



MIT Open Access Articles

3D Interpreter Networks for Viewer-Centered Wireframe Modeling

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

Citation	Wu, Jiajun, et al., "3D Interpreter Networks for Viewer-Centered Wireframe Modeling." International Journal of Computer Vision 126 (2018): p. 1009-26 doi 10.1007/S11263-018-1074-6 ©2018 Author(s)
As Published	10.1007/S11263-018-1074-6
Publisher	Springer Nature
Version	Author's final manuscript
Citable link	https://hdl.handle.net/1721.1/124536
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

3D Interpreter Networks for Viewer-Centered Wireframe Modeling

Jiajun Wu¹ · Tianfan Xue² · Joseph J. Lim³ · Yuandong Tian⁴ ·
Joshua B. Tenenbaum¹ · Antonio Torralba¹ · William T. Freeman^{1,5}

Received: date / Accepted: date

Abstract Understanding 3D object structure from a single image is an important but challenging task in computer vision, mostly due to the lack of 3D object annotations to real images. Previous research tackled this problem by either searching for a 3D shape that best explains 2D annotations, or training purely on synthetic data with ground truth 3D information.

In this work, we propose 3D INterpreter Networks (3D-INN), an end-to-end trainable framework that sequentially estimates 2D keypoint heatmaps and 3D object skeletons and poses. Our system learns from both 2D-annotated real images and synthetic 3D data. This is made possible mainly by two technical innovations. First, heatmaps of 2D keypoints serve as an intermediate representation to connect real and synthetic data. 3D-INN is trained on real images to estimate 2D keypoint heatmaps from an input image; it then

predicts 3D object structure from heatmaps using knowledge learned from synthetic 3D shapes. By doing so, 3D-INN benefits from the variation and abundance of synthetic 3D objects, without suffering from the domain difference between real and synthesized images, often due to imperfect rendering. Second, we propose a Projection Layer, mapping estimated 3D structure back to 2D. During training, it ensures 3D-INN to predict 3D structure whose projection is consistent with the 2D annotations to real images.

Experiments show that the proposed system performs well on both 2D keypoint estimation and 3D structure recovery. We also demonstrate that the recovered 3D information has wide vision applications, such as image retrieval.

Keywords 3D skeleton · Single image 3D reconstruction · Keypoint estimation · Neural network · Synthetic data

Jiajun Wu
E-mail: jiajunwu@mit.edu

Tianfan Xue
E-mail: tianfan@google.com

Joseph J. Lim
E-mail: limjj@usc.edu

Yuandong Tian
E-mail: yuandong@fb.com

Joshua B. Tenenbaum
E-mail: jbt@mit.edu

Antonio Torralba
E-mail: torralba@csail.mit.edu

William T. Freeman
E-mail: billf@mit.edu

¹ Massachusetts Institute of Technology, MA, USA

² Google Research, Mountain View, CA, USA

³ University of Southern California, CA, USA

⁴ Facebook Inc., CA, USA

⁵ Google Research, Cambridge, MA, USA

Jiajun Wu and Tianfan Xue contributed equally to this work.

1 Introduction

Deep networks have achieved impressive performance on image recognition (Russakovsky et al 2015). Nonetheless, for any visual system to parse objects in the real world, it needs not only to assign category labels to objects, but also to interpret their intra-class variation. Figure 2 shows an example: for a chair, we are interested in its intrinsic properties such as its *style*, *height*, *leg length*, and *seat width*, and extrinsic properties such as its *pose*.

In this paper, we recover these object properties from a single image by jointly estimating the object's 3D wireframe and the viewpoint. We choose to use a single image as input primarily for three reasons. First, this problem has great scientific value: humans can easily recognize 3D structure from a single image, and we want to build machines that replicate such an ability. Second, by starting with a single image, instead of multi-view images or videos, our model can be directly applied to images taken in the wild, *e.g.*,

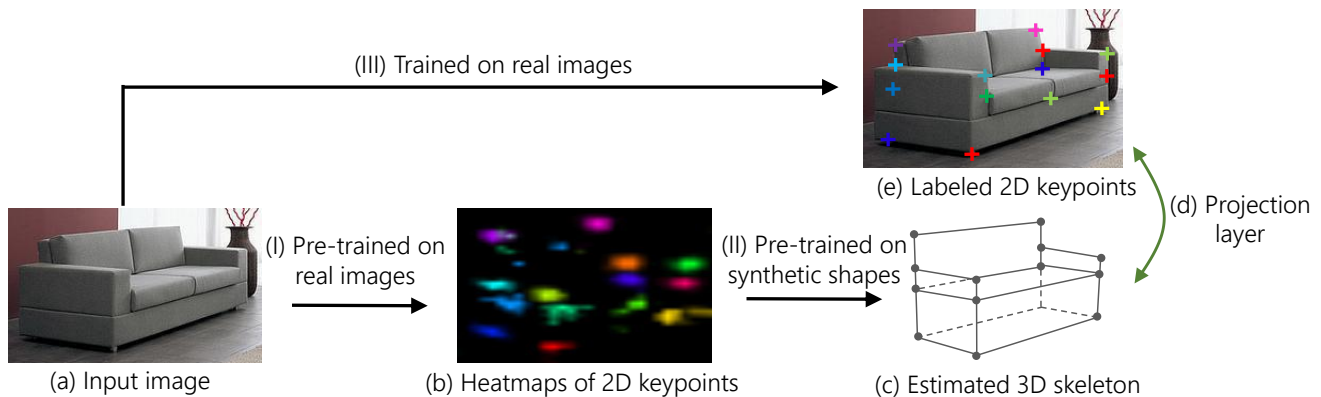


Fig. 1: An overview of our model. For an image (a) with a category-level label (sofa), the system first estimates its 2D keypoint heatmaps (b), and then recovers the 3D skeleton of the object (c). During training, through the projection layer (d), it also enforces the consistency between annotated 2D keypoint locations (e) and projected 2D locations of estimated 3D keypoints.



Fig. 2: Given an image of a chair, we are interested in its intrinsic properties such as its height, leg length, and seat width, and extrinsic properties such as its pose.

from the web, and can cluster them based on their structure or pose. This offers extensive practical usages: social media and e-commerce companies can better understand and analyze user-uploaded data, and household robots can efficiently interact with objects after modeling them in 3D. Third, using a single image as input enables online inference: for moving objects like cars, the system can reconstruct their geometry and viewpoint on the fly. This is crucial for real-time applications such as autonomous driving.

We represent an object via a 3D skeleton (Torresani et al 2003) (Figure 1c), instead of a 3D mesh or a depth map (Doso-vitskiy et al 2015; Aubry et al 2014; Prasad et al 2010; Kar et al 2015; Su et al 2014; Vicente et al 2014; Huang et al 2015), because skeletons are simpler and preserve the structural properties that we are interested in. We refer readers to Section 6 for detailed discussions on our design choices. We assume that a 3D skeleton consists of keypoints and the connections between them, and manually pre-define the skeleton model for each object category (e.g., chair, sofa, and

car). Then, our task is to estimate 3D keypoint locations of an object from a single RGB image.

The main challenge of single image 3D estimation is the difficulty in obtaining training images with ground truth 3D structure, as manually annotating 3D object structure in real images is labor-intensive and often inaccurate. Previous methods tackled this problem mostly in two ways. One is to directly solve for a 3D skeleton from estimated 2D keypoint locations by minimizing its reprojection error (Zhou et al 2015), without any 3D annotations. Most algorithms in this category do not encode 3D shape priors within the model, and thus they are not robust to inaccurate keypoint estimation, as shown in experiments in Section 4. The other is to train on synthetically rendered images of 3D objects (Li et al 2015; Su et al 2015); in synthetic data, complete 3D structure is available. But the statistics of synthesized images are often different from those of real images, due to changes in lighting, occlusion, and shape details. This makes it hard for models trained only on synthetic data to generalize well to real images.

In this paper, we propose 3D INterpreter Networks (3D-INN), an end-to-end trainable framework for 3D skeleton and viewpoint estimation. In contrast to prior art, our model learns from both 2D-labeled real images and synthetic 3D objects. This is made possible by two major innovations. First, we use heatmaps of 2D keypoints as an intermediate representation to connect real and synthetic data. 3D-INN is trained on real images to estimate 2D keypoint heatmaps from an input image (Figure 1-I); it then learns from synthetic data to estimate 3D structure from heatmaps (Figure 1-II). By doing so, 3D-INN benefits from the variations in abundant synthetic 3D objects, without suffering from the domain difference between real and synthesized data.

Second, we introduce a *Projection Layer*, a rendering function that calculates projected 2D keypoint positions given

a 3D skeleton and camera parameters. We attach it at the end of the framework (Figure 1d). This enables 3D-INN to predict 3D structural parameters that minimize errors in 2D on real images with keypoint labels, without requiring 3D object annotations. Our training paradigm therefore consists of three steps: we first train the keypoint estimation component (Figure 1-I) on 2D-annotated real images; we then train the 3D interpreter (Figure 1-II) on synthetic 3D data; we finally fine-tune the entire framework end-to-end with the projection layer (Figure 1-III).

