

CoCoNets: Continuous Contrastive 3D Scene Representations

Shamit Lal*, Mihir Prabhudesai*, Ishita Mediratta, Adam W. Harley, Katerina Fragkiadaki
Carnegie Mellon University

{shamitl, mprabhud, imedirat, aharley, katef} @cs.cmu.edu

Abstract

This paper explores self-supervised learning of amodal 3D feature representations from RGB and RGB-D posed images and videos, agnostic to object and scene semantic content, and evaluates the resulting scene representations in the downstream tasks of visual correspondence, object tracking, and object detection. The model infers a latent 3D representation of the scene in the form of 3D feature points, where each continuous world 3D point is mapped to its corresponding feature vector. The model is trained for contrastive view prediction by rendering 3D feature clouds in queried viewpoints and matching against the 3D feature point cloud predicted from the query view. Notably, the representation can be queried for any 3D location, even if it is not visible from the input view. Our model brings together three powerful ideas of recent exciting research work: 3D feature grids as a neural bottleneck for view prediction, implicit functions for handling resolution limitations of 3D grids, and contrastive learning for unsupervised training of feature representations. We show the resulting 3D visual feature representations effectively scale across objects and scenes, imagine information occluded or missing from the input viewpoints, track objects over time, align semantically related objects in 3D, and improve 3D object detection. We outperform many existing state-of-the-art methods for 3D feature learning and view prediction, which are either limited by 3D grid spatial resolution, do not attempt to build amodal 3D representations, or do not handle combinatorial scene variability due to their non-convolutional bottlenecks.

1. Introduction

Understanding the three-dimensional structure of objects and scenes may be a key for success of machine perception and control in object detection, tracking, manipulation and navigation. Exciting recent works have explored learn-

ing representations of objects and scenes from multiview imagery and capture the three-dimensional scene structure implicitly or explicitly with 3D binary or feature grids [46, 45, 40], 3D point feature clouds [50], implicit functions that map continuous world coordinates to 3D point occupancy [4, 9, 41, 25, 31, 3], as well as 1D or 2D feature maps [6]. These methods typically evaluate the accuracy of the inferred 3D scene occupancy [4, 30, 46, 25, 31, 3] and the fidelity of image views rendered from the 3D representation [22, 26, 6, 50, 41], as opposed to the suitability of representations for downstream semantic tasks. Methods that indeed focus on rendering photo-realistic images often give up on cross-scene generalization [26, 36], or focus on single-object scenes [41]. Methods that instead focus on learning semantically relevant scene representations are expected to generalize across scenes, and handle multi-object scenes. In the 2D image space, contrastive predictive coding has shown to generate state-of-the-art visual features for correspondence and recognition [49, 12], but does not encode 3D scene structure. In 3D voxel feature learning methods [14, 13], convolutional latent 3D feature grids encode the 3D structure and a view contrastive objective learns semantically useful 3D representations, but the grid resolution limits the discriminability of the features learnt. Recent exciting works combine 3D voxel grids and implicit functions and learn to predict 3D scene and object 3D occupancy from a single view with unlimited spatial resolution [32, 33]. The model proposed in this work brings together these two powerful ideas: 3D feature grids as a 3D-informed neural bottleneck for contrastive view prediction [14], and implicit functions for handling the resolution limitations of 3D grids [32].

We propose Continuous Contrastive 3D Networks (CoCoNets), a model that learns to map RGB-D images to infinite-resolution 3D scene feature representations by contrastively predicting views, in an object and scene agnostic way. Our model is trained to predict views of static scenes given 2.5D (color and depth; RGB-D) video streams as input, and is evaluated on its ability to detect and recognize objects in 3D. CoCoNets map the 2.5D input streams into 3D feature grids of the depicted scene. Given a target view

*Equal contribution

Project page: https://mihirp1998.github.io/project_pages/coconets/

and its viewpoint, the model first warps its inferred 3D feature map from the input view to a target view, then queries point features using their continuous coordinates, and pulls these features closer to the point features extracted from the target view at the same 3D locations (Figure 1). We use a contrastive loss to measure the matching error, and back-propagate gradients end-to-end to our differentiable modular architecture. At test time, our model forms plausible 3D completions of the scene given a *single RGB-D image* as input: it learns to fill in information behind occlusions, and infer the 3D extents of objects.

We demonstrate the advantages of combining 3D neural bottleneck, implicit functions and contrastive learning for 3D representation learning by comparing our model against state-of-the-art self-supervised models, such as i) contrastive learning for pointclouds [51], which shares a similar loss but not the amodal predictive ability of our model, ii) contrastive neural mapping [13], which can amodally inpaint a 3D discrete feature grid but suffers from limited spatial grid resolution, and iii) dense ObjectNets [7], which self-learns 2D (instead of 3D) feature representations with a triangulation-driven supervision similar to (i). Our experimental results can be summarized as follows: **(1)** 3D object tracking and re-identification (Figure 3): We show that scene representations learnt by CoCoNets can detect objects in 3D across large frame gaps better than the baselines [51, 13, 7]. **(2)** Supervised 3D object detection: Using the learnt 3D point features as initialization boosts the performance of the state-of-the-art Deep Hough Voting detector of [34]. **(3)** 3D cross-view and cross-scene object 3D alignment: We show that the learnt 3D feature representations can infer 6DoF alignment between the same object in different viewpoints, and across different objects of the same category, better than [51, 13, 7]. We further show that our model can predict image views (with or without depth as input) and 3D occupancies that outperform or are on par with the state-of-the-art view and occupancy prediction models [33, 47, 6].

