

# Cuboids Revisited: Learning Robust 3D Shape Fitting to Single RGB Images

Florian Kluger<sup>1</sup>, Hanno Ackermann<sup>1</sup>, Eric Brachmann<sup>2</sup>, Michael Ying Yang<sup>3</sup>, Bodo Rosenhahn<sup>1</sup>

<sup>1</sup>Leibniz University Hannover, <sup>2</sup>Niantic, <sup>3</sup>University of Twente

## Abstract

Humans perceive and construct the surrounding world as an arrangement of simple parametric models. In particular, man-made environments commonly consist of volumetric primitives such as cuboids or cylinders. Inferring these primitives is an important step to attain high-level, abstract scene descriptions. Previous approaches directly estimate shape parameters from a 2D or 3D input, and are only able to reproduce simple objects, yet unable to accurately parse more complex 3D scenes. In contrast, we propose a robust estimator for primitive fitting, which can meaningfully abstract real-world environments using cuboids. A RANSAC estimator guided by a neural network fits these primitives to 3D features, such as a depth map. We condition the network on previously detected parts of the scene, thus parsing it one-by-one. To obtain 3D features from a single RGB image, we additionally optimise a feature extraction CNN in an end-to-end manner. However, naively minimising point-to-primitive distances leads to large or spurious cuboids occluding parts of the scene behind. We thus propose an occlusion-aware distance metric correctly handling opaque scenes. The proposed algorithm does not require labour-intensive labels, such as cuboid annotations, for training. Results on the challenging NYU Depth v2 dataset demonstrate that the proposed algorithm successfully abstracts cluttered real-world 3D scene layouts.

## 1. Introduction

Humans tend to create using simple geometric forms. For instance, a house is made from bricks and squared timber, and a book is a cuboidal assembly of rectangles. Consequently, it appears that humans also visually abstract environments by decomposing them into arrangements of cuboids, cylinders, ellipsoids and other simple volumetric primitives [5]. Such an abstraction of the 3D world is also very useful for machines with visual perception. Scene representation based on geometric shape primitives has been an active topic since the very beginning of Computer Vision. In 1963, *Blocks World* [50] from Larry Roberts was one of the earliest approaches for qualitative 3D shape recovery from 2D images using generic shape primitives. In

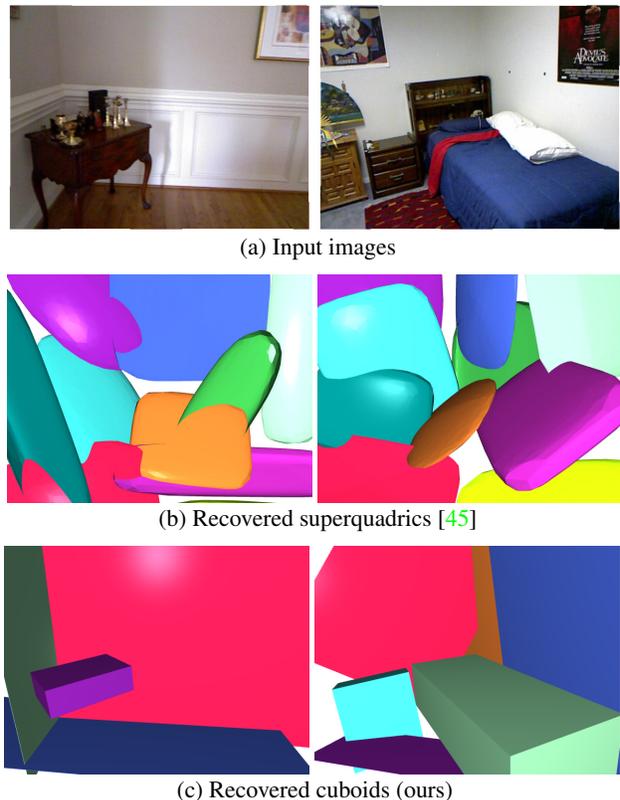


Figure 1: **Primitive-based Scene Abstractions:** We parse images of real-world scenes (a) in order to generate abstractions of their 3D structure using cuboids (c). Our method is capable of capturing scene structure more accurately than previous work [45] based on superquadrics (b).

recent years, with rapid advances in the field of deep learning, high-quality 3D reconstruction from single images has become feasible. Most approaches recover 3D information such as depth [13] and meshes [56] from RGB images. Fewer works consider more parsimonious 3D shape descriptions such as cuboids [54] or superquadrics [45, 44]. These 3D shape parsers work well for isolated objects, but do not generalise to complex real-world scenes (cf. Fig. 1).

Robust model fitting algorithms such as RANSAC [14] and its many derivatives [2, 9, 47] have been used to fit

low-dimensional parametric models, such as plane homographies, fundamental matrices or geometric primitives, to real-world noisy data. Trainable variants of RANSAC [6, 7, 28] use a neural network to predict sampling weights from data. They require fewer samples and are more accurate.

Leveraging advances in the fields of single image 3D reconstruction and robust multi-model fitting, we present a novel approach for robustly parsing real-world scenes using 3D shape primitives, such as cuboids. A trainable RANSAC estimator fits these primitives to 3D features, such as a depth map. We build upon the estimator proposed in [28], and extend it by predicting multiple sets of RANSAC sampling weights concurrently. This enables our method to distinguish between different structures in a scene more easily. We obtain 3D features from a single RGB image using a CNN, and show how to optimise this CNN in an end-to-end manner. Our training objective is based on geometrical consistency with readily available 3D sensory data.

During primitive fitting, a naive maximisation of inlier counts considering point-to-primitive distances causes the algorithm to detect few but excessively large models. We argue that this is due to parts of a primitive surface correctly representing some parts of a scene, while other parts of the same primitive wrongly occlude other parts of the scene. Thus, points should not be assigned to primitive surfaces which cannot be seen by the camera due to occlusion or self-occlusion. We therefore propose an occlusion-aware distance and a corresponding occlusion-aware inlier count.

As no closed-form solution exists to calculate cuboid parameters, we infer them by numerical optimisation. However, backpropagation through this optimisation is numerically unstable and computationally costly. We therefore analytically derive the gradient of primitive parameters w.r.t. the features used to compute them. Our gradient computation allows for end-to-end training without backpropagation through the minimal solver itself. We demonstrate the efficacy of our method on the challenging real-world NYU Depth v2 dataset [52].

In summary, our **contributions**<sup>1</sup> are as follows:

- A 3D scene parser which can process more complex real-world scenes than previous works on 3D scene abstraction.
- An occlusion-aware distance metric for opaque scenes.
- Analytical derivation of the gradient of cuboids w.r.t. input features, in order to circumvent infeasible backpropagation through our minimal solver, thus enabling end-to-end training.
- Our method does not require labour-intensive labels, such as cuboid or object annotations, and can be trained on readily available sensory data instead.