Both innovations are essential for the system to exploit the complementary richness of real 2D data and synthetic 3D data. Without the heatmap representation and synthetic 3D data, the system can still be trained on real images with 2D annotations, using the projection layer. But it does not perform well: because of intrinsic ambiguities in 2D-to-3D mapping, the algorithm recovers unnatural 3D geometries, though their projections may perfectly align with 2D annotations, as explored in Lowe (1987). Synthetic 3D shapes help the network resolve this ambiguity by learning prior knowledge of “plausible shapes”. Without the projection layer, the system becomes two separately trained networks: one trained on real images to estimate 2D keypoint heatmaps, and the other trained on synthetic data to estimate 3D structure. As shown in Section 4, the 3D predictions in this case are not as accurate due to the domain adaptation issue.

Several experiments demonstrate the effectiveness of 3D-INN. First, the proposed network achieves good performance on various keypoint localization datasets, including FLIC (Sapp and Taskar 2013) for human bodies, CUB-200-2011 (Wah et al 2011) for birds, and our new dataset, Keypoint-5, for furniture. We then evaluate our network on IKEA (Lim et al 2013), a dataset with ground truth 3D object structure and viewpoints. We augmented the original IKEA dataset with additional 2D keypoint labels. On 3D structure estimation, 3D-INN shows its advantage over an optimization-based method (Zhou et al 2015) when keypoint estimation is imperfect. On 3D viewpoint estimation, it also performs better than the state of the art (Su et al 2015). We further evaluate 3D-INN, in combination with an object detection framework, R-CNN (Girshick et al 2014), on PASCAL 3D+ benchmark (Xiang et al 2014) for joint detection and viewpoint estimation. 3D-INN also achieves results comparable to the state of the art (Su et al 2015; Tulsiani and Malik 2015). At last, we show that 3D-INN has wide vision applications including 3D object retrieval.

Our contributions are three-fold. First, we introduce 3D INterpreter Networks (3D-INN); by incorporating 2D keypoint heatmaps to connect real and synthetic worlds, we strengthen the generalization ability of the network. Second, we propose a projection layer, so that 3D-INN can be trained to predict 3D structural parameters using only 2D-annotated images. Third, our model achieves state-of-the-art perfor-

mance on both 2D keypoint and 3D structure and viewpoint estimation.

2 Related work

Single Image 3D Reconstruction Previous 3D reconstruction methods mainly modeled objects using either dense representations such as depth or meshes, or sparse representations such as skeletons or pictorial structure. Depth-/mesh-based models can recover detailed 3D object structure from a single image, either by adapting existing 3D models from a database (Aubry et al 2014; Satkin et al 2012; Su et al 2014; Huang et al 2015; Zeng et al 2016; Wu et al 2016; Hu and Zhu 2015; Bansal and Russell 2016; Shrivastava and Gupta 2013; Choy et al 2016), or by inferring from its detected 2D silhouette (Kar et al 2015; Soltani et al 2017; Vicente et al 2014; Prasad et al 2010; Wu et al 2017).

In this paper, we choose to use a skeleton-based representation, exploiting the power of abstraction. The skeleton model can capture geometric changes of articulated objects (Torresani et al 2003; Yasin et al 2016; Akhter and Black 2015), like a human body or the base of a swivel chair. Typically, researchers recovered a 3D skeleton from a single image by minimizing its projection error on the 2D image plane (Lowe 1987; Leclerc and Fischler 1992; Hejrati and Ramanan 2014; Xue et al 2012; Ramakrishna et al 2012; Zia et al 2013). Recent work in this line (Akhter and Black 2015; Zhou et al 2015) demonstrated state-of-the-art performance. In contrast to them, we propose to use neural networks to predict a 3D object skeleton from its 2D keypoints, which is more robust to imperfect detection results and can be jointly learned with keypoint estimators.

Our work also connects to the traditional field of vision as inverse graphics (Hinton and Ghahramani 1997; Kulkarni et al 2015b) and analysis by synthesis (Yuille and Kersten 2006; Kulkarni et al 2015a; Bever and Poeppl 2010; Wu et al 2015), as we use neural nets to decode latent 3D structure from images, and use a projection layer for rendering. Their approaches often required supervision for the inferred representations or made over-simplified assumptions of background and occlusion in images. Our 3D-INN learns 3D representation without using 3D supervision, and generalizes to real images well.

2D Keypoint Estimation Another line of related work is 2D keypoint estimation. During the past decade, researchers have made significant progress in estimating keypoints on humans (Sapp and Taskar 2013; Yang and Ramanan 2011) and other objects (Wah et al 2011; Shih et al 2015). Recently, there have been several attempts to apply convolutional neural networks to human keypoint estimation (Toshev and Szegedy 2014; Tompson et al 2015; Carreira et al

2016; Newell et al 2016), which all achieved significant improvement. 3D-INN uses 2D keypoints as an intermediate representation, and aims to recover a 3D skeleton from them.

3D Viewpoint Estimation 3D viewpoint estimation seeks to estimate the 3D orientation of an object from a single image (Xiang et al 2014). Some previous methods formulated it as a classification or regression problem, and aimed to directly estimate the viewpoint from an image (Fidler et al 2012; Su et al 2015). Others proposed to estimate 3D viewpoint from detected 2D keypoints or edges in the image (Zia et al 2013; Lim et al 2014; Tulsiani and Malik 2015). While the main focus of our work is to estimate 3D object structure, our method can also predict the corresponding 3D viewpoint.

Training with Synthetic Data Synthetic data are often used to augment the training set (Su et al 2014; Shakhnarovich et al 2003), especially when ground truth labels of real images are hard to obtain. This technique has found wide applications in computer vision. To name a few, Sun and Saenko (2014) and Zhou et al (2016) combined real and synthetic data for object detection and matching, respectively. Huang et al (2015) analyzed the invariance of convolutional neural networks using synthetic images. Dosovitskiy et al (2015) trained a neural network for image synthesis using synthetic images. McCormac et al (2017) rendered images for indoor scene understanding. Su et al (2014) attempted to train a 3D viewpoint estimator on both real and synthetic images.

In this paper, we combine real 2D-annotated images and synthetic 3D data for training 3D-INN to recover a 3D skeleton. We use heatmaps of 2D keypoints, instead of (often imperfectly) rendered images, from synthetic 3D data, so that our algorithm has better generalization ability as the effects of imperfect rendering are minimized. Yasin et al (2016) also proposed to use both 2D and 3D data for training, but they used keypoint locations, instead of heatmaps, as the intermediate representation that connects 2D and 3D. While their focus is on estimating human poses, we study the problem of recovering the 3D structure of furniture and cars.

3 Method

We design a deep convolutional network to recover 3D object structure. The input to the network is a single image with the object of interest at its center, which can be obtained by state-of-the-art object detectors. The output of the network is a 3D object skeleton, including its 2D keypoint locations, 3D structural parameters, and 3D pose (see Figure 4). In the following sections, we will describe our 3D skeleton representation and the camera model (Section 3.1), network design (Section 3.2), and training strategy (Section 3.3).

3.1 3D Skeleton Representation

We use skeletons as our 3D object representation. A skeleton consists of a set of keypoints as well as their connections. For each object category, we manually design a 3D skeleton characterizing its abstract 3D geometry.

There exist intrinsic ambiguities in recovering 3D keypoint locations from a single 2D image. We resolve this issue by assuming that objects can only have constrained deformations (Torresani et al 2003). For example, chairs may have various leg lengths, but for a single chair, its four legs are typically of equal length. We model these constraints by formulating 3D keypoint locations as a weighted sum of a set of base shapes (Kar et al 2015). The first base shape is the mean shape of all objects within the category, and the rest define possible deformations and intra-class variations. Figure 3a shows our skeleton representation for chairs: the first is the mean shape of chairs, the second controls how the back bends, and the last two are for legs. Figure 3b shows base shapes for cars. The weight for each base shape determines how strong the deformation is, and we denote these weights as the *structural parameters* of an object.

Formally, let $\mathbf{Y} \in \mathbb{R}^{3 \times N}$ be a matrix of 3D coordinates of all N keypoints. Our assumption is that the 3D keypoint locations are a weighted sum of base shapes $B_k \in \mathbb{R}^{3 \times N}$, or

$$\mathbf{Y} = \sum_{k=1}^K \alpha_k B_k, \quad (1)$$

where $\{\alpha_k\}$ is the set of structural parameters of this object, and K is the number of base shapes.



Further, let $\mathbf{X} \in \mathbb{R}^{2 \times N}$ be the corresponding 2D coordinates. Then the relationship between the observed 2D coordinates \mathbf{X} and the structural parameters $\{\alpha_k\}$ is

$$\mathbf{X} = P(R\mathbf{Y} + T) = P\left(R \sum_{k=1}^K \alpha_k B_k + T\right), \quad (2)$$

where $R \in \mathbb{R}^{3 \times 3}$ (rotation) and $T \in \mathbb{R}^3$ (translation) are the external parameters of the camera, and $P \in \mathbb{R}^{3 \times 4}$ is the camera projection matrix which we will discuss soon.

