

Probabilistic Vehicle Reconstruction Using a Multi-Task CNN

Max Coenen Franz Rottensteiner
 Leibniz University Hannover, Germany

{coenen, rottensteiner}@ipi.uni-hannover.de

Abstract

The retrieval of the 3D pose and shape of objects from images is an ill-posed problem. A common way to object reconstruction is to match entities such as keypoints, edges, or contours of a deformable 3D model, used as shape prior, to their corresponding entities inferred from the image. However, such approaches are highly sensitive to model initialisation, imprecise keypoint localisations and/or illumination conditions. In this paper, we present a probabilistic approach for shape-aware 3D vehicle reconstruction from stereo images that leverages the outputs of a novel multi-task CNN. Specifically, we train a CNN that outputs probability distributions for the vehicle’s orientation and for both, vehicle keypoints and wireframe edges. Together with 3D stereo information we integrate the predicted distributions into a common probabilistic framework. We believe that the CNN-based detection of wireframe edges reduces the sensitivity to illumination conditions and object contrast and that using the raw probability maps instead of inferring keypoint positions reduces the sensitivity to keypoint localisation errors. We show that our method achieves state-of-the-art results, evaluating our method on the challenging KITTI benchmark and on our own new ‘Stereo-Vehicle’ dataset.

1. Introduction

The highly dynamic nature of street environments is one of the biggest challenges for autonomous driving applications. The precise reconstruction of moving objects, especially of other cars, are fundamental to ensure safe navigation and to enable applications such as interactive motion planning and collaborative positioning. To this end, cameras provide a cost-effective solution to deliver perceptive data of a vehicle’s surroundings. However, the projection from 3D to 2D images leaves many ambiguities about 3D objects, causing the retrieval of their pose and shape to be ill-posed and difficult to solve. To confine the parameter space, deformable models can be used as shape prior and are aligned with the objects in the image to recover their pose and shape. In this paper, we make use of such a deformable



Figure 1: Qualitative results of our method. Left: Heatmaps for vehicle keypoints (top) and vehicle wireframes (bottom) superimposed to the input image. Right: Input image and backprojected 3D wireframe of a reconstructed vehicle (top) and 3D view on the reconstructed scene (bottom).

vehicle model and present a method that fully reconstructs vehicles in 3D given street level stereo image pairs, allowing the derivation of precise 3D pose and shape parameters. Earlier approaches used model edges and contours to align them with image edges, usually derived from gradient information [21, 31, 33]. However, these approaches are highly sensitive to illumination, reflections, contrast, object color and model initialisation, because these factors can cause erroneous edge-to-edge correspondences, thus prohibiting a correct model alignment. Instead, recent approaches leveraged keypoint detections to be used for model alignment [1, 9, 27, 30]. However, compared to edges, keypoints are less stable in the way that already small localisation errors are likely to cause large errors in 3D space. We start from the idea that if it is possible to detect keypoints, it should also be possible to detect model edges and contours without the explicit dependency on good image gradients. Based on initially detected vehicles, we make the following contributions in this paper: **(1)** We propose a novel multi-task convolutional neural network (CNN) that simultaneously detects vehicle keypoints and vehicle wireframe edges (cf. Fig. 1) and also outputs a probability distribution for the vehicle’s orientation. For the orientation estimation, we expand the overlapping viewpoint class definition proposed

in [25] by defining a novel hierarchical class and classifier structure. We define a novel loss for the detection of keypoints and wireframes, which we believe allows the usage of a single encoder-decoder (“hourglass”) CNN and makes the repeated hourglass architecture proposed in [28, 26] unnecessary. **(2)** Instead of relying on gradient based edge representations [7, 21, 33], which highly depend on illumination, contrast, and object color, we incorporate the wireframe predictions given by our CNN into our reconstruction approach. Since we train our CNN to distinguish between wireframe edges belonging to different sides of the vehicle, we not only avoid the problems due to low contrast of vehicle edges and silhouettes, but also achieve better prospects for edge-to-edge correspondences. **(3)** For the purpose of vehicle reconstruction, we build upon our previous probabilistic model and optimisation procedure [7] and significantly extend the model by adding state prior terms, leading to major improvements of the reconstruction results. To avoid the error source of incorrect keypoint localisations, in contrast to [1, 30, 26], we build our probabilistic model directly on the raw keypoint heatmaps obtained by our CNN. **(4)** We propose a new dataset for vehicle reconstruction, which exceeds the famous KITTI dataset [11] as it not only delivers 2D and 3D reference bounding boxes, but precisely fitted vehicle CAD models, which allow the evaluation of shape reconstructions and vehicle categorisation or identification.