---

<sup>1</sup>Source code is available at: [https://github.com/fkluger/cuboids\\_revisited](https://github.com/fkluger/cuboids_revisited)

## 2. Related Work

**Monocular Depth Estimation.** Monocular depth estimation has become feasible with advances in deep learning and progressed steadily over recent years. Instead of fully supervised learning [13, 35, 30], later works approach the problem via semi-supervised learning [29], unsupervised learning [15, 38, 18, 58] or self-supervised learning [19, 36]. These methods predict dense 3D information instead of the more parsimonious primitive based descriptions we are interested in. However, we show how to leverage monocular depth estimation to this end.

**Single Image 3D Reconstruction.** The authors of [33, 32] propose methods for piece-wise planar reconstruction from single images. Older works [10] rely on geometric image features, such as vanishing points [26] or horizon lines [27]. Like depth estimation, however, these methods can only describe the visible surfaces of a scene and do not capture volumetric characteristics. Although some methods [43, 20, 56, 39] perform 3D mesh reconstruction, they demonstrably only work on images of isolated single objects. The authors of [41, 59, 17] present mesh reconstructions for multiple objects in real-world scenarios. However, the latter methods require ground-truth object meshes and object class labels for training, which are both costly to obtain. Our approach, on the other hand, requires neither.

**3D Bounding Box Regression.** For predefined object classes, 3D bounding box regression is a well investigated topic in 3D object detection for RGB images [49, 12, 24, 40], RGB-D images [48], and 3D point clouds [51, 53]. These works not only detect bounding boxes but also classify objects. During training, target 3D boxes as well as category labels are needed. In contrast, our proposed method does not require such annotations.

**Robust Multi-Model Fitting.** Robust fitting of parametric models, for instance using RANSAC [14], is a key problem in Computer Vision. Several approaches [3, 4, 46, 1] fit *multiple instances* of a model simultaneously by optimising an energy-based functional. Moreover, the algorithms in [3, 4, 37] approach the *multi-class* problem, *i.e.* when models of multiple types may fit the data. Sampling efficiency is improved in [6, 7, 28, 47]. In this work, we extend [28] for robustly parsing real-world scenes using 3D shape primitives. Unlike [28], we not just learn the parameters of the sampling weight from data, but train it conjointly with a depth estimation network in an end-to-end manner.

**3D Primitive Fitting.** Although an old topic in Computer Vision (cf. Blocks World [50]), 3D shape recovery from 2D images using volumetric primitives is still considered non-trivial [21]. The methods of [57] and [22, 31] localise cuboids in images and RGB-D data, respectively. While [22, 31] use the challenging NYU Depth v2 dataset [52], all three works require ground truth cuboids.

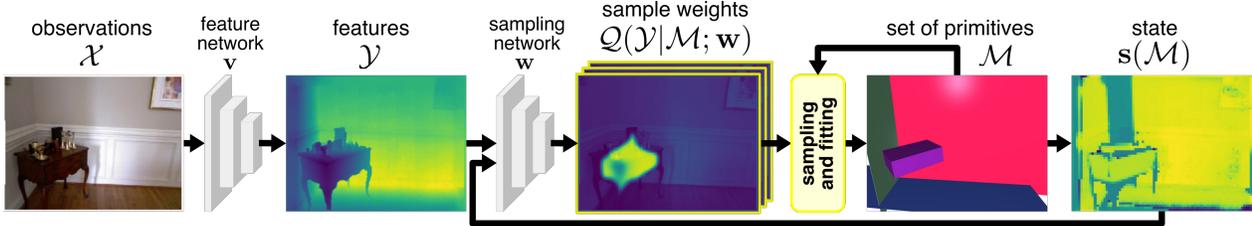


Figure 2: **Overview:** Given observations  $\mathcal{X}$  (RGB image), we predict 3D features  $\mathcal{Y}$  (depth map) using a neural network with parameters  $\mathbf{v}$ . Conditioned on a state  $\mathbf{s}$ , a second neural network with parameters  $\mathbf{w}$  predicts sampling weights  $p(\mathbf{y}|\mathbf{s}; \mathbf{w}) \in \mathcal{Q}$  for each feature  $\mathbf{y} \in \mathcal{Y}$ . Using these weights, a RANSAC-based estimator samples minimal sets of features, and generates primitive (cuboid) hypotheses  $\mathcal{H}$ . It selects the best hypothesis  $\hat{\mathbf{h}} \in \mathcal{H}$  and appends it to the set of previously recovered primitives  $\mathcal{M}$ . We update the state  $\mathbf{s}$  based on  $\mathcal{M}$  and repeat the process in order to recover all primitives step-by-step.

The works most related to ours are [54, 45, 44]. In [54], the authors propose a method for 3D shape abstraction using cuboids. It is based on a neural network, which directly regresses cuboid parameters from either an RGB image or a 3D input, and is trained in an unsupervised fashion. Similarly, the authors of [45] decompose objects into sets of superquadric surfaces. They extend this to hierarchical sets of superquadrics in [44]. However, these approaches are only evaluated on relatively simple shapes, such as those in ShapeNet [8]. While they yield good results for such objects, they are unable to predict reasonable primitive decompositions for complex real-world scenes (cf. Sec. 4).

### 3. Method

We formulate the problem of abstracting a target scene  $\mathcal{Z}$  by fitting a set of shape primitives  $\mathcal{M}$  to features  $\mathcal{Y}$ , which may either be provided directly, or extracted from observations  $\mathcal{X}$ . Here,  $\mathcal{Z}$  represents the ground truth depth or its corresponding point cloud, and  $\mathcal{X}$  is an RGB image of the scene aligned with  $\mathcal{Z}$ . Features  $\mathcal{Y}$  are either equal to  $\mathcal{Z}$ , or estimated via a function  $\mathcal{Y} = f_{\mathbf{v}}(\mathcal{X})$  if  $\mathcal{Z}$  is unknown. We implement function  $f_{\mathbf{v}}$  as a neural network with parameters  $\mathbf{v}$ . Each primitive  $\mathbf{h} \in \mathcal{M}$  is a cuboid with variable size and pose. See Fig. 2 for an example of  $\mathcal{X}$ ,  $\mathcal{Y}$  and  $\mathcal{M}$ .

For primitive fitting, we build upon the robust multi-model estimator of Kluger et al. [28]. This estimator predicts sampling weights  $\mathbf{p} = f_{\mathbf{w}}(\mathcal{Y}, \mathbf{s})$  from observations  $\mathcal{Y}$  and a state  $\mathbf{s}$  via a neural network with parameters  $\mathbf{w}$ , which are learnt from data. It samples minimal sets of features from  $\mathcal{Y}$  according to  $\mathbf{p}$ , and fits primitive hypotheses  $\mathcal{H}$  via a minimal solver  $f_h$ . From these hypotheses, it selects the best primitive  $\hat{\mathbf{h}} \in \mathcal{H}$  according to an inlier criterion, and adds it to the current set of primitives  $\mathcal{M}$ . Based on  $\mathcal{M}$ , it then updates the state  $\mathbf{s}$  and predicts new sampling weights  $\mathbf{p}$  in order to sample and select the next primitive. This process, as visualised in Fig. 2, is repeated until all primitives have been found one by one. Unlike [28], we learn the parameters of the sampling weight predictor  $f_{\mathbf{w}}$  jointly with the feature extractor  $f_{\mathbf{v}}$  in an end-to-end manner. We show how to achieve this despite backpropagation through the

minimal solver  $f_h$  being numerically unstable and computationally costly. In addition, we generalise  $f_{\mathbf{w}}$  so that it predicts multiple sets of sampling weights  $\mathcal{Q} = \{\mathbf{p}_1, \dots, \mathbf{p}_Q\}$  concurrently, which enables it to distinguish between different primitive instances more effectively.