Therefore, to recover the 3D structural information of an object in a 2D image, we only need to estimate its structural parameters ($\{\alpha_k\}$) and the external viewpoint parameters (R , T , and f). We supply the detailed camera model below. In Section 3.2 and Section 3.3, we discuss how we design a neural network for this task, and how it can be jointly trained with real 2D images and synthetic 3D objects.

Camera Model We use perspective projection in order to model the perspective distortion in 2D images. We assume that the principal point is at the origin, all pixels are square, and there is no axis skew. In this way, we only need to estimate the focal length in the camera projection matrix P .

(a) Base shapes for chairs. B_1 is the mean shape of chairs, and the others characterize possible variations of the structure.

(b) Base shapes for cars

Fig. 3: Our skeleton model and base shapes for chairs (a) and cars (b).

For the ease of inference in neural networks, we rewrite the normal perspective projection as follows. Let $x_i \in \mathbb{R}^2$ be a column vector of the 2D coordinates of the i -th keypoint and y_i be the corresponding 3D coordinates to be recovered. We assume that the camera center is at $(0, 0, f)$ instead of the origin. The perspective projection is written as (using projective coordinates):

$$\begin{pmatrix} x_i^1 \\ x_i^2 \\ 1 \end{pmatrix} = \begin{pmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} y_i^1 \\ y_i^2 \\ y_i^3 + f \\ 1 \end{pmatrix}, \quad (3)$$

where f is the focal length, x_i^1 and x_i^2 are the x- and y-components of x_i , and y_i^1, y_i^2, y_i^3 are x-, y-, and z-components of y_i .

When $f^{-1} \rightarrow 0$, Equation 3 converges to the formulation of parallel projection. To see that, based on Equation 3, we get the Euclidean coordinates of the 2D projection as (we abuse the notation of x_i^1 and x_i^2 for both Euclidean coordinates and projective coordinates)

$$\begin{cases} x_i^1 = \frac{fy_i^1}{y_i^3 + f} = \frac{y_i^1}{f^{-1}y_i^3 + 1}, \\ x_i^2 = \frac{fy_i^2}{y_i^3 + f} = \frac{y_i^2}{f^{-1}y_i^3 + 1}. \end{cases} \quad (4)$$

Then when $f \rightarrow \infty$, we have

$$\begin{cases} x_i^1 = y_i^1, \\ x_i^2 = y_i^2, \end{cases} \quad (5)$$

which is the formulation of parallel projection. Therefore, Equation 3 models the perspective projection when $f^{-1} \neq 0$ and models the parallel projection when $f^{-1} = 0$.

3.2 3D INterpreter Networks (3D-INN)

Our network consists of three components: first, a keypoint estimator, which localizes 2D keypoints of objects from 2D images by regressing their heatmaps (Figure 4a and b, in blue); second, a 3D interpreter, which infers internal 3D structural and viewpoint parameters from the heatmaps (Figure 4c, in red); third, a projection layer, mapping 3D skeletons to 2D keypoint locations so that real 2D-annotated images can be used as supervision (Figure 4d, in yellow).

Keypoint Estimation The keypoint estimation consists of two steps: initial estimation (Figure 4a) and keypoint refinement (Figure 4b).

The network architecture for initial keypoint estimation is inspired by the pipeline proposed by [Tompson et al \(2014, 2015\)](#). The network takes multi-scaled images as input and estimates keypoint heatmaps. Specifically, we apply Local Contrast Normalization (LCN) on each image, and then scale it to 320×240 , 160×120 , and 80×60 as input to three separate scales of the network. The output is k heatmaps, each with resolution 40×30 , where k is the number of keypoints of the object in the image.

At each scale, the network has three sets of 5×5 convolutional (with zero padding), ReLU, and 2×2 pooling layers, followed by a 9×9 convolutional and ReLU layer. The final outputs for the three scales are therefore images with resolution 40×30 , 20×15 , and 10×7 , respectively. We then up-sample the outputs of the last two scales to ensure they have the same resolution (40×30). The outputs from the three scales are later summed up and sent to a Batch Normalization layer and three 1×1 convolution layers, whose goal is to regress target heatmaps. We found that Batch Normalization is critical for convergence, while Spatial Dropout, proposed in [Tompson et al \(2015\)](#), does not affect performance.

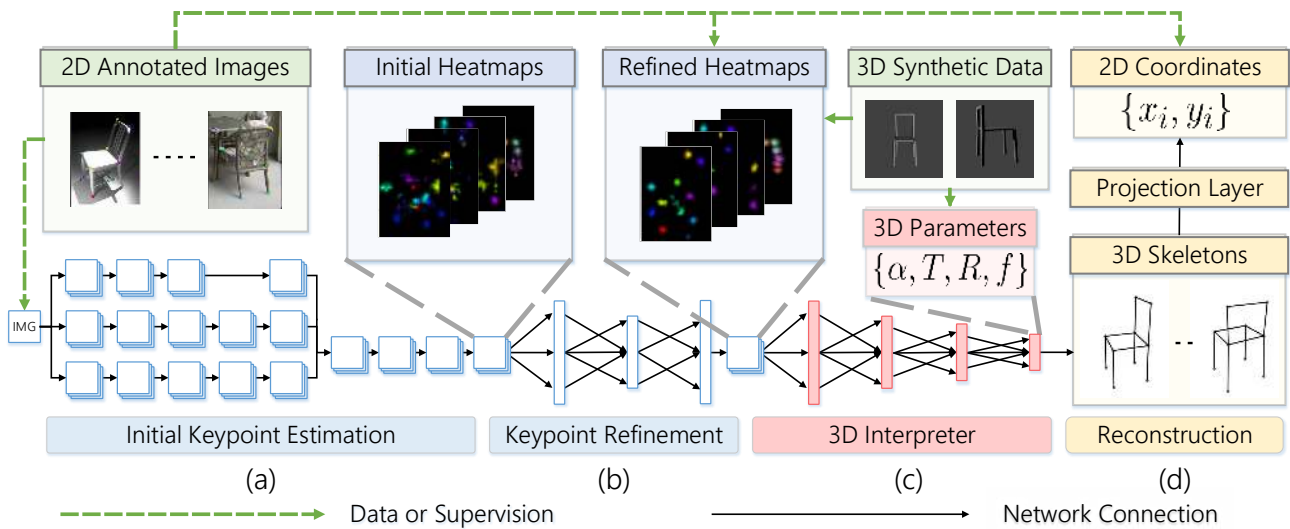


Fig. 4: 3D-INN takes a single image as input and reconstructs the detailed 3D structure of the object in the image (*e.g.*, human, chair, *etc.*). The network is trained independently for each category, and here we use chairs as an example. **(a)** Estimating 2D keypoint heatmaps with a multi-scale CNN. **(b)** Refining keypoint locations by considering the structural constraints between keypoints. This is implicitly enforced with an information bottleneck and yields cleaner heatmaps. **(c)** Recovered 3D structural and camera parameters $\{\alpha, T, R, f\}$. **(d)** The projection layer maps reconstructed 3D skeletons back to 2D keypoint coordinates.

The second step of keypoint estimation is keypoint refinement, whose goal is to implicitly learn category-level structural constraints on keypoint locations after the initial keypoint localization. The motivation is to exploit the contextual and structural knowledge among keypoints (*e.g.*, arms cannot be too far from the torso). We design a mini-network which, like an autoencoder, has information bottleneck layers, enforcing it to implicitly model the relationship among keypoints. Some previous works also use this idea and achieve better performance with lower computational cost in object detection (Ren et al 2015) and face recognition (Taigman et al 2015).

In the keypoint refinement network, we use three fully connected layers with widths 8,192, 4,096, and 8,192, respectively. After refinement, the heatmaps of keypoints are much cleaner, as shown in Figure 6 and Section 4.

3D Interpreter The goal of our 3D interpreter is to infer 3D structure and viewpoint parameters, using estimated 2D heatmaps from earlier layers. While there are many different ways of solving Equation 2, our deep learning approach has clear advantages. First, traditional methods (Hejrati and Ramanan 2012; Torresani et al 2003) that minimize the re-projection error consider only one keypoint hypothesis, and is therefore not robust to noisy keypoint detection. In contrast, our framework uses soft heatmaps of keypoint locations as input, as shown in Figure 4c, which is more robust when some keypoints are invisible or incorrectly located.

Second, our algorithm only requires a single forward propagation during testing, making it more efficient than the most previous optimization-base methods.

As discussed in Section 3.1, the set of 3D parameters we estimate is of $S = \{\alpha_i, R, T, f\}$, with which we are able to recover the 3D object structure using Equation 2. In our implementation, the network predicts f^{-1} instead of f for better numerical stability. As shown in Figure 4c, we use four fully connected layers as our 3D interpreter, with widths 2,048, 512, 128, and $|S|$, respectively. Spatial Transformer Networks (Jaderberg et al 2015) also explored the idea of learning rotation parameters R with neural nets, but our network can also recover structural parameters $\{\alpha_i\}$.