2. Related Work

One of the biggest challenges in image based reconstruction and pose estimation of vehicles is the enormous variability of appearance, caused by the intra-class variability of vehicles and by different viewpoints of the acquired images. One strategy to overcome these problems is to train viewpoint and/or category specific classifiers [15, 19, 29] or to learn viewpoint specific shape templates [31] to reason about vehicle pose and/or shape. However, in these approaches the viewpoint usually is discretised into a set of viewpoint bins and thus only coarse viewpoint estimates are delivered as output. 3D vehicle detection approaches such as [4] and [25] deliver oriented bounding box estimates in 3D using CNNs. However, describing objects by a box only gives a very coarse representation of their shape. Due to the ambiguous representation of 3D object information in images, the usage of 3D models as shape prior can be extremely beneficial for the task of object reconstruction. In this context, CAD vehicle models can be used directly to guide the vehicle reconstruction [13]. However, finding the best suitable among the vast amount of existing CAD models becomes intractable quickly. Instead, deformable shape representations learned from a set of reference shapes, such as signed distance functions (SDF) [10] or active shape models (ASM) [42] are more flexible rep-

resentations to cope with the intra-class variability of vehicles. Using such prior models, a shape aware reconstruction is conducted, in which an instance of the deformed and transformed shape model is fitted to the observations to derive the target pose and shape parameters. Matching a shape prior with image observations can be done via curve alignment of backprojected model edges and image edges [6, 21, 31, 33] or via alignment of the backprojected model silhouette with an instance segmentation mask in the image [8, 17, 32, 39]. However, matching image edges highly depends on the illumination and contrast of the objects and on a good model initialisation to be able to establish correct edge-to-edge correspondences. On the other hand, silhouette matching inherits pose and shape ambiguities and neglects object details and structures. To counteract these problems, we define a vehicle wireframe that contains edges not only representing the silhouette but also boundaries between vehicle parts and learn a CNN to detect the wireframe, differentiating between edges belonging to different sides of the vehicle (left, right, front, back). This enables an illumination/contrast invariant, finer-grained and more robust establishment of correspondences for model fitting.

[6, 10, 40] propose methods for 3D vehicle reconstruction solely based on 3D points obtained from stereo images or laserscanning, respectively. In contrast to images, 3D points provide explicit 3D information, but when used as the only data source they deliver a rather sparse and incomplete representation of the object and valuable image cues are disregarded completely for model fitting.

Another strategy to shape-aware vehicle reconstruction is to match semantic model keypoints with their corresponding keypoints detected in the image. Traditionally, handcrafted features are used in [2, 7, 22, 23, 42] for the model to keypoint alignment, while recently CNNs are applied to detect keypoints [1, 3, 9, 27, 30]. Typically, to infer the object pose and shape, the backprojection error of model keypoints and detected keypoints is minimized which makes these approaches prone to imprecise, incorrect and missing keypoint localisations. In contrast, we do not infer pointwise keypoint locations but instead build our model fitting approach on probabilistic keypoint heatmaps generated by our CNN.

End-to-end approaches that take the image as input and infer pose and shape of the target vehicle directly are presented in [20, 24]. These methods require a vast amount of expensive training data, which is why often synthetically generated data is used for learning [20, 24]. However, training CNNs on synthetic images usually leads to a drop of performance when applying the CNN to real world images. Besides, explicit scene and/or model knowledge is disregarded in [20, 24]. In contrast, we derive scene knowledge from the data and incorporate it as prior information.

3. Shape aware vehicle reconstruction

3.1. Overview

The goal of our method is to recover the precise pose (i.e. position and orientation) as well as the shape of vehicles detected from street-level stereo images. For this purpose, we fit a 3D vehicle model to the detected vehicles. We learn a deformable model as shape prior and formulate a probabilistic model to find the best fitting vehicle model by making use of a new vehicle CNN, trained to predict vehicle keypoints and wireframe edges, as well as the vehicle’s viewpoint. A schematic overview of our method is shown in Fig. 2. The input to our method are stereo image pairs, incl. their interior and relative orientation parameters. To derive 3D information, we make use of the ELAS matcher [12] to calculate a dense disparity map for every stereo pair and reconstruct 3D points \mathbf{X} via triangulation for every pixel of the reference image (the left stereo partner).