In summary, the main contribution of this paper is a model that learns infinite-resolution 3D scene representations from RGB-D posed images, useful for tracking and corresponding objects in 3D, pre-training 3D object detectors, and predicting views and 3D occupancies. We set a new state-of-the-art in self-supervision of 3D feature representations.

2. Related work

Learning to 3D reconstruct objects and scenes Learning to infer 3D reconstructions of objects and scenes from single images or videos has been the goal of recent deep geometrical methods, that have explored a variety of explicit 3D representations, such as 3D point clouds [29, 23], 3D binary voxel occupancies [46, 52], or 3D meshes [18, 10].

Since detailed 3D supervision is only possible in large scale in simulation, many approaches attempt supervision from RGB or depth map prediction through differentiable rendering of the inferred 3D reconstruction [18, 10]. To handle limitations of spatial resolution of 3D voxel grids, recent approaches represent 3D object occupancy with implicit functions parameterized by deep neural networks trained to map continuous world 3D coordinates to the corresponding point 3D occupancy values [4, 25, 31, 3]. While these methods train one 3D shape implicit function per object, by assigning a 1D latent embedding to each object, some works [33, 32] train 3D grids of functions where each function takes care of estimating the occupancy of points in the vicinity of the corresponding 3D voxel centroid. These latter methods enjoy the generalization of 3D convolutions and can scale to multi-object scenes, while implicit functions parametrized by non-convolutional, fully connected networks are mostly limited to single object scenes [4]. Our approach also employs 3D grids of functions which predict feature embeddings for the corresponding continuous 3D world coordinates as opposed to merely occupancy.

Neural image synthesis with 3D inductive biases Deep image generative networks have shown compelling results in generating photorealistic images that match the image statistics of the unlabelled image collections they are trained on [11]. They are based on variational autoencoders [20], generative adversarial networks [1], generative flows [19], or autoregressive image pixel generators [48]. However, these 2D generative models learn to parameterize the manifold of 2D natural images, and struggle to generate images that are multi-view consistent, since the underlying 3D scene structure cannot be exploited. Generative models trained from multi-view data can render arbitrary views of an input scene [43, 17, 40, 42, 37]. Such multi-view generative models often restrict themselves to single-object scenes [41, 28, 39] or to a single complex scene without aiming at cross-scene generalization [26, 36], with the goal of generating high fidelity photorealistic images, replacing hand-engineered graphics engines. Their architectures incorporate many inductive biases of graphics engines, such as 3D-to-2D rendering modules [43] and explicit feature transformations to handle viewpoint changes [47]. Their lack of cross-scene generalization or their limitation to single object scenes makes it hard to adopt their inferred feature representations for visual recognition.

Learning visual feature representations by self-supervised view prediction Recent methods learn neural scene representations by predicting views of a scene under known egomotion [6, 41, 14, 40]. View prediction, as opposed to the related and very effective objective of feature learning via triangulation [51], results in explicitly

or implicitly *amodal* representations, i.e., representations that can predict information missing from the input observations [27], as opposed to simply featurizing the visible image pixels [7]. Different view prediction methods for learning representations vary with respect to the amount of their reasoning regarding geometry and the underlying 3D structure of the scene [44]. The generative query network (GQN) of Eslami *et al.* [6] showed that it can predict alternative views of toy simulated scenes without explicit 3D structure inference, and demonstrated the usefulness of the inferred representations as pre-training for reinforcement behaviour learning tasks. Geometry-aware recurrent networks of Tung *et al.* [47] use a latent 3D feature map in their bottleneck for view prediction and demonstrate superior generalization, granted from the 3D convolutional inductive bias. Harley *et al.* [14] uses a similar 3D latent feature grid but optimizes for contrastive prediction as opposed to RGB regression, and demonstrate its usefulness for 3D object tracking and 3D moving object segmentation. Our work optimizes a contrastive view prediction objective similar to [14] but uses a 3D grid of implicit functions as its latent bottleneck. We empirically show that the emergent 3D feature representations are more accurate in 3D object tracking and visual correspondence than the features obtained from existing state-of-the-art 3D feature learning methods.

3. Continuous Contrastive 3D Networks (CoCoNets) for Learning Amodal Visual Representations

We consider a mobile agent that can move about the scene and has access to its egomotion. The agent has a color camera with known intrinsics, and a depth sensor registered to the camera’s coordinate frame. We use groundtruth depth provided by the simulation environment, and we will show in Sec. 4 that the learned models generalize to the real world, where (sparser) depth is provided by a LiDAR unit.