Projection Layer The last component of the network is a projection layer (Figure 4d). The projection layer takes estimated 3D parameters as input, and computes projected 2D keypoint coordinates $\{x_i, y_i\}$ using Equation 2. As all operations are differentiable, the projection layer enables us to use 2D-annotated images as ground truth, and run back-propagation to update the entire network.

3.3 Training Strategy

A straightforward training strategy is to use real 2D images as input, and their 2D keypoint locations as supervision for the output of the projection layer. Unfortunately, experi-

ments show that the network can hardly converge using this training scheme, due to the high-dimensional search space and the ambiguity in the 3D to 2D projection.

We therefore adopt an alternative three-step training strategy: first, training the keypoint estimator (Figure 4a and 4b) using real images with 2D keypoint heatmaps as supervision; second, training the 3D interpreter (Figure 4c) using synthetic 3D data as there are no ground truth 3D annotations available for real images; and third, training the whole network using real 2D images with supervision on the output of the projection layer at the end.

To generate synthetic 3D objects, for each object category, we first randomly sample structural parameters $\{\alpha_i\}$ and viewpoint parameters P , R and T . We calculate 3D keypoint coordinates using Equation 2. To model deformations that cannot be captured by base shapes, we add Gaussian perturbation to 3D keypoint locations of each synthetic 3D object, whose variance is 1% of its diagonal length. Examples of synthetic 3D shapes are shown in Figure 4c. In experiments, we do not render synthesized shapes; we use heatmaps of keypoints, rather than rendered images, as training input.

4 Evaluation

We evaluate our entire framework, 3D-INN, as well as each component within. In this section, we present both qualitative and quantitative results on 2D keypoint estimation (Section 4.1) and 3D structure and viewpoint recovery (Section 4.2).

4.1 2D Keypoint Estimation

Data For 2D keypoint estimation, we evaluate our algorithm on three image datasets: FLIC (Sapp and Taskar 2013) for human bodies, CUB-200-2011 (Wah et al 2011) for birds, and a new dataset Keypoint-5 for furniture. Specifically, FLIC is a challenging dataset containing 3,987 training images and 1,016 test images, each labeled with 10 keypoints of human bodies. The CUB-200-2011 dataset was originally proposed for fine-grained bird classification, but with labeled keypoints of bird parts. It has 5,994 images for training and 5,794 images for testing, each coming with up to 15 keypoints.

We also introduce a new dataset, Keypoint-5, which contains five categories: bed, chair, sofa, swivel chair, and table. There are 1,000 to 2,000 images in each category, where 80% are for training and 20% for testing. For each image, we asked three workers on Amazon Mechanical Turk to label locations of a pre-defined category-specific set of keypoints; we then, for each keypoint, used the median of the three labeled locations as ground truth.

Table 1: Keypoint estimation results on CUB-200-2011, measured in PCP (%) and AE. Our method is comparable to Mdshift (Shih et al 2015) (better in AE but worse in PCP), and better than all other algorithms.

Method	PCP (%)	Average Error
Poselets (Bourdev et al 2010)	27.47	2.89
Consensus (Belhumeur et al 2013)	48.70	2.13
Exemplar (Liu and Belhumeur 2013)	59.74	1.80
Mdshift (Shih et al 2015)	69.1	1.39
3D-INN (ours)	66.7	1.36
Human	84.72	1.00

Table 2: Keypoint estimation results of 3D-INN and Tompson et al (2015) on Keypoint-5, measured in PCP (%) and AE. 3D-INN is consistently better in both measures. We re-trained the network in Tompson et al (2015) on Keypoint-5.

	Method	Bed	Chair	Sofa	Swivel Chair
PCP	3D-INN (ours)	77.4	87.7	77.4	78.5
	Tompson et al	76.2	85.3	76.9	69.2
AE	3D-INN (ours)	1.16	0.92	1.14	1.19
	Tompson et al	1.20	1.02	1.19	1.54

Metrics To quantitatively evaluate the accuracy of estimated keypoints on FLIC (human body), we use the standard Percentage of Correct Keypoints (PCK) measure (Sapp and Taskar 2013) to be consistent with previous works (Sapp and Taskar 2013; Tompson et al 2014, 2015). We use the evaluation toolkit and results of competing methods released by Tompson et al (2015). On CUB-200-2011 (bird) and the new Keypoint-5 (furniture) dataset, following the convention (Liu and Belhumeur 2013; Shih et al 2015), we evaluate results in Percentage of Correct Parts (PCP) and Average Error (AE). PCP is defined as the percentage of keypoints localized within 1.5 times of the standard deviation of annotations. We use the evaluation code from (Liu and Belhumeur 2013) to ensure consistency. Average error is computed as the mean of the distance, bounded by 5, between a predicted keypoint location and ground truth.

Results For 2D keypoint detection, we only train the keypoint estimator in our 3D-INN (Figure 4a and 4b) using the training images in each dataset. Figure 5 shows the accuracy of keypoint estimation on the FLIC dataset. On this dataset, we employ a fine-level network for post-processing, as suggested by Tompson et al (2015). Our method performs better than all previous methods (Sapp and Taskar 2013; Tompson et al 2014, 2015; Yang and Ramanan 2011; Toshev and Szegedy 2014) at all precisions. Moreover, the keypoint refinement step (Figure 4c) improves results significantly (about 2% for a normalized distance ≥ 0.15), without which our framework has similar performance with Tompson et al




Fig. 5: (a) PCK curves on the FLIC dataset (Sapp and Taskar 2013). 3D-INN performs consistently better than other methods. Without keypoint refinement, it is comparable to Tompson et al (2015). (b) A zoomed view of the dashed rectangle in (a).




Fig. 6: 2D keypoint predictions from a single image, where each color corresponds to a keypoint. The keypoint refinement step cleans up false positives and produces more regulated predictions.

(2015). Such improvement is also demonstrated in Figure 6, where the heatmaps after refinement are far less noisy.

The accuracy of keypoint estimation on CUB-200-201 dataset is listed in Table 1. Our method is better than Liu and Belhumeur (2013) in both metrics, and is comparable to the state-of-the-art (Shih et al 2015). Specifically, compared with Shih et al (2015), our model more precisely estimates the keypoint locations for correctly detected parts (a lower AE), but misses more parts in the detection (a lower PCP). On our Keypoint-5 dataset, our model achieves higher PCPs and lower AEs compared to the state-of-the-art (Tompson et al 2015) for all categories, as shown in Table 2. These

experiments in general demonstrate the effectiveness of our model on keypoint detection.

4.2 3D Structure and Viewpoint Estimation

For 3D structural parameter estimation, we evaluate 3D-INN from three different perspectives. First, we evaluate our 3D interpreter (Figure 4c alone) against the optimization-based method (Zhou et al 2015). Second, we test our full pipeline on the IKEA dataset (Lim et al 2013), where ground truth 3D labels are available, comparing with both baselines and the state-of-the-art. Third, we show qualitative results on four




Fig. 7: Plots comparing our method against an analytic solution on synthetic heatmap. (a) The accuracy of 3D structure estimation; (b) The accuracy of 3D viewpoint estimation.




Fig. 8: Evaluation on chairs in the IKEA dataset (Lim et al 2013). The network trained with our paradigm (3D-INN) is significantly better than the one trained from *scratch* on both 3D structure (a) and viewpoint estimation (b).

datasets: Keypoint-5, IKEA, the SUN database (Xiao et al 2010), and PASCAL 3D+ Xiang et al (2014).

Comparing with an optimization-based method. First, we compare our 3D interpreter (Figure 4c) with the state-of-the-art optimization-based method that directly minimizes re-projection error (Equation 2) on the synthetic data. As most optimization based methods only consider the parallel projection, while we model perspective projection for real images, we extend the one by Zhou et al (2015) as follows: we first use their algorithm to get an initial guess of internal parameters and viewpoints, and then applying a simple gradient descent method to refine it considering perspective distortion.

We generate synthetic data for this experiment, using the scheme described in Section 3.3. Each data point contains the 2D keypoint heatmaps of an object, and its corresponding 3D keypoint locations and viewpoint, which we would like to estimate. We also add different levels of salt-and-pepper noise to heatmaps to evaluate the robustness of

both methods. We generated 30,000 training and 1,000 testing cases. Because the analytic solution only takes keypoint coordinates as input, we convert heatmaps to coordinates using an argmax function.

For both methods, we evaluate their performance on both 3D structure recovery and 3D viewpoint estimation. To evaluate the estimated 3D structure, we compare their accuracies on 3D keypoint estimation (\mathbf{Y} in Section 3.1); for 3D viewpoint estimation, we compute errors in azimuth angle, following previous work (Su et al 2015). As the original algorithm by Zhou et al (2015) was mainly designed for the parallel projection and comparatively clean heatmaps, our 3D interpreter outperforms it in the presence of noise and perspective distortion, as shown in Figure 7. Our algorithm is also efficient, taking less than 50 milliseconds for each test image.