The scenes we are dealing with in this work have been captured with an RGB-D camera. We therefore do not have full 3D shapes available as ground-truth for the scenes. Instead, we only have 2.5D information, *i.e.* 3D information for visible parts of the scene, and no information about the occluded parts. If not taken into account, this fact can lead to spurious, oversized or ill-fitting primitives being selected, as visualised in Fig. 3. For this reason we present occlusion-aware distance and inlier metrics.

#### 3.1. Feature Extraction

In order to fit 3D shapes such as cuboids to an RGB image  $\mathcal{X}$ , we have to extract 3D features  $\mathcal{Y}$  from  $\mathcal{X}$ . We employ a depth estimator  $f_{\mathbf{v}}$  which gives us the desired features  $\mathcal{Y} = f_{\mathbf{v}}(\mathcal{X})$  in form of a pixel-wise depth map. The depth estimator is realised as a convolutional neural network with parameters  $\mathbf{v}$ . We then convert  $\mathcal{Y}$  into a point cloud via backprojection using known camera intrinsics  $\mathbf{K}$ .

#### 3.2. Cuboid Parametrisation

A generic cuboid is described by its shape  $(a_x, a_y, a_z)$  and pose  $(\mathbf{R}, \mathbf{t})$ . The shape corresponds to its width, height and length in a cuboid-centric coordinate system, while its pose translates the latter into a world coordinate system. We represent the rotation  $\mathbf{R}$  in angle-axis notation  $\mathbf{r} = \theta \mathbf{u}$ . Each cuboid thus has nine degrees of freedom, and we require minimal sets of  $C = 9$  points to estimate them.

##### 3.2.1 Point-to-Cuboid Distance

When computing the distance between point  $\mathbf{y} = (x, y, z)^T$  and cuboid  $\mathbf{h} = (a_x, a_y, a_z, \mathbf{R}, \mathbf{t})$ , we first translate  $\mathbf{y}$  into the cuboid-centric coordinate frame:  $\hat{\mathbf{y}} = \mathbf{R}(\mathbf{y} - \mathbf{t})$ . We then compute its squared distance to the cuboid surface:

$$d(\mathbf{h}, \mathbf{y})^2 = \max(\min(a_x - |\hat{x}|, a_y - |\hat{y}|, a_z - |\hat{z}|, 0)^2 + \max(|\hat{x}| - a_x, 0)^2 + \max(|\hat{y}| - a_y, 0)^2 + \max(|\hat{z}| - a_z, 0)^2).$$

Similarly, we can compute the distances to any of the six individual sides of the cuboid. Defining, for example, the plane orthogonal to the  $x$ -axis in its positive direction as the first side, we define the distance of a point  $\mathbf{y}$  to it as:

$$d_1^p(\mathbf{h}, \mathbf{y})^2 = (\hat{x} - a_x)^2 + \max(|\hat{y}| - a_y, 0)^2 + \max(|\hat{z}| - a_z, 0)^2. \quad (1)$$

Distances  $d_2^p, \dots, d_6^p$  to the other sides of the cuboid are calculated accordingly.

### 3.2.2 Occlusion Handling

When dealing with 2.5D data which only represents visible parts of the scene, simply using minimal point-to-cuboid distances as in Sec. 3.2.1 is not adequate. Fig. 3 gives an intuitive example: The mean distance of all points to their closest surface of either cuboid A or B is the same. However, the visible surfaces of cuboid B occlude all points while not representing any structure present in the scene. Cuboid A, on the other hand, does not occlude any points, and its visible surfaces fit very well to an existing structure. Although other parts of its surface do not represent any structure either, they are self-occluded and thus of no interest. So, cuboid A is a much better fit than cuboid B.

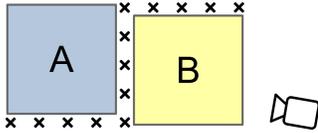


Figure 3: **Occlusion:** Given are a point cloud (X), two cuboids (A and B) and a camera observing the scene. Cuboid A is a better fit since it does not occlude any points.

**Occlusion Detection.** In order to detect whether a cuboid  $\mathbf{h}$  occludes a point  $\mathbf{y}$  from the perspective of a camera with centre  $\mathbf{c}$ , we must first translate  $\mathbf{c}$  into the cuboid-centric coordinate frame. We parametrise the line of sight for a point  $\mathbf{y}$  and determine its intersections with each of the six cuboid planes. Via this procedure we define the following indicator function:

$$\chi_o(\mathbf{h}, \mathbf{y}, i) = \begin{cases} 1 & \text{if } i\text{-th plane of } \mathbf{h} \text{ occludes } \mathbf{y}, \\ 0 & \text{else, with } i \in \{1, \dots, 6\}. \end{cases} \quad (2)$$

**Occlusion-Aware Point-to-Cuboid Distance.** In order to correctly evaluate scenarios with occlusions, we propose an *occlusion-aware* point-to-cuboid distance: Given a point  $\mathbf{y}$  and set of cuboids  $\mathcal{M} = \{\mathbf{h}_1, \dots, \mathbf{h}_{|\mathcal{M}|}\}$ , we compute its distance to the most distant occluding surface. In other words, we need to know the minimal distance  $\mathbf{y}$  would have to travel in order to become visible:

$$d_o(\mathcal{M}, \mathbf{y}) = \max_{\mathbf{h} \in \mathcal{M}, i \in \{1, \dots, 6\}} (\chi_o(\mathbf{h}, \mathbf{y}, i) \cdot d_i^p(\mathbf{h}, \mathbf{y})).$$

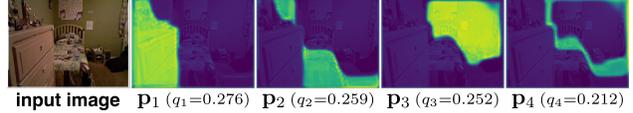


Figure 4: **Sampling Weights:** For each image, we predict multiple sets of sampling weights  $\mathcal{Q} = \{\mathbf{p}_1, \dots, \mathbf{p}_Q\}$  and corresponding selection probabilities  $\mathbf{q} = [q_1, \dots, q_Q]$ . In this example, the first three sampling weight sets roughly cover distinct parts of the scene. The fourth set  $\mathbf{p}_4$  does not, but also has the lowest selection probability.