CoCoNets learn 3D visual feature representations by collecting posed images in static scenes and doing contrastive view prediction. We describe the architecture in Sec. 3.1. We then evaluate the correspondability of the resulting 3D feature representations in 3D object re-identification and tracking in dynamic scenes (Section 4.1), as pre-training for 3D object detection (Section 4.2), and cross-object semantic visual correspondence (Section 4.3).

3.1. Continuous Contrastive 3D Networks (CoCoNets)

Our model’s architecture is illustrated in Figure 1. It is a neural network with a three-dimensional neural bottleneck $\mathbf{M} \in \mathbb{R}^{w \times h \times d \times c}$, which has three spatial dimensions (width w , height h , and depth d) and a feature dimension (c

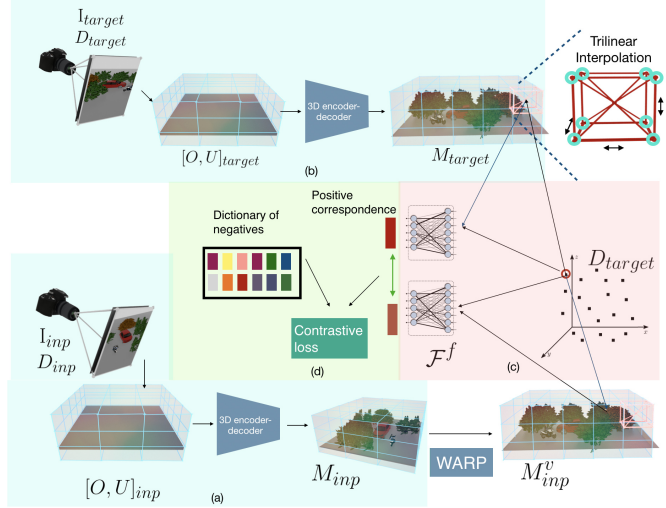


Figure 1: **Continuous Contrastive 3D Networks (CoCoNets)** are trained to lift 2.5D images to 3D feature function grids of the scene by optimizing for view-contrastive prediction. (a) In the top-down path, the model encodes RGB-D images into a 3D feature map $\mathbf{M} \in \mathbb{R}^{w \times h \times d \times c}$, and uses explicit 3D feature transformations (translation and 3D rotation) to account for changes of viewpoint between the input and target views. (b) In the bottom-up path, we encode the RGB-D of the target viewpoint into a 3D feature cloud. (c) Given continuous 3D world coordinates (X, Y, Z) and its embedded code $v_{(X, Y, Z)}$ inferred via trilinear interpolation, a fully connected network maps the coordinates and the embedded code, to the feature vector of the 3D point at location (X, Y, Z) . (d) Metric learning losses in 3D tie the two point cloud representations together.

channels per grid location).

The latent state aims to capture an informative and geometrically-consistent 3D deep feature map of the world space. Therefore, the spatial extent corresponds to a large cuboid of world space, defined with respect to the camera’s position at the first timestep. Each voxel in the 3D feature map \mathbf{M} corresponds to a cuboid in the 3D scene depicted in the RGB-D image.

To be able to generate features at infinite spatial resolution, i.e., to featurize continuous 3D physical points within a voxel, we use implicit function parametrization. Our model’s architecture of interpolating within a voxel grid resembles that of Peng *et al.* [32]. Let (X, Y, Z) denote the continuous world coordinate of a 3D point whose feature we wish to infer. First, we trilinearly interpolate the feature grid to obtain a c -dimensional feature vector at point (X, Y, Z) . Denoting this trilinearly interpolated feature vector as p for input point (X, Y, Z) , we further obtain a

3d location conditioned feature vector using $\phi(p, x)$, where ϕ is a small fully-connected network. Finally we obtain our 32-dimensional embedding vector using $f_\theta(p, \phi(p, x))$, where f_θ represents a multi-block fully connected ResNet. Further details about the network architecture can be found in the supplementary.

A similar fully-connected ResNet parametrization provides a point’s binary occupancy $o_{(X,Y,Z)}$, and its RGB color value $c_{(X,Y,Z)}$. The three ResNets that predict features, occupancy and RGB values of 3D points do not share weights. During training, we use the point cloud of the target viewpoint to query our model for features, occupancies and colors in the corresponding visible 3D point locations and propagate gradients in an end-to-end manner to the same feature voxel 3D map \mathbf{M} . We denote the operation of obtaining point features, occupancies and colors by querying the feature map at continuous locations by $\mathcal{F}^f(\mathbf{M}, (X, Y, Z))$, $\mathcal{F}^o(\mathbf{M}, (X, Y, Z))$, and $\mathcal{F}^c(\mathbf{M}, (X, Y, Z))$. At test time, we query the model in both visible and non visible 3D point locations, to obtain an amodal completed 3D point feature cloud.

CoCoNets is made up of differentiable modules that go back and forth between 3D feature space and 2D image space. It can take as input a variable number of RGB-D images both at training time and test time. For simplicity, in our experiments we encode only a single view input at both train and test time. More details on each of the following modules are included in the supplementary file.