Evaluating the full pipeline. We now evaluate 3D-INN on estimating 3D structure and 3D viewpoint. We use the IKEA dataset (Lim et al 2013) for evaluation, as it provides ground

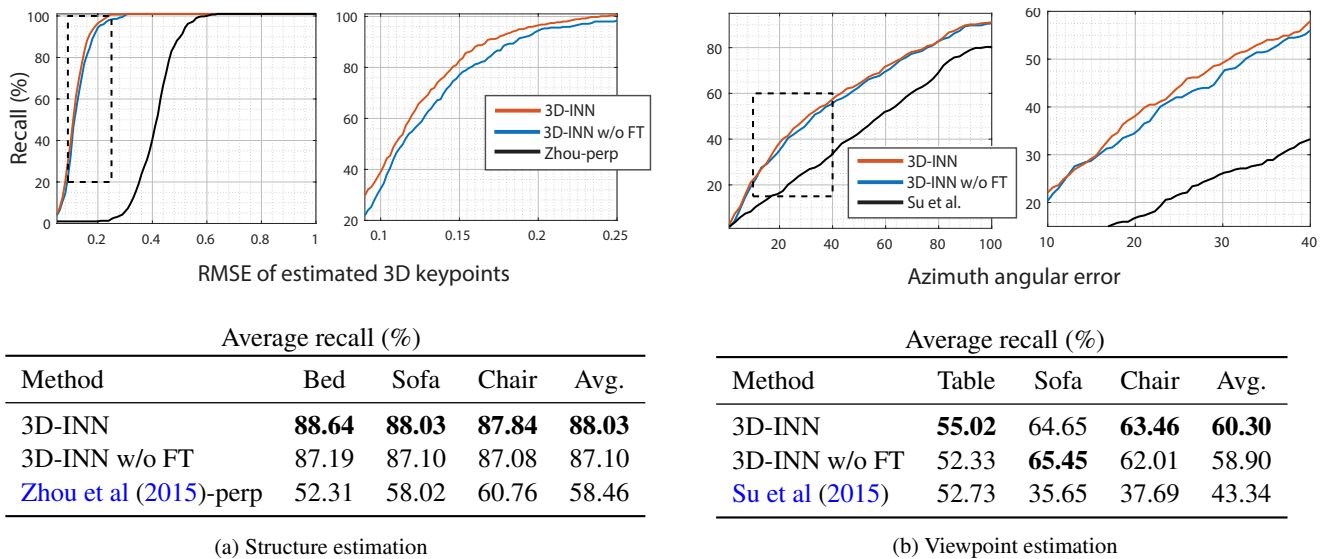


Fig. 9: Evaluation on the IKEA dataset (Lim et al 2013). (a) The accuracy of structure estimation. RMSE-Recall curved is shown in the first row, and zoomed-views of the dashed rectangular regions are shown on the right. The third row shows the average recall on all thresholds. (b) The accuracy of viewpoint estimation.

Table 3: Joint object detection and viewpoint estimation on PASCAL 3D+ (Xiang et al 2014). Following previous work, we use Average Viewpoint Precision (AVP) as our measure, which extends AP so that a true positive should have both a correct bounding box and a correct viewpoint (here we use a 4-view quantization). Both 3D-INN and V&K (Tulsiani and Malik 2015) use R-CNN (Girshick et al 2014) for object detection, precluding the influence of object detectors. The others use their own detection algorithm. VDPM (Xiang et al 2014) and DPM-VOC+VP (Pepik et al 2012) are trained on PASCAL VOC 2012, V&K (Tulsiani and Malik 2015) is trained on PASCAL 3D+, Su et al (2015) is trained on PASCAL VOC 2012, together with synthetic 3D CAD models, and 3D-INN is trained on Keypoint-5.

Category	VDPM (Xiang et al 2014)	DPM-VOC+VP (Pepik et al 2012)	Su et al (2015)	V&K (Tulsiani and Malik 2015)	3D-INN
Chair	6.8	6.1	15.7	25.1	23.1
Sofa	5.1	11.8	18.6	43.8	45.8
Car	20.2	36.9	41.8	55.2	52.2

truth 3D mesh models and the associated viewpoints for testing images. We manually label ground truth 3D keypoint locations on provided 3D meshes, and calculate the root-mean-square error (RMSE) between estimated and ground truth 3D keypoint locations.

As IKEA only has no more than 200 images per category, we instead train 3D-INN on our Keypoint-5, as well as one million synthetic data points, using the strategy described in Section 3.3. Note that, first, we are only using no more than 2,000 real images per category for training and, second, we are testing the trained model on different datasets, avoiding possible dataset biases (Torralba and Efros 2011).

We first compare with a baseline method to evaluate our training paradigm: we show quantitative comparisons between 3D-INN trained using our paradigm proposed in Section 3.3, and the same network but only end-to-end trained

with real images, without having the two pre-training stages. We called it the *scratch* model.

As shown in the RMSE-Recall curve in Figure 8, 3D-INN performs much better than *scratch* on both 3D structure and viewpoint estimation. The average recall of 3D-INN is about 20% higher than *scratch* in 3D structure estimation, and about 40% higher in 3D pose estimation. This shows the effectiveness of the proposed training paradigm.

We then compare our full model with the state-of-the-art methods. The left half of Figure 9 shows RMSE-Recall curve of both our algorithm and the optimization-based method described above (Zhou et al (2015)-perp). The y-axis shows the recall — the percentage of testing samples under a certain RMSE threshold. We test two versions of our algorithm: with fine-tuning (3D-INN) and without fine-tuning (3D-INN w/o FT). Both significantly outperform the optimization-based method (Zhou et al 2015). This is because the method from Zhou et al (2015) was not designed to handle noisy key-