If  $\mathbf{y}$  is not occluded at all, this distance naturally becomes zero. We hence define the occlusion-aware distance of a point to a set of cuboids as:

$$d_{oa}(\mathcal{M}, \mathbf{y}) = \max \left( \min_{\mathbf{h} \in \mathcal{M}} d(\mathbf{h}, \mathbf{y}), d_o(\mathcal{M}, \mathbf{y}) \right). \quad (3)$$

### 3.3. Robust Fitting

We seek to robustly fit a set of cuboids  $\mathcal{M} = \{\mathbf{h}_1, \dots, \mathbf{h}_{|\mathcal{M}|}\}$  to the possibly noisy 3D features  $\mathbf{y} \in \mathcal{Y}$ . To this end, we build upon the robust multi-model fitting approach of Kluger et al. [28]:

1. We predict sets of sampling weights  $\mathbf{p}$  from data  $\mathcal{Y}$  using a neural network  $f_w$ .
2. Using these sampling weights, a RANSAC-based [14] estimator generates a cuboid instance  $\mathbf{h}$ , which we append to  $\mathcal{M}$ .
3. Conditioned on  $\mathcal{M}$ , we update the sampling weights  $\mathbf{p}$  via  $f_w$  and generate the next cuboid instance.

We repeat these steps multiple times, until all cuboids have been recovered one-by-one. Fig. 2 gives an overview of the algorithm, while Fig. 5 depicts the sampling and fitting stage in more detail.

#### 3.3.1 Sampling

In [28], one set of sampling weights  $\mathbf{p}(\mathcal{Y} | \mathcal{M})$  is predicted at each step. Optimally, these weights should highlight a single coherent structure in  $\mathcal{Y}$  and suppress the rest, in order to maximise the likelihood of sampling an all-inlier set of features. However, we often have multiple important structures present in a scene. Which structure to emphasise and sample first may therefore be ambiguous, and the approach proposed in [28] struggles with primitive fitting for that reason. In order to deal with this, we allow the network to instead predict several sets of sampling weights  $\mathcal{Q} = \{\mathbf{p}_1, \dots, \mathbf{p}_Q\}$  and corresponding selection probabilities  $\mathbf{q}(\mathcal{Y} | \mathcal{M}) \in \mathbb{R}^Q$ , as shown in Fig. 4. This results in a two-step sampling procedure: First, we randomly select one of the sampling weight sets  $\mathbf{p} \in \mathcal{Q}$  according to  $\mathbf{q}$ . Then we sample a minimal set of features according to the selected sampling weights  $\mathbf{p}$  in order to generate a cuboid hypothesis. This allows the neural network  $(\mathcal{Q}, \mathbf{q}) = f_w(\mathcal{Y} | \mathcal{M})$

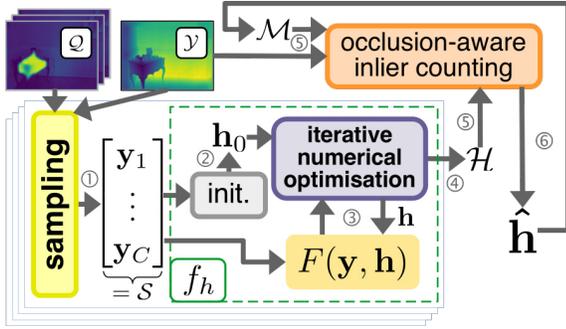


Figure 5: **Sampling and Fitting:** (1) We sample a minimal set of features  $\mathcal{S} \subset \mathcal{Y}$  using sampling weights  $\mathcal{Q}$  (Sec. 3.3.1). (2) Within minimal solver  $f_h$  (Sec. 3.3.2), we initialise cuboid parameters  $\mathbf{h}_0$ . (3) We optimise these parameters iteratively using Eq. 4, resulting in a hypothesis  $\mathbf{h}$ . (4) We compute multiple cuboid hypotheses concurrently, resulting in a set of hypotheses  $\mathcal{H}$ . (5) Using occlusion-aware inlier counting (Sec. 3.3.3), we select the best hypothesis  $\hat{\mathbf{h}}$  and (6) add it to the set of recovered cuboids  $\mathcal{M}$ .

to highlight multiple structures at once, without them interfering with each other. Ideally,  $\mathbf{q}$  contains non-zero values for all sampling weight sets highlighting valid structures, and zero values otherwise. At worst, it would degenerate to  $\mathbf{q} = \mathbf{e}_i$ , with  $i \in \{1, \dots, Q\}$ , effectively setting  $Q = 1$ . We avoid this via regularisation during training (cf. Sec. 3.4).

### 3.3.2 Fitting

Given a minimal set of sampled features  $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_C\} \subset \mathcal{Y}$ , we want to find parameters of a cuboid  $\mathbf{h} = f_h(\mathcal{S})$  such that it fits these features optimally. We define the minimal solver  $f_h$  via this objective function:

$$F(\mathbf{y}, \mathbf{h}) = d(\mathbf{h}, \mathbf{y})^2 \cdot (a_x + a_y + a_z). \quad (4)$$

As the size of the cuboid may be ambiguous when dealing with 2.5D data, we add the regularisation term  $(a_x + a_y + a_z)$  in order to favour smaller cuboids. Unfortunately, no closed form solution for  $\mathbf{h} \in \arg \min_{\mathbf{h}} \|F(\mathcal{S}, \mathbf{h})\|_1$  exists. Starting from an initial solution<sup>2</sup>  $\mathbf{h}_0$ , we therefore approximate it via an iterative numerical optimisation method such as L-BFGS [34] or Adam [25].

### 3.3.3 Occlusion-Aware Inlier Counting

In order to select the cuboid hypothesis  $\mathbf{h} \in \mathcal{H}$  which fits best to features  $\mathcal{Y}$ , given a set of existing cuboids  $\mathcal{M}$ , we need to define an inlier function  $f_I(\mathbf{y}, \mathbf{h})$ . We could naively take  $f_I(\mathbf{y}, \mathbf{h}) \in [0, 1]$ , with  $f_I(\mathbf{y}, \mathbf{h}) = 1$  if feature  $\mathbf{y}$  is well represented by cuboid  $\mathbf{h}$ , and  $f_I(\mathbf{y}, \mathbf{h}) = 0$  otherwise. However, as described in Sec. 3.2.2, we want to avoid cuboids which create occlusions. Hence, we define an occlusion-aware inlier function  $f_{\text{OAI}}(\mathbf{y}, \mathcal{M}) \in [-1, 1]$  with the additional property  $f_{\text{OAI}}(\mathbf{y}, \mathcal{M}) = -1$  if  $\mathbf{y}$  is occluded by

<sup>2</sup>Please refer to the supplementary (Sec. A.5) for details.