2D-to-3D unprojection (Figure 1 (a, b)) This module converts the input RGB image $\mathbf{I} \in \mathbb{R}^{w \times h \times 3}$ and depth map $D \in \mathbb{R}^{n \times 3}$ where, n is the size of the pointcloud) into 3D tensors using available camera intrinsics. The RGB is “unprojected” into a 3D tensor $\mathbf{U} \in \mathbb{R}^{w \times h \times d \times 3}$ by filling each 3D grid location with the RGB value of its corresponding subpixel. The pointcloud is converted to a 3D occupancy grid $\mathbf{O} \in \mathbb{R}^{w \times h \times d \times 1}$, by assigning each voxel a value of 1 or 0, depending on whether or not a point lands in the voxel. We then convert the concatenation of these tensors into a 3D feature tensor $\mathbf{M} \in \mathbb{R}^{w \times h \times d \times c}$, via a 3D convolutional encoder-decoder network with skip connections. We L_2 -normalize the feature in each grid cell.

3D-to-2D projection This module first warps the 3D feature map \mathbf{M} to align it to the target viewpoint, yielding \mathbf{M}^V . It then generates a 3D feature point cloud by querying the model at the 3D point locations visible from the target viewpoint, provided by the target depth map D_{target} . It also generates a 2D feature map by computing for each visible 3D point the corresponding 2D pixel location using the camera intrinsics, $(x, y) = f_{\frac{X}{Z}}, f_{\frac{Y}{Z}}$, and copying the 3D point feature into that location in the 2D feature map. Note that each 3D feature point is mapped independently to 2D, in contrast to neural renderers that convolutionally map feature maps to image views [40, 6]. This independent

point-by-point rendering is closer to the spirit of graphics operations [41].

3D contrastive learning of point features Given a pair of input RGB images $(\mathbf{I}_{inp}, \mathbf{I}_{target})$, depth maps (D_{inp}, D_{target}) , and the camera pose change between input and target views \mathbf{V} , we consider two types of representations for the target view:

- a top-down one, \mathbf{M}_{inp}^V (Figure 1 (a)) by encoding the input RGB-D image $(\mathbf{I}_{inp}, D_{inp})$ and orienting the feature map to the target viewpoint, and predicting the features for the target 3D points in D_{target} by querying the functions in \mathbf{M}_{inp}^V , obtaining the feature cloud $\{(X, Y, Z, \mathcal{F}(\mathbf{M}_{inp}^V, (X, Y, Z)))\}$ for $(X, Y, Z) \in D_{target}$ (Figure 1 (c)).
- a bottom-up one, \mathbf{M}_{target} (Figure 1 (b)) by simply encoding the target RGB-D image $\mathbf{I}_{target}, D_{target}$ and predicting the features for the target 3D points in D_{target} , obtaining the feature cloud $\{(X, Y, Z, \mathcal{F}(\mathbf{M}_{target}, (X, Y, Z)))\}$ for $(X, Y, Z) \in D_{target}$ (Figure 1 (c)).

We use a contrastive InfoNCE loss [49] (Figure 1 (d)) to pull corresponding top-down and bottom-up point features close together in embedding space and push non-corresponding ones farther away:

$$\mathcal{L} = -\log \frac{\exp(\frac{q \cdot k_+}{\tau})}{\sum_{i=0}^K \exp \frac{q \cdot k_i}{\tau}}, \quad (1)$$

where τ is a temperature hyper-parameter and the sum in the denominator is over one positive and K negative samples. In the numerator, q represents $\mathcal{F}^f(\mathbf{M}_{inp}, (i, j, k))$, and k_+ is its corresponding positive sample $\mathcal{F}^f(\mathbf{M}_{target}, (m, n, o))$. The dot product computes a similarity. In practice, we randomly sample corresponding points from the top-down or bottom-up feature clouds at each training iteration. We also maintain a large pool of negative pairs in a dictionary, using the approach proposed by He *et al.* [15].

Our metric learning loss, if applied only on **3D points visible from both input and target views**, coincides with the point contrastive metric learning of Xie *et al.* [51], a state-of-the-art 3D point feature learning method. CoCoNets can handle input and target views that actually have few or no points in common. We compare against Xie *et al.*’s state-of-the-art 3D feature learning model in our experiments, and demonstrate the importance of amodal completion for feature learning.

Occupancy prediction and RGB view regression We train a separate CoCoNet model to predict occupancy and RGB in novel viewpoints, which we denote CoCoNets-*OccRGB*. We use a standard binary cross entropy loss