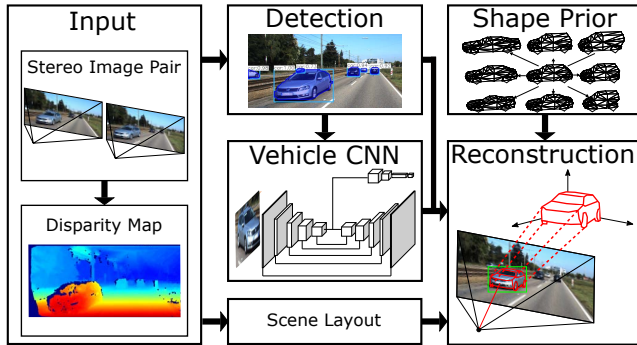


Figure 2: Overview of our framework.

3.1.1 Notation

Given a set of detected vehicles $\mathbf{v}_k \in \mathcal{V}$, our goal is to associate each vehicle with its state vector $\mathbf{s}_k = (\mathbf{t}_k, \theta_k, \gamma_k)$, comprising its pose and shape parameters. After determining the ground plane $\Omega \in \mathbb{R}^3$, we describe the vehicle pose by its 2D position \mathbf{t}_k on the ground plane and its heading θ_k , i.e. the rotation angle about the normal vector of the ground plane; γ_k is a vector of shape parameters determining the shape of the 3D deformable ASM representing each vehicle (cf. Sec. 3.1.4).

3.1.2 Scene layout

We use the stereo data to derive knowledge about the 3D layout of the scene, represented by the 3D ground plane and a probabilistic free-space grid map.

Ground plane Ω : We apply RANSAC to the stereo point cloud \mathbf{X} to find the ground plane Ω as plane of maximum support. All inliers of the final RANSAC consensus

set are stored as ground points $\mathbf{X}_\Omega \subset \mathbf{X}$. Requiring vehicles always to be located on the ground plane, we are able to fix three of the 6 DoF vehicle pose parameters (1 translational + 2 rotational parameters) and, thus, to constrain the parameter space of the model fitting approach.

Probabilistic free-space grid map Φ : Based on the ground plane points \mathbf{X}_Ω and all the points not belonging to the ground plane, thus representing arbitrary objects $\mathbf{X}_{Obj} = \mathbf{X} \setminus \mathbf{X}_\Omega$, it is possible to reason about free space, i.e. non occupied areas, in the observed scene. To represent the free space areas we create a probabilistic free space grid map Φ . For this purpose, we create a grid in the ground plane consisting of square cells with a side length l_Φ . For each grid cell Φ_g with $g \in [1, G]$ we count the number of ground points n_Ω^g and the number of object points n_{Obj}^g whose vertical projection is within the respective cell. We define the probability ρ_g of each cell to be free space as the ratio of n_Ω^g and the sum of n_Ω^g and n_{Obj}^g . Grid cells without projected points are marked as *unknown*.

3.1.3 Vehicle detection

To initially detect vehicles we apply the pretrained mask R-CNN (mRCNN) [14] to the reference image. Besides its good performance it has the advantage of not only delivering bounding boxes but also an instance segmentation mask for every vehicle. To obtain a list of $k \in [1, K]$ detected vehicles $\mathbf{v}_k = (\mathbf{X}_k, {}^l\mathcal{B}_k, {}^r\mathcal{B}_k)$, we associate each detection with the object points \mathbf{X}_k reconstructed from the pixels belonging to the respective segmentation mask, as well as with its left and right image bounding boxes ${}^l\mathcal{B}_k$ and ${}^r\mathcal{B}_k$, the latter being derived from the dense stereo correspondences.

3.1.4 Shape prior

Similar to [42] we use a 3D ASM as vehicle shape prior. The ASM is learned by applying principal component analysis (PCA) to a set of manually annotated keypoints \mathcal{K} of 3D CAD vehicle models. A deformed vehicle ASM is defined by the deformed vertex positions $\nu(\gamma)$, which can be obtained by the linear combination

$$\nu(\gamma) = \mathbf{m} + \sum_j \gamma^{(j)} \sigma_j \mathbf{e}_j \quad (1)$$

of the mean model \mathbf{m} and the eigenvectors \mathbf{e}_j , weighted by the square root of their corresponding eigenvalues σ_j^2 and scaled by the object specific shape parameters $\gamma^{(j)}$. A fully parametrised instance of a 3D vehicle ASM on the ground plane, denoted by $M(\mathbf{s})$, can be created according to the state vector \mathbf{s} by computing the deformed keypoints using the shape vector γ and subsequently shifting and rotating the whole model on the ground plane according to the translation vector \mathbf{t} and the heading angle θ .