cuboids in  $\mathcal{M}$ , but *only* if it is occluded by cuboid sides to which it is not also an inlier. We define:

$$f_{\text{IO}}(\mathbf{y}, \mathbf{h}, i) = f_I(\mathbf{y}, \mathbf{h}, i) - \chi_o(\mathbf{y}, \mathbf{h}, i) \cdot (1 - f_I(\mathbf{y}, \mathbf{h}, i)),$$

to determine whether  $\mathbf{y}$  is an inlier to the  $i$ -th side of cuboid  $\mathbf{h}$  ( $f_{\text{IO}} > 0$ ), occluded by it ( $f_{\text{IO}} < 0$ ), or a regular outlier ( $f_{\text{IO}} = 0$ ). If  $\mathbf{y}$  is occluded by any cuboid side in  $\mathcal{M}$ , it must be marked as occluded. Otherwise it should be marked as an inlier to its closest cuboid, or as an outlier:

$$f_{\text{OAI}}(\mathbf{y}, \mathcal{M}) = \begin{cases} \min_{\mathbf{h}, i} f_{\text{IO}}(\mathbf{y}, \mathbf{h}, i) & \text{if } (\cdot) < 0, \\ \max_{\mathbf{h}, i} f_{\text{IO}}(\mathbf{y}, \mathbf{h}, i) & \text{else,} \end{cases} \quad (5)$$

with  $\mathbf{h} \in \mathcal{M}$  and  $i \in \{1, \dots, 6\}$ , *i.e.* minimising or maximising over all sides of all cuboids. This implies that features which are occluded indeed reduce the inlier count  $I_c$ , which we use to determine which cuboid hypothesis  $\mathbf{h}$  shall be added to our current set of cuboids  $\mathcal{M}$ :

$$I_c(\mathcal{Y}, \mathcal{M} \cup \{\mathbf{h}\}) = \sum_{\mathbf{y} \in \mathcal{Y}} f_{\text{OAI}}(\mathbf{y}, \mathcal{M} \cup \{\mathbf{h}\}). \quad (6)$$

## 3.4. Training

As in [28], we want to optimise parameters  $\mathbf{w}$  of the sampling weight estimation network in order to increase the likelihood of sampling all-inlier minimal sets of features. In order to achieve this, we minimise the expectation of a task loss  $\ell(\mathbf{h}, \mathcal{M})$  which measures how well the resulting cuboid  $\mathbf{h}$  and previously determined cuboids  $\mathcal{M}$  fit to a scene:

$$\mathcal{L}(\mathbf{w}) = \mathbb{E}_{\mathcal{H} \sim p(\mathcal{H}|\mathcal{M}; \mathbf{w})} [\ell(\hat{\mathbf{h}}, \mathcal{M})], \quad (7)$$

with  $\hat{\mathbf{h}} \in \mathcal{H}$  being the cuboid hypothesis selected according to the inlier criterion:

$$\hat{\mathbf{h}} \in \arg \max_{\mathbf{h} \in \mathcal{H}} I_c(\mathcal{Y}, \mathcal{M} \cup \{\mathbf{h}\}). \quad (8)$$

However, as described by Brachmann et al. [6, 7], this discrete hypothesis selection prohibits also learning parameters  $\mathbf{v}$  of the feature extraction network. We must therefore turn hypothesis selection into a probabilistic action:

$$\hat{\mathbf{h}} \sim p(\hat{\mathbf{h}}|\mathcal{H}, \mathcal{M}) = \frac{\exp I_c(\mathcal{Y}, \mathcal{M} \cup \{\hat{\mathbf{h}}\})}{\sum_{\mathbf{h} \in \mathcal{H}} \exp I_c(\mathcal{Y}, \mathcal{M} \cup \{\mathbf{h}\})}. \quad (9)$$

This allows us to compute the expected loss over cuboid hypotheses  $\mathcal{H}$ , which is differentiable:

$$\mathcal{L}(\mathbf{v}) = \mathbb{E}_{\mathbf{h} \sim p(\mathbf{h}|\mathcal{H}, \mathcal{M})} [\ell(\mathbf{h}, \mathcal{M})]. \quad (10)$$

Combining Eq. 7 and Eq. 10, we can train both networks together end-to-end:

$$\mathcal{L}(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{\mathcal{H} \sim p(\mathcal{H}|\mathcal{M}; \mathbf{w})} \mathbb{E}_{\mathbf{h} \sim p(\mathbf{h}|\mathcal{H}, \mathcal{M})} [\ell(\mathbf{h}, \mathcal{M})]. \quad (11)$$

Computing the exact expectation in Eq. 7 is intractable, so we approximate its gradient by drawing  $K$  samples of  $\mathcal{H}$ :

$$\frac{\partial \mathcal{L}(\mathbf{v}, \mathbf{w})}{\partial (\mathbf{v}, \mathbf{w})} \approx \frac{1}{K} \sum_{k=1}^K \left[ \mathbb{E}_{\mathbf{h}} [\ell] \frac{\partial \log p(\mathcal{H}_k; \mathbf{w})}{\partial (\mathbf{v}, \mathbf{w})} + \frac{\partial \mathbb{E}_{\mathbf{h}} [\ell]}{\partial (\mathbf{v}, \mathbf{w})} \right]. \quad (12)$$

**Differentiable Solver.** The above requires that the gradient for a particular 3D feature can be computed w.r.t. a resulting cuboid. However, cuboid parameters are computed via iterative numerical optimisation (cf. Sec. 3.3.2). Tracking the operations through this step results in inaccurate gradients and prohibitively high computational costs. In the following, we describe a solution which does not necessitate tracking the computations. Instead, it relies on the implicit function theorem to directly compute the desired gradients.

Given a set of features  $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_C\}$  and an optimally fit cuboid  $\mathbf{h}$ , we seek to compute the partial derivatives  $\partial \mathbf{h} / \partial \mathcal{S}$ . From Eq. 4, we have  $F(\mathcal{S}, \mathbf{h}) = \mathbf{0}$ . Via the implicit function theorem, we can therefore obtain:

$$\frac{\partial \mathbf{h}}{\partial \mathcal{S}} = - \left( \frac{\partial F}{\partial \mathbf{h}}(\mathcal{S}, \mathbf{h}) \right)^{-1} \cdot \frac{\partial F}{\partial \mathcal{S}}(\mathcal{S}, \mathbf{h}). \quad (13)$$

In practice, the partial derivatives by the size parameters  $\partial F / \partial a_{\{x,y,z\}}$  are mostly zero or close to zero. Inversion of the Jacobian  $\frac{\partial F}{\partial \mathbf{h}}(\mathcal{S}, \mathbf{h})$  is not possible or numerically unstable in this case. We thus approximate it by masking out the possibly zero derivatives and using the pseudo-inverse:

$$\frac{\partial \mathbf{h}}{\partial \mathcal{S}} \approx - \left( \frac{\partial F}{\partial (\mathbf{R}, \mathbf{t})}(\mathcal{S}, \mathbf{h}) \right)^+ \cdot \frac{\partial F}{\partial \mathcal{S}}(\mathcal{S}, \mathbf{h}). \quad (14)$$

This allows us to perform end-to-end backpropagation even without a differentiable minimal solver  $\mathbf{h} = f_h(\mathcal{S})$ .

**Task Loss.** During training, we aim to maximise the expected inlier counts of the sampled cuboids, *i.e.* we minimise the following loss:

$$\ell(\mathbf{h}, \mathcal{M}) = -I_c(\mathcal{Y}, \mathcal{M} \cup \{\mathbf{h}\}) \quad (15)$$

In order enable gradient computation according to Eq. 12,  $\ell$  must also be differentiable. Instead of a hard threshold, we thus use a soft inlier measure derived from [28]:

$$f_1(\mathbf{y}, \mathbf{h}, i) = 1 - \sigma(\beta(\frac{1}{\tau} d_i^p(\mathbf{y}, \mathbf{h})^2 - 1)), \quad (16)$$

with softness parameter  $\beta$ , inlier threshold  $\tau$ , and  $\sigma(\cdot)$  being the sigmoid function. This loss is based on geometrical consistency only, and does not need any additional labels.