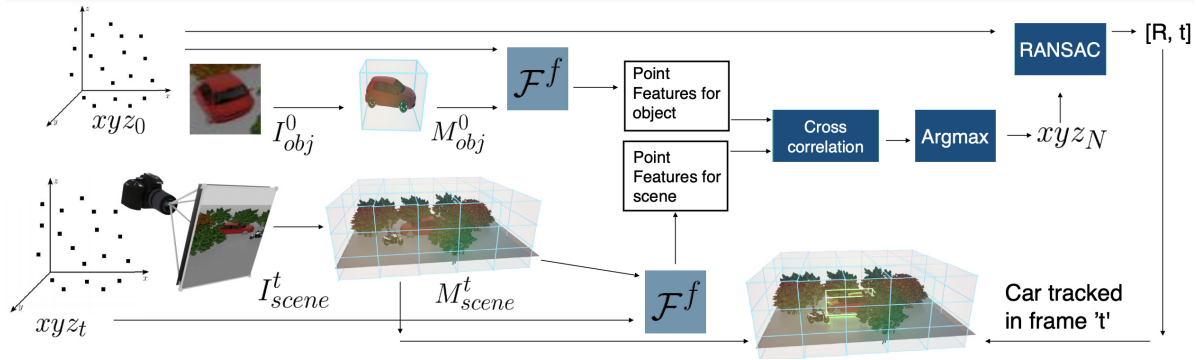


Figure 2: **3D object tracking using CoCoNets.** Given the cropped RGB-D image I_{obj} of the object to track at $t = 0$, our model infers the 3D object feature map M_{obj} , and queries it using xyz_0 , the point cloud of the object, to obtain object point features. Similarly, it obtains the point features of the entire scene at timestep t . Finally, it does cross correlation between these features to get xyz_N , where each i^{th} point in xyz_N is the point from the scene whose feature matched best with the feature for i^{th} point in xyz_0 . We then apply RANSAC on xyz_0 and xyz_N to obtain the location of the car at timestep t .

for the occupancy, and a regression loss for RGB. The model predicts point clouds in target viewpoints and infers their occupancy and color from the implicit function grids $\mathcal{F}^o, \mathcal{F}^c$. Specifically, to predict a target RGB view and its occupancy given a source RGB-D image, after encoding and orienting the 3D feature map M to the target viewpoint, we predict the point occupancies of the target view, and predict the point RGB colors for the occupied points. These are then projected to the image space in the target view.

Jointly training RGB and occupancy prediction with contrastive view prediction did not improve discriminability of our feature representations. The novelty of our paper is in using voxel-implicit architectures for amodal contrastive 3D feature learning as opposed to 3D occupancy prediction [33, 32], or in addition to it.

RGB view regression without depth input Additionally, we also make use of NeRF’s volumetric renderer [26], with features extracted from M_{input} as priors, to render novel views using just a single RGB image (no depth information) as input. We call this $CoCoNets^{NoDepth, RGB}$ and explain the setup in the supplementary file. Qualitative results from this model can be seen in Figure 5.

4. Experiments

Useful visual feature representations are expected to correspond visual entities across variations in pose and appearance, as well as across intra-category variability. We evaluate the correspondability of our 3D scene representations in 3D object tracking and object re-identification, 3D pose estimation and cross-object correspondence. We further evaluate our features as pre-training for 3D object detector for the state-of-the-art method of [34]. Our experiments aim to answer the following questions:

1. Is amodal completion important for learning scene representations useful for visual correspondence, across time and across scenes? To this end, we compare 3D feature representations learned by CoCoNets against contrastive point clouds [51] that do not consider amodal completion but share a similar contrastive objective for 3D points visible in both scenes.
2. Is our 3D function parameterization important for learning discriminative visual representations? To this end, we compare with 3D contrastive neural mapping of Harley *et al.* [14] that does not consider continuous functions, but rather uses discrete 3D feature maps as the neural bottleneck.
3. Does the performance of 3D object detectors in point-clouds improve when using CoCoNets as initialization of the point features, and by how much? To this end, we compare with VoteNet[34], a model that uses PointNet++ [35] as their feature extractor and deep Hough voting for 3D object detection in point clouds.
4. Is the continuous 3D convolutional bottleneck of CoCoNets useful for view prediction? To this end, we compare image views rendered from our model in novel scenes with images from GQN [6], which does not consider 3D neural bottleneck, and GRNN [47], which has a discrete 3D convolutional bottleneck.

We train and test CoCoNets in the simulated datasets of CARLA [5], ShapeNet [2] and test them further—without training—in the real-world KITTI dataset [8]. CARLA is an open-source photorealistic simulator of urban driving scenes, which permits moving the camera to any desired viewpoint in the scene. ShapeNet is a large scale 3D object repository which again permits camera placement at will.

KITTI is an urban driving scene dataset collected from a camera mounted on a car. More information on datasets is available in the supplementary file.

4.1. Self-supervised 3D object tracking and re-identification

Setup We evaluate the 3D feature representations of CoCoNets in their ability to track a 3D object over time, as well as across large frame gaps, a task known as object re-identification. Given the 3D bounding box of an object in the first frame, we are interested in tracking the object in subsequent frames. We first infer 3D features for object points visible in the first frame. For subsequent timesteps, we predict the amodal *complete* 3D feature cloud of the scene from a single view by querying our model in a set of visible and invisible 3D point locations. We found this uniform random sampling of query point set strategy to suffice for our tasks, though more elaborate strategies that focus on querying points close to object surfaces can be used, such as Marching Cubes [24], which has been used by previous 3D implicit shape function models for learning object occupancy to obtain a watertight mesh [9, 32, 33]. Given the object feature point cloud and the amodal scene feature cloud, we obtain a rigid body transformation (3D rotation and translation) by solving the orthogonal Procrustes problem [38], and use RANSAC to find the transformation that satisfies most inliers as shown in Figure 2, where inlier point matches are obtained by thresholding inner products at 0 between the object 3D point features and 3D point features of the target scene. We find the object 3D box in the target scene by warping the box b_0 using the obtained transformation. Note that we do not restrict our model to search locally only for a matching object, but search over the whole scene which makes the task harder. Such object matching across frame gaps is important for re-detecting an object after occlusions. It is also a good test for the discriminability of the features, since they are not assisted by temporal continuity or local search regions.

