.: Perception as an inference problem. In: *The Cognitive Neurosciences*. MIT Press (2013)
31. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv:1807.03748 (2018)
32. Patla, A.E.: Visual control of human locomotion. In: *Advances in psychology*, vol. 78, pp. 55–97. Elsevier (1991)
33. Pinto, Y., van Gaal, S., de Lange, F.P., Lamme, V.A., Seth, A.K.: Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision* **15**(8), 13–13 (2015)
34. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv:1704.00675 (2017)
35. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: *Advances in neural information processing systems*. pp. 1177–1184 (2008)
36. Rao, R.P., Ballard, D.H.: Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* **2**(1), 79 (1999)
37. Roberts, L.: Machine perception of three-dimensional solids. Ph.D. thesis, MIT (1965)
38. Schultz, W., Dayan, P., Montague, P.R.: A neural substrate of prediction and reward. *Science* **275**(5306), 1593–1599 (1997)
39. Sohn, K.: Improved deep metric learning with multi-class N-pair loss objective. In: NIPS. pp. 1857–1865 (2016)
40. Tatarchenko, M., Dosovitskiy, A., Brox, T.: Single-view to multi-view: Reconstructing unseen views with a convolutional network. In: ECCV (2016)
41. Tomasi, C., Kanade, T.: Shape and motion from image streams under orthography: A factorization method. *Int. J. Comput. Vision* **9**(2), 137–154 (Nov 1992). <https://doi.org/10.1007/BF00129684>, <http://dx.doi.org/10.1007/BF00129684>
42. Tulsiani, S., Zhou, T., Efros, A.A., Malik, J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: CVPR (2017)
43. Tung, H.Y.F., Cheng, R., Fragkiadaki, K.: Learning spatial common sense with geometry-aware recurrent networks. In: CVPR (2019)

44. Tung, H.F., Harley, A.W., Seto, W., Fragkiadaki, K.: Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation with unpaired supervision. *ICCV* (2017)
45. Vijayanarasimhan, S., Ricco, S., Schmid, C., Sukthankar, R., Fragkiadaki, K.: Sfmnet: Learning of structure and motion from video. arXiv:1704.07804 (2017)
46. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 391–408 (2018)
47. Walker, J., Doersch, C., Gupta, A., Hebert, M.: An uncertain future: Forecasting from static images using variational autoencoders. In: *ECCV* (2016)
48. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: *ICCV* (2015)
49. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: *CVPR* (2019)
50. Wiskott, L., Sejnowski, T.J.: Slow feature analysis: Unsupervised learning of invariances. *Neural computation* **14**(4), 715–770 (2002)
51. Wu, J., Xue, T., Lim, J.J., Tian, Y., Tenenbaum, J.B., Torralba, A., Freeman, W.T.: Single image 3d interpreter network. In: *ECCV*. pp. 365–382 (2016)
52. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: *CVPR*. pp. 1912–1920. *IEEE Computer Society* (2015)
53. Yuille, A., Kersten, D.: Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences* **10**, 301–308 (2006)
54. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *CVPR* (2017)