**Regularisation.** In order to prevent mode collapse of the sampling weights  $\mathcal{Q}$  (cf. Sec. 3.3.1), we apply a regularisation term during training. We minimise the correlation coefficients between individual sets of sampling weights  $\mathbf{p} \in \mathcal{Q}$ :

$$\ell_{\text{corr}}(\mathcal{Q}) = \sum_{\substack{\mathbf{p}_i, \mathbf{p}_j \in \mathcal{Q} \\ i \neq j}} \frac{\text{cov}(\mathbf{p}_i, \mathbf{p}_j)}{\sigma_{\mathbf{p}_i} \sigma_{\mathbf{p}_j}}, \quad (17)$$

with covariances  $\text{cov}(\cdot, \cdot)$  and standard deviations  $\sigma$ . For the same reason, we also maximise the entropy of selection probabilities  $\mathbf{q}$ , *i.e.*  $\ell_{\text{entropy}} = -H(\mathbf{q})$ .

## 4. Experiments

**Dataset.** We provide qualitative and quantitative results on the NYU Depth v2 dataset [52]. It contains 1449 images of indoor scenes with corresponding ground truth depth recorded with a Kinect camera. We use the same dataset split as [30]: 654 images for testing and 795 images for training, of which we reserved 195 for validation.

**Implementation Details.** For experiments with RGB image input, we use the BTS [30] depth estimator pre-trained on NYU as our feature extraction network. BTS achieves state-of-the-art results on NYU with a publicly available implementation<sup>3</sup>, and is thus a sensible choice. We implemented the sample weight estimator as a variant of the fully convolutional neural network used in [7]. We pretrain the sample weight estimation network with ground truth depth as input for 20 epochs. This is also the network we use for experiments with depth input. We then continue training both networks end-to-end with RGB input for another 25 epochs. For cuboid fitting (cf. Sec. 3.3.2), we implement the minimal solver by applying the Adam [25] optimiser to perform gradient descent w.r.t. cuboid parameters for 50 iterations. We provide additional implementation details and a listing of hyperparameters in the supplementary.

**Baselines.** While multiple works in the field of primitive based 3D shape parsing have been published in recent years, not all of them can be used for comparison. We cannot use methods which require ground truth shape annotations [42, 60] or watertight meshes [16, 11] for training, as the data we are concerned with does not provide these. Unfortunately, we also cannot compare against [44], as they also evaluated on different data and have not yet provided their source code. For the cuboid based approach of [54], source code is available<sup>4</sup>. However, following the instructions they provided, we were unable to obtain sensible results on NYU. The neural network of [54] appears to degenerate when trained on NYU, predicting similar unreasonable cuboid configuration for all scenes (cf. qualitative results in the supplementary material). We therefore did not include it in our evaluation. On the other hand, we were able to train the superquadric-based approach of [45] (SQ-Parsing) on NYU using their provided source code<sup>5</sup>. Since SQ-Parsing is designed for meshes as input, we preprocess the ground truth depth point clouds of NYU by applying Poisson surface reconstruction [23]. We also compare against a variant of [45] which directly operates on RGB images (SQ-Parsing RGB). We further compare our method against a variant of Sequential RANSAC [55] from depth input. Lastly, we evaluate the results of SQ-Parsing when applied to the prediction of the same depth estimation

<sup>3</sup><https://github.com/cogaplex-bts/bts>

<sup>4</sup><https://github.com/shubhtuls/volumetricPrimitives>

<sup>5</sup>[https://github.com/paschalidoud/superquadric\\_parsing](https://github.com/paschalidoud/superquadric_parsing)

	occlusion-aware $L_2$ -distance					$L_2$
	AUC@50 cm	AUC@20 cm	AUC@10 cm	AUC@5 cm	mean (cm)	mean (cm)
RGB input						
<b>Ours</b>	<b>57.0%</b> $\pm 0.20$	<b>33.1%</b> $\pm 0.17$	<b>18.9%</b> $\pm 0.09$	<b>10.0%</b> $\pm 0.06$	<b>34.5</b> $\pm 0.24$	<b>30.1</b> $\pm 0.34$
SQ-Parsing [45] + BTS [30]	30.1%	11.6%	4.3%	1.1%	65.9	33.9
SQ-Parsing (RGB) [45]	19.4%	6.3%	1.9%	0.4%	87.4	42.4
depth input						
<b>Ours</b>	<b>77.2%</b> $\pm 0.08$	<b>62.7%</b> $\pm 0.07$	<b>49.1%</b> $\pm 0.12$	<b>34.3%</b> $\pm 0.14$	<b>20.8</b> $\pm 0.51$	<b>17.9</b> $\pm 0.50$
Sequential RANSAC [55]	59.9% $\pm 0.17$	42.8% $\pm 0.09$	30.1% $\pm 0.84$	18.8% $\pm 0.10$	37.9 $\pm 0.74$	32.7 $\pm 0.72$
SQ-Parsing [45]	50.3%	22.4%	8.8%	2.5%	34.7	20.9

Table 1: **Quantitative Results on NYU:** We evaluate on NYU Depth v2 [52], for both RGB and depth inputs. We compare our method against variants of SQ-Parsing [45] and Sequential RANSAC [55]. We present AUC values (higher is better) for various upper bounds of the occlusion-aware (OA)  $L_2$  distance (cf. Sec. 3.2.2). We also report mean OA- $L_2$  and regular  $L_2$  distances (lower is better). See Sec. 4 for details.

network (BTS [30]) which we employ for our method. We used the same parameters as [45] for these experiments.

**Metrics.** Previous works [45, 44] evaluated their results using Chamfer distance and volumetric IoU. These works, however, deal with full 3D shapes of watertight objects, while we fit 3D primitives to scenes where only a 2.5D ground truth is available. As we explain in Sec. 3.2.2, we therefore evaluate using the occlusion-aware distance metric, OA- $L_2$  for short, instead. We calculate the occlusion-aware distances of all ground-truth points to the recovered primitives per scene. Using these distances, we then compute the relative area under the recall curve (AUC, in percent) for multiple upper bounds: 50 cm, 20 cm, 10 cm, 5 cm. The AUC values are less influenced by outliers than the mean distance, and gauge how many points are covered by the primitives within the upper bound. In addition, we also report mean OA- $L_2$  as well as the mean of the regular  $L_2$  distance. As our method is random sampling based, we report the mean and variance over five runs for all metrics.