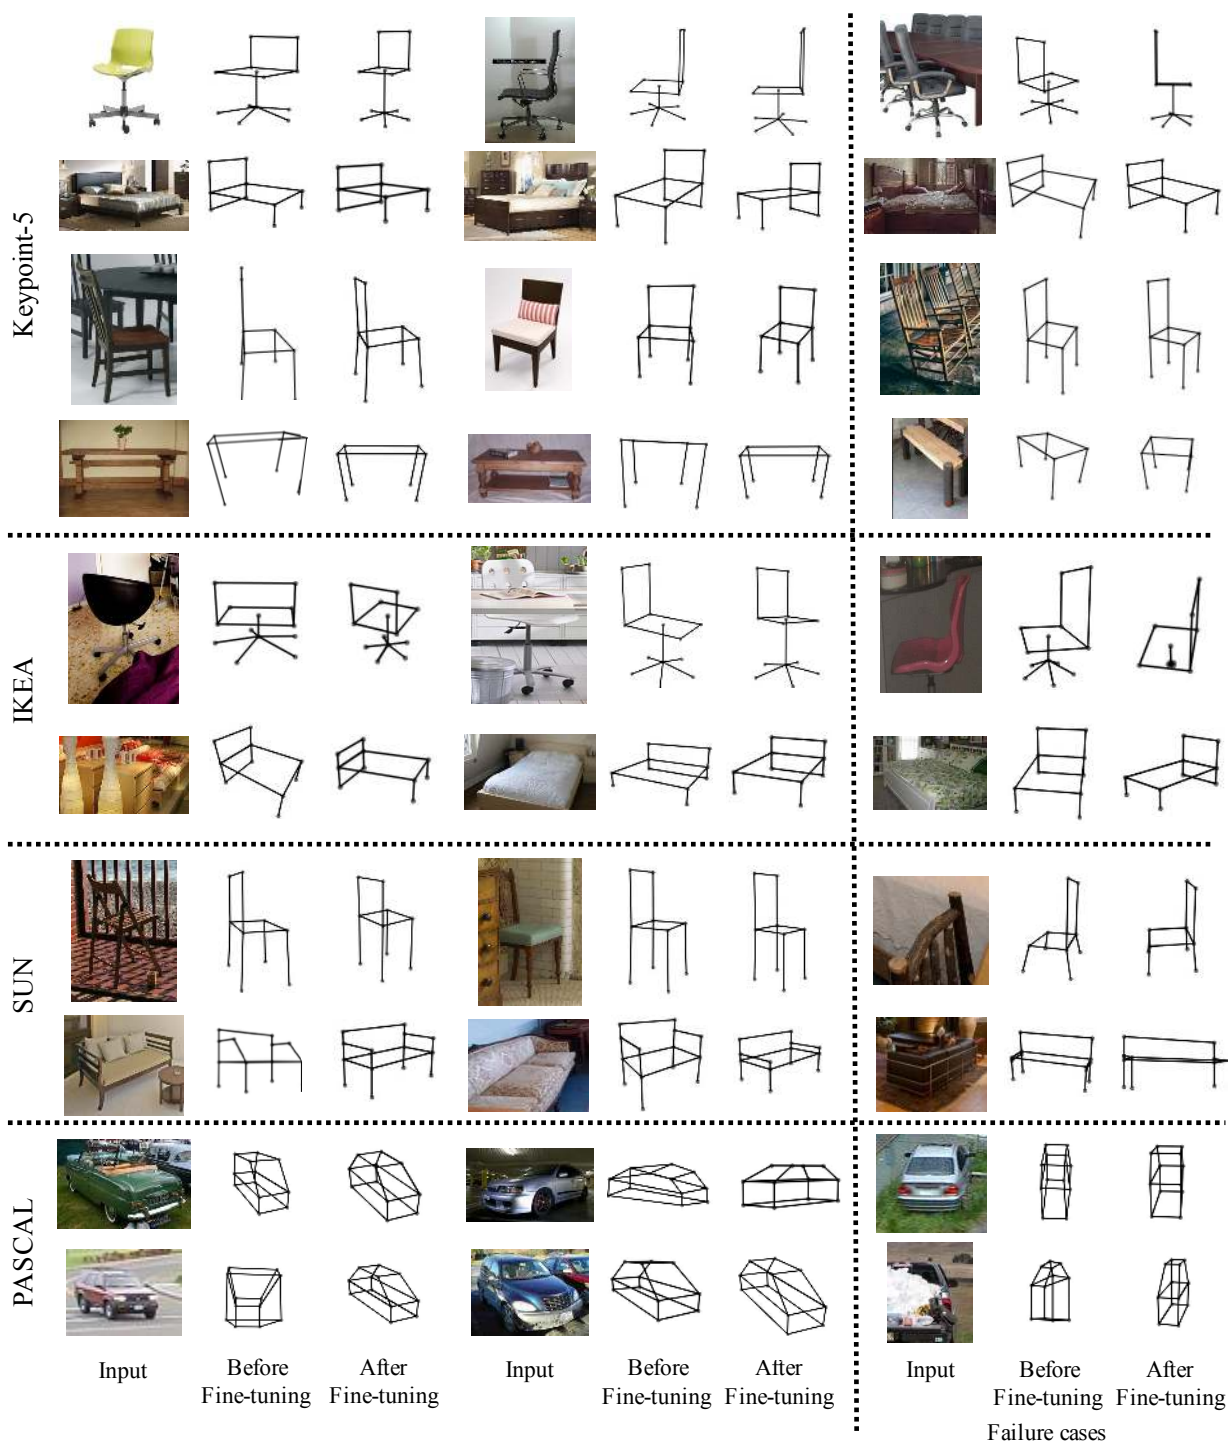


Fig. 10: Qualitative results on Keypoint-5, IKEA, and SUN databases. For each example, the first one is the input image, the second one is the reconstruct 3D skeleton using the network before fine-tuning, and third one is using the network after fine-tuning. The last column shows failure cases.

point estimation and perspective distortions, while our 3D-INN can deal with them. Also, fine-tuning improves the accuracy of keypoint estimation by about 5% under the RMSE threshold 0.15.

Though we focus on recovering 3D object structure, as an extension, we also evaluate 3D-INN on 3D viewpoint estimation. We compare it with the state-of-the-art viewpoint estimation algorithm by [Su et al \(2015\)](#). The right




Fig. 11: Car structure estimation on images from the PASCAL 3D+ dataset



(a) Training: beds, Test: chairs

(b) Training: sofas, Test: chairs

Fig. 12: Qualitative results on chairs using networks trained on sofas or beds. In most cases models provide reasonable output. Mistakes are often due to the difference between the training and test sets, *e.g.*, in the third example, the model trained on beds fails to estimate chairs facing backward.

half of Figure 9 shows the results (recall) in azimuth angle. As shown in the table, 3D-INN outperforms Su et al (2015) by about 40% (relative), measured in average recall. This is mainly because it is not straightforward for Su et al (2015), mostly trained on (cropped) synthesized images, to deal with the large number of heavily occluded objects in the IKEA dataset.

Although our algorithm assumes a centered object in an input image, we can apply it, in combination with an object detection algorithm, on images where object locations are unknown. We evaluate the results of joint object detection and viewpoint estimation on PASCAL 3D+ dataset (Xiang et al 2014). PASCAL 3D+ and Keypoint-5 has two overlapping categories: chair and sofa, and we evaluate on both. We also study an additional object category, car, for which 3D-INN is trained on 1,000 car images from ImageNet (Rusakovsky et al 2015) with 2D keypoint annotations. Following Tulsiani and Malik (2015), we use non-occluded and non-truncated objects for testing. We use the standard R-CNN (Girshick et al 2014) for object detection, and our 3D-INN for viewpoint estimation.

Table 3 shows that 3D-INN is comparable with Viewpoints and Keypoints (V&K by Tulsiani and Malik (2015)), and outperforms other algorithms with a significant margin. Both 3D-INN and V&K use R-CNN (Girshick et al 2014) for object detection (we use the R-CNN detection results provided by Tulsiani and Malik (2015)); this rules out the influence of object detectors. Further, while all the other algorithms are trained on either PASCAL VOC or PASCAL 3D+, ours is trained on Keypoint-5 or ImageNet. This indi-

cates our learned model transfers well across datasets, without suffering much from the domain adaptation issue.

Qualitative results on benchmarks. We now show qualitative results on Keypoint-5, IKEA, the SUN database (Xiao et al 2010), and the PASCAL 3D+ dataset (Xiang et al 2014) in Figure 10. When the image is clean and objects are not occluded, our algorithm can recover 3D object structure and viewpoint with high accuracy. Fine-tuning further helps to improve the results (see chairs at row 1 column 1, and row 4 column 1). Our algorithm is also robust to partial occlusion, demonstrated by the IKEA bed at row 5 column 1. We show failure cases in the last column: one major failure case is when the object is heavily cropped in the input image (the last column, row 4 to 7), as the 3D object skeleton becomes hard to infer. Figure 11 shows more results on car structure recovery.

When 3D-INN is used in combination with detection models, it needs to deal with imperfect detection results. Here, we also evaluate 3D-INN on noisy input, specifically, on images with an object from a different but similar category. Figure 12 shows the recovered 3D structures of chairs using a model trained either on sofas or beds. In most cases 3D-INN still provides reasonable output, and the mistakes are mostly due to the difference between training and test sets, *e.g.*, the model trained on beds does not perform well on chairs facing backward, because there are almost no beds with a similar viewpoint in the training set.

At the end of the manuscript, we supply more results on chair and sofa images randomly sampled from the test set


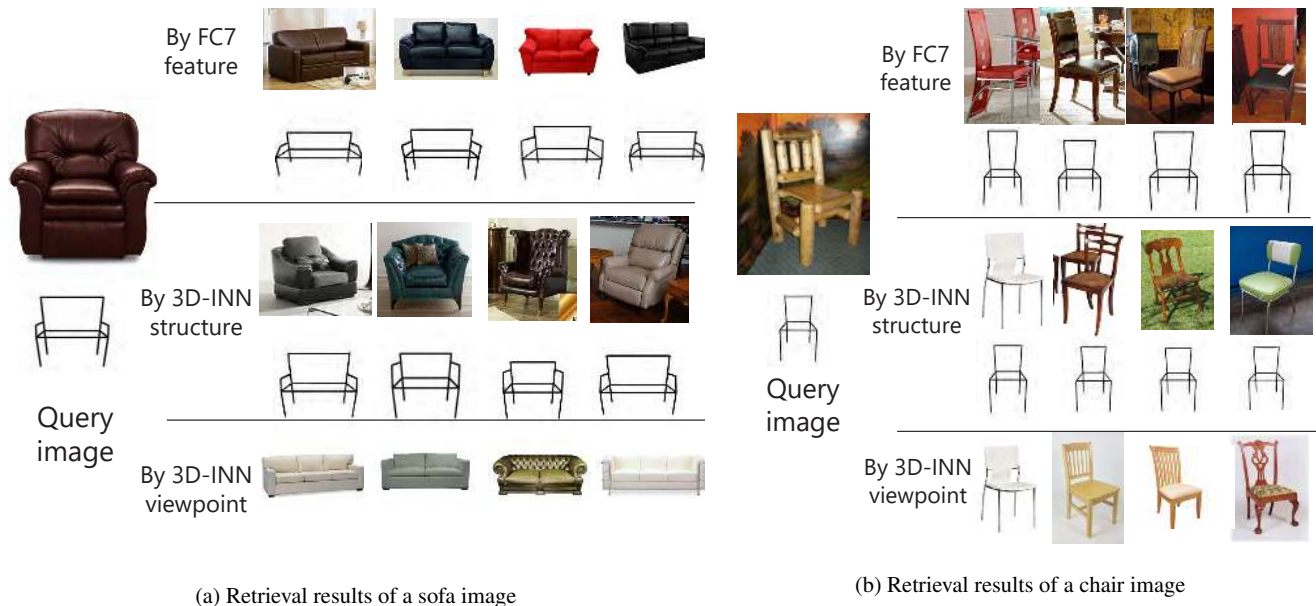


Fig. 13: Visualization of 3D reconstruction results. We render objects using *Blender*.