Evaluation and comparison with baselines We compare CoCoNets against the following baselines (we use names indicative to the differences and characteristics): i) 3DVoxContrast of Harley *et al.* [13], ii) 3DPointContrast of Xie *et al.* [51] iii) 2.5DDenseObjectNets of Florence *et al.* [7], which uses triangulation supervision to train 2D deep feature maps, then lifts them to 3D point feature clouds using the available depth map.

We follow the experimental setup of Harley *et al.* [13], which uses video sequences of 10 timesteps. For 3DVoxContrast, we first use multiview RGB-D data with known egomotion to train a model using 3D contrastive losses. Upon training, 3DVoxContrast maps a single RGB-D image to a discrete 3D feature map. Then, 3DVoxContrast

uses this model for tracking by solving orthogonal procrustes problem with RANSAC between voxel features as opposed to point features. For 3DPointContrast, upon training the model, we compute RANSAC on the cross-correlation scores between point features from the visible 3D points on the object at frame 0 and the *visible* point cloud at each subsequent frame. Our method also considers amodal completed point cloud at each frame. For 2.5DDenseObjectNets, upon training the model, we apply RANSAC much the same way as in the 3DPointContrast model.

We show quantitative tracking results for our model and baselines on the CARLA and KITTI datasets in Table 1. Figure 3 visualizes 3D object trajectories over time from an overhead view. Note that we do not use any locality search, i.e., each model searches across the whole target frame for the configuration of the object. Our model achieves superior performance on 3D object tracking than all the baselines. It dramatically outperforms 3DVoxContrast, which highlights the importance of high spatial resolution for 3D feature learning and correspondence. It also outperforms 3DPointContrast, which highlights the importance of amodal completion supported by our model. Lastly, all 3D methods outperform 2.5DDenseObjectNets which highlights the effectiveness of learning features in a metric 3D feature space.

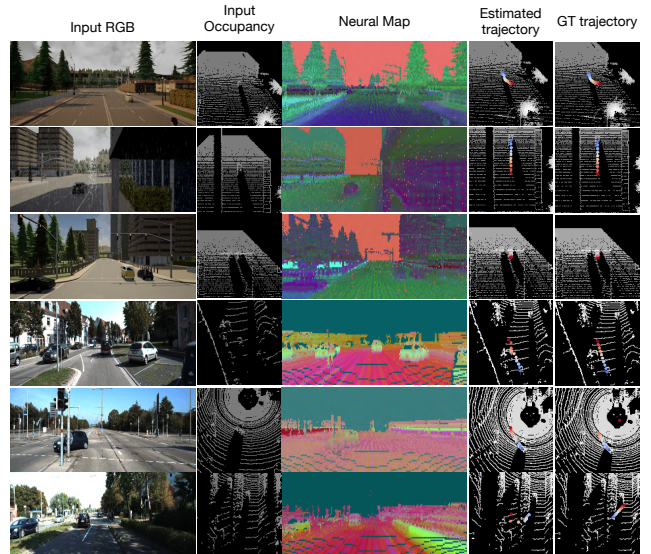


Figure 3: **Self-supervised 3D object tracking.** In the 1st and 2nd column we visualize the RGB and depth from the first frame, which is given as input to our model, along with a 3D box specifying the object to track. In the 3rd column we visualize our inferred point features by projecting them to the same RGB image and then doing PCA compression. In the last 2 columns we show the estimated and ground truth trajectories. The top three rows show our results on CARLA; the bottom three rows show our KITTI results.

Method	CARLA	KITTI
CoCoNets (ours)	0.61	0.54
3DPointContrast [51]	0.55	0.48
3DVoxContrast [13]	0.37	0.23
2.5DDenseObjectNets [7]	0.24	0.19

Table 1: **3D object tracking mean IOU across timesteps.**

Method	0.25	0.30	0.40	0.50
VoteNet [34]	0.32	0.26	0.24	0.20
CoCoNets- <i>VoteNet</i> (Ours)	0.51	0.47	0.41	0.32

Table 2: **Comparing VoteNet with CoCoNets-*VoteNet* on mAP for different IOUs.** Pre-training the 3D point cloud features with CoCoNets significantly boosts performance.