**Qualitative Results.** We present qualitative examples from the NYU dataset comparing the results of [45] against our method. In Fig. 6, rows 3-5, we show renderings of the cuboids predicted by by our method (RGB input in row 1) for scenes from the NYU test set. Row three is rendered from the same perspective as the original image, while row four and five show top and side views, respectively. As these examples show, our method is able to recover cuboids for key elements of the scene, such as walls, floors, cupboards, counters or tables. It often recovers cuboids covering larger parts of the image first. It does not always succeed in capturing volumetric properties of all objects, such as the fridge in column five, which is represented by two almost planar cuboids instead. In addition, spurious, usually very thin cuboids appear occasionally, such as the cuboid intersecting the table in the last column. For comparison, we show renderings of superquadrics obtained with SQ-Parsing + BTS – which is our best performing competi-

	OA- $L_2$			
	AUC@20	AUC@10	AUC@5	mean (cm)
$Q = 1$	40.2%	29.7%	19.8%	45.3
$Q = 2$	57.5%	43.6%	29.7%	22.3
$Q = 4$	<b>62.7%</b>	<b>49.1%</b>	<b>34.3%</b>	<b>20.8</b>

Table 2: **Ablation Study:** We evaluate our approach with varying numbers  $Q$  of sampling weight sets with depth input. We present AUC values for various upper bounds of the OA- $L_2$  distance and the mean OA- $L_2$  distance.

tor for RGB input – in row 2. Unlike our cuboid representations, these superquadrics bear little resemblance to the original scene. Our method is thus able to abstract the depicted indoor scenes in a much more sensible way than [45].

**Quantitative Evaluation.** Table 1 gives quantitative results on the NYU Depth v2 [52] test set for both RGB and depth inputs. For depth input, our method achieves significantly higher AUC values than Sequential RANSAC [55], with margins between 15.5 and 19.9 percentage points. Sequential RANSAC outperforms SQ-Parsing on AUC, yet is inferior w.r.t. the means. Our method also achieves significantly higher AUC values than SQ-Parsing [45] across the whole range, with margins between 26.9 and 40.3 percentage points. We present lower mean  $L_2$  and occlusion-aware  $L_2$  distances as well, with improvements of 3.0 cm (14.3%) and 13.9 cm (40.0%) respectively. For RGB input, Tab. 1 shows that SQ-Parsing + BTS, *i.e.* the combination of [45] for depth input with the monocular depth estimator from [30], performs better than the RGB variant of [45]. This implies that the image based encoder network of [45] is not able to extract 3D information from images as well as [30] can. Our approach, however, also outperforms SQ-Parsing + BTS by large margins, with AUC improvements of 8.9 to 26.9 percentage points. We improve mean  $L_2$  and OA- $L_2$  distances by 4.0 cm (11.8%) and 31.5 cm (47.8%) respectively. In summary, our method outperforms [45] on NYU in all settings on all metrics by significant margins, presenting a new state-of-the-art for this challenging task.



Figure 6: **Qualitative Results.** First row: Input RGB images. Second and third rows: Cuboids obtained with our proposed method (same views as the original images and top views, respectively). Colours convey the order in which the cuboids have been selected: red, blue, green, purple, cyan, orange. Cuboids covering large parts of the image are often selected early. Fourth row: OA- $L_2$  distance (Sec. 3.2.2) between the cuboids and the ground-truth depth. Very large distances are due to missing or occluding cuboids. Last row: Superquadrics obtained using [45]. These abstractions are very cluttered and hardly represent the original scenes. Our algorithm infers abstractions that represent the original scenes more closely.

**Ablation Study.** To demonstrate the efficacy of predicting multiple sets of sampling weights at once (cf. Sec. 3.3.1), we trained our approach with varying numbers  $Q$  of sampling weight sets. We evaluated these variants with depth input and present the results in Tab. 2. Predicting multiple sampling weight sets ( $Q = 4$ ) performs significantly better than predicting just one ( $Q = 1$ ) or two ( $Q = 2$ ) sets. Please refer to the supplementary for an ablation study w.r.t. our proposed occlusion-aware inlier counting.

## 5. Conclusion

We present a 3D scene parser which abstracts complex real-world scenes into ensembles of simpler volumetric primitives. It builds upon a learning-based robust estimator, which we extend in order to recover cuboids from RGB images. To this end, we propose an occlusion-aware distance metric which enables us to correctly handle opaque scenes. We facilitate end-to-end training by circumventing

backpropagation through our minimal solver and deriving the gradient of primitive parameters w.r.t. the input features analytically. Our algorithm neither requires known ground truth primitive parameters nor any other costly annotations. It can thus be straight-forwardly applied to other datasets which lack this information. Results on the challenging real-world NYU Depth v2 dataset demonstrate that the proposed method successfully parses and abstracts complex and cluttered 3D scenes. In future work, we plan to address common failure cases, *i.e.* planar and spurious cuboids, by replacing the depth estimation with more expressive 3D features and using matching data for training.

**Acknowledgements.** This work was supported by the BMBF grant *LeibnizAILab* (01DD20003), by the DFG grant *COVMAP* (RO 2497/12-2), by the DFG Cluster of Excellence *PhoenixD* (EXC 2122), and by the Center for Digital Innovations (ZDIN).