(a) Retrieval results of a sofa image

(b) Retrieval results of a chair image

Fig. 14: Retrieval results for a sofa (a) and a chair (b) in different feature spaces. 3D-INN helps to retrieve objects with similar 3D structure or pictured in a similar viewpoint.

of Keypoint-5. Figure 16 and Figure 17 show the estimated skeletons for chairs and sofas, respectively.

5 Applications

The inferred latent parameters, as a compact and informative representation of objects in images, have wide applications. In this section, we demonstrate representative ones including 3D object rendering, image retrieval, and object graph visualization.

3D Object Rendering Given an estimated 3D object structure, we can render it in a 3D graphics engine like Blender, as shown in Figure 13.

Image Retrieval Using estimated 3D structural and viewpoint information, we can retrieve images based on their 3D configurations. Figure 14 shows image retrieval results using FC7 features from AlexNet (Krizhevsky et al 2012) and using the 3D structure and viewpoint learned by 3D-INN. Our retrieval database includes all testing images of chairs and sofas in Keypoint-5. In each row, we sort the best matches of the query image, measured by Euclidean distance in a



Fig. 15: Object graph visualization based on learned object representations: we visualize images using t-SNE (Van der Maaten and Hinton 2008) on 3D viewpoints predicted by 3D-INN.

specific feature space. We retrieve images in two ways: *by structure* uses estimated internal structural parameters ($\{\alpha_i\}$ in Equation 2), and *by viewpoint* uses estimated external viewpoint parameters (R in Equation 2).

Object Graph Similar to the retrieval task, we visualize all test images for chairs in Keypoint-5 in Figure 15, using t-SNE (Van der Maaten and Hinton 2008) on estimated 3D viewpoints. Note the smooth transition from the chairs facing left to those facing right.

6 Discussions

We have introduced 3D INterpreter Networks (3D-INN). From a single image, our model recovers the 2D keypoints and 3D structure of a (possibly deformable) object, as well as camera parameters. To achieve this goal, we used 3D skeletons as an abstract 3D representation, incorporated a projection layer to the network for learning 3D parameters from 2D labels, and employed keypoint heatmaps to connect real and synthetic data. Empirically, we showed that 3D-INN performs well on both 2D keypoint estimation and 3D structure and viewpoint recovery, comparable to or better than the state of the art. Further, various applications demonstrated the potential of the skeleton representation learned by 3D-INN.

We choose to model objects via 3D skeletons and the corresponding 2D keypoints, as opposed to other dense 3D representations such as voxels, meshes, and point clouds, because skeletons offer unique advantages. First, given an RGB image, its sparse 2D annotations like keypoints are easier and more affordable to acquire, and can be used as 2D supervision for 3D skeleton and viewpoint estimation; in

comparison, it is prohibitively challenging to obtain dense annotations like a depth map to constrain 3D reconstructions in voxels or meshes. Second, the employed base shapes carry rich category-specific shape priors, with which 3D-INN can encode an object skeleton with a few parameters. This feature is particularly useful on platforms with severe memory and computational constraints, such as on autonomous cars and on mobile phones.

That being said, skeletons have their own limitations. The most significant is on its generalization power: there are many real-world objects whose keypoints are hard to define, such as trees, flowers, and deformable shapes like ropes; in those cases, there lacks a straightforward way to apply 3D-INN to model these objects. Recent research on 3D reconstruction via richer, generic intermediate representations like intrinsic images (Barrow and Tenenbaum 1978) suggests a potential solution to the problem, though as discussed above it is much harder to obtain annotated intrinsic images, compared to keypoints (Wu et al 2017).

In this work, we focus on single-view 3D reconstruction. As discussed in Section 1, requiring only a single image as input has unique practical advantages, in addition to its scientific value. First, our algorithm can be directly applied to cases where only in-the-wild images are available, not multi-view images or videos. Second, taking a single image as input enables online inference and therefore fits real-time applications; in contrast, most multi-view reconstruction algorithms are offline. It is also possible to our 3D-INN to use multi-view data when they are available (Kar et al 2017), and more generally, to integrate viewer-centered and object-centered representations in a principled manner (Hinton 1981).



Fig. 16: Estimated 3D skeletons on more Keypoint-5 chair images. Images are randomly sampled from the test set.



Fig. 17: Estimated 3D skeletons on more Keypoint-5 sofa images. Images are randomly sampled from the test set.

3D-INN estimates the 3D skeleton and pose of an object from an RGB image, and can therefore be applied to

the enormous existing RGB data. But we are also aware that depth sensors have recently become affordable to end

users (Newcombe et al 2011), and large-scale RGB-D datasets are being built (Song et al 2017; McCormac et al 2017). Depth data help to resolve the ambiguity in the projection from 3D shapes to 2D images, allow object structure prediction in metric scale, and enable wide applications (Chen et al 2012). Hence, a promising future research topic would be to extend the current framework to handle depth data, while enforcing the 2D-3D differentiable consistencies in various forms (Tulsiani et al 2017; Wu et al 2017).

Acknowledgement

This work is supported by NSF Robust Intelligence 1212849 and NSF Big Data 1447476 to W.F., NSF Robust Intelligence 1524817 to A.T., ONR MURI N00014-16-1-2007 to J.B.T., Shell Research, the Toyota Research Institute, and the Center for Brain, Minds and Machines (NSF STC award CCF-1231216). The authors would like to thank Nvidia for GPU donations. Part of this work was done when Jiajun Wu was an intern at Facebook AI Research, and Tianfan Xue was a graduate student at MIT CSAIL.

References

- Akhter I, Black MJ (2015) Pose-conditioned joint angle limits for 3d human pose reconstruction. In: IEEE Conference on Computer Vision and Pattern Recognition 3
- Aubry M, Maturana D, Efros A, Russell B, Sivic J (2014) Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In: IEEE Conference on Computer Vision and Pattern Recognition 2, 3
- Bansal A, Russell B (2016) Marr revisited: 2d-3d alignment via surface normal prediction. In: IEEE Conference on Computer Vision and Pattern Recognition 3
- Barrow HG, Tenenbaum JM (1978) Recovering intrinsic scene characteristics from images. *Computer vision systems* 14
- Belhumeur PN, Jacobs DW, Kriegman DJ, Kumar N (2013) Localizing parts of faces using a consensus of exemplars. *IEEE Transactions on Pattern Analysis and Machine intelligence* 35(12):2930–2940 7
- Bever TG, Poeppel D (2010) Analysis by synthesis: a (re-) emerging program of research for language and vision. *Biolinguistics* 4(2-3):174–200 3
- Bourdev L, Maji S, Brox T, Malik J (2010) Detecting people using mutually consistent poselet activations. In: *European Conference on Computer Vision* 7
- Carreira J, Agrawal P, Fragkiadaki K, Malik J (2016) Human pose estimation with iterative error feedback. In: *IEEE Conference on Computer Vision and Pattern Recognition* 3
- Chen J, Izadi S, Fitzgibbon A (2012) Kinètre: animating the world with the human body. In: *ACM symposium on User Interface Software and Technology* 16
- Choy CB, Xu D, Gwak J, Chen K, Savarese S (2016) 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: *European Conference on Computer Vision* 3
- Dosovitskiy A, Tobias Springenberg J, Brox T (2015) Learning to generate chairs with convolutional neural networks. In: *IEEE Conference on Computer Vision and Pattern Recognition* 2, 4
- Fidler S, Dickinson SJ, Urtasun R (2012) 3d object detection and view-point estimation with a deformable 3d cuboid model. In: *Advances in Neural Information Processing Systems* 4
- Girshick R, Donahue J, Darrell T, Malik J (2014) Rich feature hierarchies for accurate object detection and semantic segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition* 3, 10, 12
- Hejrati M, Ramanan D (2012) Analyzing 3d objects in cluttered images. In: *Advances in Neural Information Processing Systems* 6
- Hejrati M, Ramanan D (2014) Analysis by synthesis: 3d object recognition by object reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition* 3
- Hinton GE, Ghahramani Z (1997) Generative models for discovering sparse distributed representations. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 352(1358):1177–1190 3
- Hinton GF (1981) A parallel computation that assigns canonical object-based frames of reference. In: *International Joint Conference on Artificial Intelligence* 14
- Hu W, Zhu SC (2015) Learning 3d object templates by quantizing geometry and appearance spaces. *IEEE Transactions on Pattern Analysis and Machine intelligence* 37(6):1190–1205 3
- Huang Q, Wang H, Koltun V (2015) Single-view reconstruction via joint analysis of image and shape collections. *ACM Transactions on Graphics* 34(4):87 2, 3, 4
- Jaderberg M, Simonyan K, Zisserman A, Kavukcuoglu K (2015) Spatial transformer networks. In: *Advances in Neural Information Processing Systems* 6
- Kar A, Tulsiani S, Carreira J, Malik J (2015) Category-specific object reconstruction from a single image. In: *IEEE Conference on Computer Vision and Pattern Recognition* 2, 3, 4
- Kar A, Häne C, Malik J (2017) Learning a multi-view stereo machine. In: *Advances in Neural Information Processing Systems* 14
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems* 13
- Kulkarni TD, Kohli P, Tenenbaum JB, Mansinghka V (2015a) Picture: A probabilistic programming language for scene perception. In: *IEEE Conference on Computer Vision and Pattern Recognition* 3
- Kulkarni TD, Whitney WF, Kohli P, Tenenbaum JB (2015b) Deep convolutional inverse graphics network. In: *Advances in Neural Information Processing Systems* 3
- Leclerc YG, Fischler MA (1992) An optimization-based approach to the interpretation of single line drawings as 3d wire frames. *International Journal of Computer Vision* 9(2):113–136 3
- Li Y, Su H, Qi CR, Fish N, Cohen-Or D, Guibas LJ (2015) Joint embeddings of shapes and images via cnn image purification. *ACM Transactions on Graphics* 34(6):234 2
- Lim JJ, Pirsiavash H, Torralba A (2013) Parsing ikea objects: Fine pose estimation. In: *IEEE International Conference on Computer Vision* 3, 8, 9, 10
- Lim JJ, Khosla A, Torralba A (2014) FPM: Fine pose parts-based model with 3d cad models. In: *European Conference on Computer Vision* 4
- Liu J, Belhumeur PN (2013) Bird part localization using exemplar-based models with enforced pose and subcategory consistency. In: *IEEE International Conference on Computer Vision* 7, 8
- Lowe DG (1987) Three-dimensional object recognition from single two-dimensional images. *Artificial intelligence* 31(3):355–395 3
- Van der Maaten L, Hinton G (2008) Visualizing data using t-sne. *Journal of Machine Learning Research* 9(11):2579–2605 14
- McCormac J, Handa A, Leutenegger S, Davison AJ (2017) Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation. In: *IEEE International Conference on Computer Vision* 4, 16