4.2. Supervised 3D object detection in point clouds

Setup We test the 3D feature representations obtained by CoCoNets in their ability to boost the performance of a supervised 3D object detector when used as initialization. We use VoteNet [34], a state-of-the-art 3D object detector that combines the classic ideas of Hough voting in a deep neural network architecture that computes point features, uses them to vote for locations and sizes of the 3D object boxes, and aggregates the votes to propose 3D object bounding boxes with associated confidence, in an end-to-end differentiable framework. We modify VoteNet [34] by replacing their point cloud feature extractor backbone with a trained CoCoNet; we will refer to this as CoCoNets-*VoteNet*. The original backbone in VoteNet uses PointNet++ [35] to compute point features. Both the VoteNet backbone and CoCoNets-*VoteNet* backbone sample 1024 points from the dense point cloud using farthest point sampling. The subsequent voting and proposal modules of VoteNet then operate on these sampled point features to make the final predictions. In both cases, upon initialization, the 3D point features are trained end-to-end supervised for the 3D object detection task. We show below that our pre-training gives a significant boost to 3D object detection performance.

Evaluation and comparison with baselines We evaluate both models on the CARLA dataset in the task of predicting 3D axis aligned car bounding boxes. We show the mean average precision (mAP) scores for both the models for different IoU thresholds in Table 2. CoCoNets-*VoteNet* outperforms VoteNet across all IoU thresholds.

4.3. Cross-view and cross-object 3D alignment

Setup We evaluate the correspondability of our 3D feature representations by testing how well they can align dif-

ferent views of the same instance, and different instances that belong to the same object category. Given the amodal 3D feature point clouds extracted for two objects in random viewpoints, we use orthogonal Procrustes problem with RANSAC to estimate the 3D rigid body transformation that aligns them, similar to tracking a car in Section 4.1. Estimating fine grained correspondences between objects, once their rigid alignment has been estimated, is obtained using iterative closest point method [21], and we show such results in the supplementary file.

Evaluation and comparison with baselines We investigate two different setups: i) Given two viewpoints of the same object, our goal is to estimate the relative 3D rotation between them. ii) Given two objects of the same category in random viewpoints, our goal is to estimate a rigid 3D alignment between them. We compare our model against baselines 3DVoxContrast and 3DPointContrast, for which we use RANSAC on voxel and point features, respectively, to estimate a 3D rigid transformation for the two objects.

Method	cross-object	cross-view
3DVoxContrast [13]	0.18	0.21
3DPointContrast [51]	0.09	0.14
CoCoNets (Ours)	0.18	0.58

Table 3: **Cross-object and cross-view 3D alignment accuracies in Shapenet dataset** (mean over 4 classes: Aeroplane, Mug, Car, Chair).

The inferred 3D alignment is considered correct if each of the yaw, roll, and pitch angles are within 10° of their respective ground truth values. In Table 3 we show the cross-view (same-object) and cross-object 3D alignment accuracy for our model and the baselines. Our model performs the same as 3DVoxContrast on cross-object alignment, and outperforms both models on cross-view alignment. Note that cross-object alignment is much harder than cross-view for all three models.

4.4. Generalization in view prediction

We evaluate CoCoNets-*OccRGB* in its ability to predict plausible images in scenes with novel number of objects, novel object appearance, and novel arrangements. We compare against state-of-the-art view prediction models. We further evaluate its ability for 3D occupancy prediction and completion. Extensive results of occupancy and RGB predictions on the ShapeNet dataset for our model can be found in the supplementary file.

We compare CoCoNets in RGB prediction against generative query network (GQN) [6] and geometry-aware recurrent networks (GRNN) [47]. We adapted GQN code to

use posed RGB-D images as input. GQN uses a 2D latent feature space to encode the RGB-D input image, and GRNN uses a discrete latent 3D feature map of the scene. Our model uses a grid of functions and does not suffer from resolution limitations of the 3D feature map.

CoCoNets predicts RGB images for target viewpoints as described in Sec. 3.1. Figure 4 shows the view generation results for CoCoNets, GQN, and GRNN. For qualitative results on RGB prediction and 3D occupancy prediction, details on the datasets used, architectural details and video results of our self-supervised tracking, please see our supplementary file.




Figure 4: **Novel view prediction.** We compare CoCoNets, GQN [6], and GRNN [47] on the CLEVR [16] dataset.



Figure 5: **RGB view regression** on the ShapeNet dataset using CoCoNets^{NoDepthRGB}. See the supplementary file for details on the experimental setup.

Limitations/Extensions CoCoNets can be extended to operate without the availability of depth maps at either train or test time using a differential rendering module [41], as shown in Figure 5. Our supplementary material further contains results from variations of our model that do not assume depth available at test time, yet use it at training time.

5. Conclusion

We present a method for learning 3D visual feature representations by self-supervised view and depth prediction from posed RGB and RGB-D images. Our networks lift input 2.5D images of objects and scenes into latent three-dimensional function grids, which can be decoded to infinite-resolution 3D occupancy and 3D feature predictions of the object or scene. Our networks are trained by predicting views using a contrastive mutual information maximization objective. We evaluate the emergent 3D visual feature representations in 3D object tracking in dynamic scenes and in cross-view and cross-object alignment estimation. We empirically demonstrate that the features are semantically meaningful and outperform popular point-based supervision [51] which does not consider 3D completion, and discrete voxel 3D latent feature maps of previous works [14, 13] which are limited by the spatial resolution of the 3D feature grid. Moreover, our models can better generalize to novel scenes with unseen number and appearance of objects than networks that do not encode 3D structure [6] or do not incorporate 3D convolutional modules [41]. They make a step towards self-supervised learning of amodal 3D feature representations, which we show are useful for 3D object tracking and correspondence. Avenues for future work include learning such representations directly from dynamic videos, and relaxing the egomotion and depth supervision.