## References

- [1] Paul Amayo, Pedro Piniés, Lina M Paz, and Paul Newman. Geometric Multi-Model Fitting with a Convex Relaxation Algorithm. In *CVPR*, 2018. 2
- [2] Daniel Barath and Jiří Matas. Graph-Cut RANSAC. In *CVPR*, 2018. 1
- [3] Daniel Barath and Jiri Matas. Multi-Class Model Fitting by Energy Minimization and Mode-Seeking. In *ECCV*, 2018. 2
- [4] Daniel Barath and Jiri Matas. Progressive-X: Efficient, Any-time, Multi-Model Fitting Algorithm. *ICCV*, 2019. 2
- [5] Irving Biederman. Recognition-by-Components: a Theory of Human Image Understanding. *Psychological Review*. 1
- [6] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC - Differentiable RANSAC for Camera Localization. In *CVPR*, 2017. 2, 5
- [7] Eric Brachmann and Carsten Rother. Neural-Guided RANSAC: Learning Where to Sample Model Hypotheses. In *ICCV*, 2019. 2, 5, 6
- [8] 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 preprint arXiv:1512.03012, 2015. 3
- [9] Tat-Jun Chin, Jin Yu, and David Suter. Accelerated Hypothesis Generation for Multi-structure Data via Preference Analysis. *TPAMI*. 1
- [10] Antonio Criminisi, Ian Reid, and Andrew Zisserman. Single view metrology. *IJCV*, 2000. 2
- [11] Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, and Andrea Tagliasacchi. CvxNet: Learnable Convex Decomposition. In *CVPR*, 2020. 6
- [12] Debidatta Dwibedi, Tomasz Malisiewicz, Vijay Badrinarayanan, and Andrew Rabinovich. Deep Cuboid Detection: Beyond 2D Bounding Boxes. *arXiv preprint arXiv:1611.10010*, 2016. 2
- [13] David Eigen, Christian Puhrsch, and Rob Fergus. Depth Map Prediction from a Single Image using a Multi-Scale Deep Network. In *NeurIPS*, 2014. 1, 2
- [14] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *ACM*, 1981. 1, 2, 4
- [15] R. Garg, V. K. Bg, G. Carneiro, and I. Reid. Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. In *ECCV*, 2016. 2
- [16] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning Shape Templates with Structured Implicit Functions. In *ICCV*, 2019. 6
- [17] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV*, 2019. 2
- [18] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation with Left-Right Consistency. In *CVPR*, 2017. 2
- [19] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging Into Self-Supervised Monocular Depth Estimation. In *ICCV*, 2019. 2
- [20] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A Papier-Mâché Approach to Learning 3D Surface Generation. In *CVPR*, 2018. 2
- [21] Abhinav Gupta, Alexei A. Efros, and Martial Hebert. Blocks World Revisited: Image Understanding Using Qualitative Geometry and Mechanics. In *ECCV*, 2010. 2
- [22] Hao Jiang and Jianxiong Xiao. A Linear Approach to Matching Cuboids in RGBD Images. In *CVPR*, 2013. 2
- [23] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson Surface Reconstruction. In *Eurographics Symposium on Geometry Processing*, 2006. 6
- [24] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab. SSD-6D: Making RGB-based 3D detection and 6D pose estimation great again. In *ICCV*, 2017. 2
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015. 5, 6
- [26] Florian Kluger, Hanno Ackermann, Michael Ying Yang, and Bodo Rosenhahn. Deep learning for vanishing point detection using an inverse gnomonic projection. In *GCPR*, 2017. 2
- [27] Florian Kluger, Hanno Ackermann, Michael Ying Yang, and Bodo Rosenhahn. Temporally consistent horizon lines. In *ICRA*, 2020. 2
- [28] Florian Kluger, Eric Brachmann, Hanno Ackermann, Carsten Rother, Michael Ying Yang, and Bodo Rosenhahn. CONSAC: Robust Multi-Model Fitting by Conditional Sample Consensus. In *CVPR*, 2020. 2, 3, 4, 5, 6
- [29] Yevhen Kuznetsov, Jörg Stückler, and Bastian Leibe. Semi-Supervised Deep Learning for Monocular Depth Map Prediction. In *CVPR*, 2017. 2
- [30] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From Big to Small: Multi-Scale Local Planar Guidance for Monocular Depth Estimation. *arXiv preprint arXiv:1907.10326*, 2019. 2, 6, 7
- [31] Dahua Lin, Sanja Fidler, and Raquel Urtasun. Holistic Scene Understanding for 3D Object Detection with RGBD Cameras. In *ICCV*, 2013. 2
- [32] Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, and Jan Kautz. PlaneRCNN: 3D Plane Detection and Reconstruction from a Single Image. In *CVPR*, 2019. 2
- [33] Chen Liu, Jimei Yang, Duygu Ceylan, Ersin Yumer, and Yasutaka Furukawa. PlaneNet: Piece-wise Planar Reconstruction from a Single RGB Image. In *CVPR*, 2018. 2
- [34] Dong Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 1989. 5
- [35] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning Depth from Single Monocular Images using Deep Convolutional Neural Fields. *TPAMI*, 2015. 2
- [36] Fangchang Ma, Guilherme Venturéli Cavalheiro, and Sertac Karaman. Self-supervised Sparse-to-Dense: Self-supervised Depth Completion from Lidar and Monocular Camera. In *ICRA*, 2019. 2

- [37] Luca Magri and Andrea Fusiello. Fitting Multiple Heterogeneous Models by Multi-Class Cascaded T-Linkage. In *CVPR*, 2019. 2
- [38] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised Learning of Depth and Ego-Motion from Monocular Video Using 3D Geometric Constraints. In *CVPR*, 2018. 2
- [39] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D Reconstruction in Function Space. In *CVPR*, 2019. 2
- [40] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka. 3D Bounding Box Estimation Using Deep Learning and Geometry. In *CVPR*, 2017. 2
- [41] Y. Nie, X. Han, S. Guo, Y. Zheng, J. Chang, and J. Zhang. Total3DUnderstanding: Joint Layout, Object Pose and Mesh Reconstruction for Indoor Scenes from a Single Image. In *CVPR*, 2020. 2
- [42] Chengjie Niu, Jun Li, and Kai Xu. Im2Struct: Recovering 3D Shape Structure from a Single RGB image. In *CVPR*, 2018. 6
- [43] Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. Deep mesh reconstruction from single RGB images via topology modification networks. In *ICCV*, 2019. 2
- [44] Despoina Paschalidou, Luc Van Gool, and Andreas Geiger. Learning Unsupervised Hierarchical Part Decomposition of 3D Objects from a Single RGB Image. In *CVPR*, 2020. 1, 3, 6, 7
- [45] Despoina Paschalidou, Ali Osman Ulusoy, and Andreas Geiger. Superquadrics Revisited: Learning 3D Shape Parsing Beyond Cuboids. In *CVPR*, 2019. 1, 3, 6, 7, 8
- [46] Trung Thanh Pham, Tat-Jun Chin, Konrad Schindler, and David Suter. Interacting Geometric Priors for Robust Multimodel Fitting. *Transactions on Image Processing (TIP)*, 2014. 2
- [47] Pulak Purkait, Tat-Jun Chin, Hanno Ackermann, and David Suter. Clustering with Hypergraphs: The Case for Large Hyperedges. In *ECCV*, 2014. 1, 2
- [48] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas. Frustum PointNets for 3D Object Detection from RGB-D Data. In *CVPR*, 2018. 2
- [49] Zhile Ren and Erik B Sudderth. Three-Dimensional Object Detection and Layout Prediction using Clouds of Oriented Gradients. In *CVPR*, 2016. 2
- [50] Lawrence G Roberts. *Machine Perception of Three-Dimensional Solids*. PhD thesis, Massachusetts Institute of Technology, 1963. 1, 2
- [51] S. Shi, X. Wang, and H. Li. PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud. In *CVPR*, 2019. 2
- [52] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor Segmentation and Support Inference from RGBD Images. In *ECCV*, 2012. 2, 6, 7
- [53] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, Francois Goulette, and Leonidas J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *ICCV*, 2019. 2
- [54] Shubham Tulsiani, Hao Su, Leonidas J Guibas, Alexei A Efros, and Jitendra Malik. Learning Shape Abstractions by Assembling Volumetric Primitives. In *CVPR*, 2017. 1, 3, 6
- [55] Etienne Vincent and Robert Laganière. Detecting planar homographies in an image pair. In *ISPA*, 2001. 6, 7
- [56] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. In *ECCV*, 2018. 1, 2
- [57] Jianxiong Xiao, Bryan Russell, and Antonio Torralba. Localizing 3D cuboids in single-view images. In *NeurIPS*, 2012. 2
- [58] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose. In *CVPR*, 2018. 2
- [59] Yinda Zhang, Mingru Bai, Pushmeet Kohli, Shahram Izadi, and Jianxiong Xiao. DeepContext: Context-Encoding Neural Pathways for 3D Holistic Scene Understanding. In *ICCV*, 2017. 2
- [60] Chuhan Zou, Ersin Yumer, Jimei Yang, Duygu Ceylan, and Derek Hoiem. 3D-PRNN: Generating Shape Primitives with Recurrent Neural Networks. In *ICCV*, 2017. 6