- Newcombe RA, Izadi S, Hilliges O, Molyneaux D, Kim D, Davison AJ, Kohi P, Shotton J, Hodges S, Fitzgibbon A (2011) Kinect-fusion: Real-time dense surface mapping and tracking. In: IEEE International Symposium on Mixed and Augmented Reality, pp 127–136 [16](#)
- Newell A, Yang K, Deng J (2016) Stacked hourglass networks for human pose estimation. In: European Conference on Computer Vision [4](#)
- Pepik B, Stark M, Gehler P, Schiele B (2012) Teaching 3d geometry to deformable part models. In: IEEE Conference on Computer Vision and Pattern Recognition [10](#)
- Prasad M, Fitzgibbon A, Zisserman A, Van Gool L (2010) Finding nemo: Deformable object class modelling using curve matching. In: IEEE Conference on Computer Vision and Pattern Recognition [2, 3](#)
- Ramakrishna V, Kanade T, Sheikh Y (2012) Reconstructing 3d human pose from 2d image landmarks. In: European Conference on Computer Vision [3](#)
- Ren S, He K, Girshick R, Sun J (2015) Faster R-CNN: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems [6](#)
- Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, et al (2015) Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252 [1, 12](#)
- Sapp B, Taskar B (2013) Modec: Multimodal decomposable models for human pose estimation. In: IEEE Conference on Computer Vision and Pattern Recognition [3, 7, 8](#)
- Satkin S, Lin J, Hebert M (2012) Data-driven scene understanding from 3D models. In: British Machine Vision Conference [3](#)
- Shakhnarovich G, Viola P, Darrell T (2003) Fast pose estimation with parameter-sensitive hashing. In: IEEE International Conference on Computer Vision [4](#)
- Shih KJ, Mallya A, Singh S, Hoiem D (2015) Part localization using multi-proposal consensus for fine-grained categorization. In: British Machine Vision Conference [3, 7, 8](#)
- Shrivastava A, Gupta A (2013) Building part-based object detectors via 3d geometry. In: IEEE International Conference on Computer Vision [3](#)
- Soltani AA, Huang H, Wu J, Kulkarni TD, Tenenbaum JB (2017) Synthesizing 3d shapes via modeling multi-view depth maps and silhouettes with deep generative networks. In: IEEE Conference on Computer Vision and Pattern Recognition [3](#)
- Song S, Yu F, Zeng A, Chang AX, Savva M, Funkhouser T (2017) Semantic scene completion from a single depth image. In: IEEE Conference on Computer Vision and Pattern Recognition [16](#)
- Su H, Huang Q, Mitra NJ, Li Y, Guibas L (2014) Estimating image depth using shape collections. *ACM Transactions on Graphics* 33(4):37 [2, 3, 4](#)
- Su H, Qi CR, Li Y, Guibas L (2015) Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In: IEEE International Conference on Computer Vision [2, 3, 4, 9, 10, 11, 12](#)
- Sun B, Saenko K (2014) From virtual to reality: Fast adaptation of virtual object detectors to real domains. In: British Machine Vision Conference [4](#)
- Taigman Y, Yang M, Ranzato M, Wolf L (2015) Web-scale training for face identification. In: IEEE Conference on Computer Vision and Pattern Recognition [6](#)
- Tompson J, Goroshin R, Jain A, LeCun Y, Bregler C (2015) Efficient object localization using convolutional networks. In: IEEE Conference on Computer Vision and Pattern Recognition [3, 5, 7, 8](#)
- Tompson JJ, Jain A, LeCun Y, Bregler C (2014) Joint training of a convolutional network and a graphical model for human pose estimation. In: Advances in Neural Information Processing Systems [5, 7](#)
- Torralba A, Efros AA (2011) Unbiased look at dataset bias. In: IEEE Conference on Computer Vision and Pattern Recognition [10](#)
- Torresani L, Hertzmann A, Bregler C (2003) Learning non-rigid 3d shape from 2d motion. In: Advances in Neural Information Processing Systems [2, 3, 4, 6](#)
- Toshev A, Szegegy C (2014) Deeppose: Human pose estimation via deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 1653–1660 [3, 7](#)
- Tulsiani S, Malik J (2015) Viewpoints and keypoints. In: IEEE Conference on Computer Vision and Pattern Recognition [3, 4, 10, 12](#)
- Tulsiani S, Zhou T, Efros AA, Malik J (2017) Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: IEEE Conference on Computer Vision and Pattern Recognition [16](#)
- Vicente S, Carreira J, Agapito L, Batista J (2014) Reconstructing pascal voc. In: IEEE Conference on Computer Vision and Pattern Recognition [2, 3](#)
- Wah C, Branson S, Welinder P, Perona P, Belongie S (2011) The Caltech-UCSD Birds-200-2011 Dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology [3, 7](#)
- Wu J, Yildirim I, Lim JJ, Freeman B, Tenenbaum J (2015) Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. In: Advances in Neural Information Processing Systems [3](#)
- Wu J, Zhang C, Xue T, Freeman WT, Tenenbaum JB (2016) Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: Advances in Neural Information Processing Systems [3](#)
- Wu J, Wang Y, Xue T, Sun X, Freeman WT, Tenenbaum JB (2017) Marnet: 3d shape reconstruction via 2.5d sketches. In: Advances in Neural Information Processing Systems [3, 14, 16](#)
- Xiang Y, Mottaghi R, Savarese S (2014) Beyond pascal: A benchmark for 3d object detection in the wild. In: IEEE Winter Conference on Applications of Computer Vision [3, 4, 9, 10, 12](#)
- Xiao J, Hays J, Ehinger K, Oliva A, Torralla A (2010) Sun database: Large-scale scene recognition from abbey to zoo. In: IEEE Conference on Computer Vision and Pattern Recognition [9, 12](#)
- Xue T, Liu J, Tang X (2012) Example-based 3d object reconstruction from line drawings. In: IEEE Conference on Computer Vision and Pattern Recognition [3](#)
- Yang Y, Ramanan D (2011) Articulated pose estimation with flexible mixtures-of-parts. In: IEEE Conference on Computer Vision and Pattern Recognition [3, 7](#)
- Yasin H, Iqbal U, Krüger B, Weber A, Gall J (2016) A dual-source approach for 3d pose estimation from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition [3, 4](#)
- Yuille A, Kersten D (2006) Vision as bayesian inference: analysis by synthesis? *Trends in cognitive sciences* 10(7):301–308 [3](#)
- Zeng A, Song S, Nießner M, Fisher M, Xiao J (2016) 3dmatch: Learning the matching of local 3d geometry in range scans. In: IEEE Conference on Computer Vision and Pattern Recognition [3](#)
- Zhou T, Krähenbühl P, Aubry M, Huang Q, Efros AA (2016) Learning dense correspondence via 3d-guided cycle consistency. In: IEEE Conference on Computer Vision and Pattern Recognition [4](#)
- Zhou X, Leonardos S, Hu X, Daniilidis K (2015) 3d shape reconstruction from 2d landmarks: A convex formulation. In: IEEE Conference on Computer Vision and Pattern Recognition [2, 3, 8, 9, 10](#)
- Zia MZ, Stark M, Schiele B, Schindler K (2013) Detailed 3d representations for object recognition and modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(11):2608–2623 [3, 4](#)