6. Acknowledgement

This work has been funded by Sony AI, DARPA Machine Common Sense, a NSF CAREER award, the Air Force Office of Scientific Research under award number FA9550-20-1-0423. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis, 2019.
- [2] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. *arXiv:1512.03012*, 2015.
- [3] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019.
- [4] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion, 2020.
- [5] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *CORL*, pages 1–16, 2017.
- [6] S. M. Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S. Morcos, Marta Garnelo, Avraham Ruderman, Andrei A. Rusu, Ivo Danihelka, Karol Gregor, David P. Reichert, Lars Buesing, Theophane Weber, Oriol Vinyals, Dan Rosenbaum, Neil Rabinowitz, Helen King, Chloe Hillier, Matt Botvinick, Daan Wierstra, Koray Kavukcuoglu, and Demis Hassabis. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [7] Peter R. Florence, Lucas Manuelli, and Russ Tedrake. Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 373–385. PMLR, 29–31 Oct 2018.
- [8] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [9] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape, 2020.
- [10] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. *CoRR*, abs/1906.02739, 2019.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *NIPS*, pages 2672–2680, 2014.
- [12] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding, 2019.
- [13] Adam W Harley, Shrinidhi K Lakshminanth, Paul Schydlo, and Katerina Fragkiadaki. Tracking emerges by looking around static scenes, with neural 3d mapping. *arXiv preprint arXiv:2008.01295*, 2020.
- [14] Adam W Harley, Fangyu Li, Shrinidhi K Lakshminanth, Xian Zhou, Hsiao-Yu Fish Tung, and Katerina Fragkiadaki. Learning from unlabelled videos using contrastive predictive neural 3d mapping. In *ICLR*, 2020.
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning, 2020.
- [16] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.
- [17] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *NIPS*, 2017.
- [18] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. *CoRR*, abs/1711.07566, 2017.
- [19] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions, 2018.
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013.
- [21] J. Minguetz L. Armesto and L. Montesano. A generalization of the metric-based iterative closest point technique for 3d scan matching. In *IEEE International Conference on Robotics and Automation*, Alaska, USA, 2010.
- [22] Hoang-An Le, Thomas Mensink, Partha Das, and Theo Gevers. Novel view synthesis from single images via point cloud transformation, 2020.
- [23] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3d object reconstruction, 2017.
- [24] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIG-GRAPH Comput. Graph.*, 21(4):163–169, Aug. 1987.
- [25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis, 2020.
- [27] Bence Nanay. The importance of amodal completion in everyday perception. *i-Perception*, 9(4):2041669518788887, 2018. PMID: 30109014.
- [28] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7588–7597, 2019.
- [29] David Novotný, Diane Larlus, and Andrea Vedaldi. Learning 3d object categories by looking around them. *CoRR*, abs/1705.03951, 2017.
- [30] David Novotny, Roman Shapovalov, and Andrea Vedaldi. Canonical 3d deformer maps: Unifying parametric and non-parametric methods for dense weakly-supervised category reconstruction, 2020.
- [31] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [32] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks, 2020.
- [33] Stefan Popov, Pablo Bauszat, and Vittorio Ferrari. Corenet: Coherent 3d scene reconstruction from a single rgb image, 2020.
- [34] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9277–9286, 2019.
- [35] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.
- [36] Gernot Riegler and Vladlen Koltun. Free view synthesis, 2020.
- [37] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. *arXiv:1905.05172*, 2019.
- [38] Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.
- [39] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *arXiv preprint arXiv:2007.02442*, 2020.
- [40] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. DeepVoxels: Learning persistent 3D feature embeddings. In *CVPR*, 2019.
- [41] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations, 2020.
- [42] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Single-view to multi-view: Reconstructing unseen views with a convolutional network. In *ECCV*, 2016.
- [43] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, Rohit Pandey, Sean Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian Theobalt, Maneesh Agrawala, Eli Shechtman, Dan B Goldman, and Michael Zollhöfer. State of the art on neural rendering, 2020.
- [44] Josh Tobin, OpenAI Robotics, and Pieter Abbeel. Geometry-aware neural rendering, 2019.
- [45] Shubham Tulsiani, Saurabh Gupta, David F. Fouhey, Alexei A. Efros, and Jitendra Malik. Factoring shape, pose, and layout from the 2d image of a 3d scene. *CoRR*, abs/1712.01812, 2017.
- [46] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. *CoRR*, abs/1704.06254, 2017.
- [47] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *CVPR*, 2019.
- [48] Aaron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *NIPS*, pages 4790–4798, 2016.
- [49] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- [50] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [51] Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas J. Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding, 2020.
- [52] Xinchen Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision, 2017.