Geometrical representation: We represent the model surface by defining a triangular mesh M_Δ for the ASM shape vertices \mathcal{K} . Further, we use a subset of \mathcal{K} to define a wireframe M_W of the vehicle model, consisting of two types of edges: *crease* edges that describe the outline of the vehicle and *semantic* edges, describing the boundaries between semantically different vehicle parts. Another subset $\mathcal{K}_A \in \mathcal{K}$ are chosen as keypoints for which we learn an image based detector described later (cf. Sec. 3.2). Fig. 3 shows the triangulations of several deformed models, their wireframes and the keypoints \mathcal{K}_A .

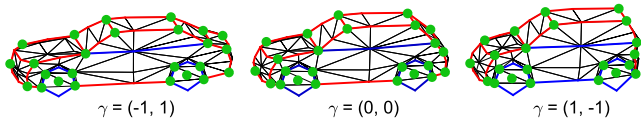


Figure 3: ASM: Triangulated surface (black), wireframe (red/blue: crease/semantic edges) and keypoints (green).

3.2. CNN structure

Our multi-task CNN consists of one common input branch and two individual output branches, each of them corresponding to one task, respectively. The overall architecture of the network can be seen in Fig. 5. The input to the network are images showing a vehicle, cropped by the bounding box. The network output consists of a probability distribution for the vehicle’s viewpoint, a probability heatmap for each of a set of vehicle keypoints, and probability heatmaps for the vehicle wireframe edges.

Input branch: This branch contains a series of shared convolutions and max pooling layers, thus acting as a shared backbone feature extractor, adopting the architecture of the VGG19 network [37].

Viewpoint branch: We design the viewpoint branch to output a probability distribution Π_ϑ for the vehicle viewpoint ϑ , which describes the aspect under which the vehicle is seen. As depicted in Fig. 4a, the viewpoint is defined as the angle between the image ray to the center of the vehicle and the vehicle’s longitudinal axis (red arrow in Fig 4a). Given the direction of the image ray ρ the vehicle orientation can directly be computed from the viewpoint via $\vartheta = 180^\circ - \theta - \rho$. In order to derive a probability distribution for the viewpoint of the vehicle, which can later be used in our probabilistic formulation for model fitting we set up a classification network rather than a regression network. The classes correspond to discretised orientation bins for the viewpoint estimation. The design of our viewpoint branch follows two assumptions regarding the general behavior of a viewpoint bin classifier: First, we expect the classification accuracy to decrease with an increasing number and therefore a finer definition of viewpoint bins. Second, we expect classification errors to primarily occur for

vehicles with viewpoint angles close to the bin borders, i.e. being distributed close to the diagonal of the confusion matrix. Inspired by [41], to consider the first assumption, we combine three hierarchical layers of classes for classification, containing different numbers (4, 8 and 16) of viewpoint classes. Thus, we believe to profit from the higher classification accuracy for the coarse class definition and the finer level of detail of the fine class definition. To consider the second assumption, we let the viewpoint bins of the individual layers overlap so that finer viewpoint classes do not share borders with coarser classes. Our hierarchical division of viewpoint classes can be seen in Fig. 4b.

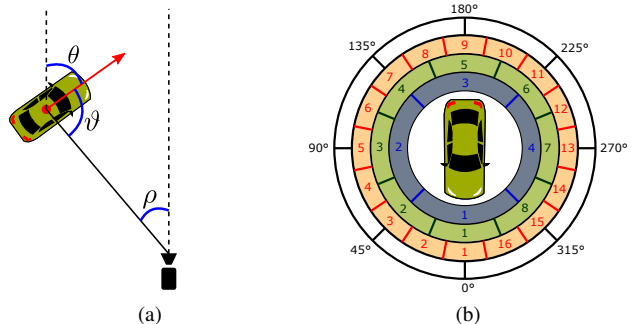


Figure 4: Definition of the viewpoint angle (a) and of the hierarchical viewpoint classes in blue, green, and red (b).

However, in contrast to [41], we do not apply conditional hierarchical classification, in which the output of the coarse classification layers decides about the execution of individual classifiers for the finer layers, because in that case, the output of images that are routed to an incorrect fine classifier cannot be corrected anymore. Instead, we split the viewpoint branch into three independent softmax classification heads, one for each layer, and establish a skip-connection between features extracted at coarser layers to the feature extraction pipeline of the next finer layers (cf. Fig. 5). In this way, the fine category classifiers can profit from the information extracted by the coarse classifiers but are less dependent on the input. The output of the classification heads is fed into a probabilistic averaging layer. Thus, our hierarchical class structure serves two purposes. On the one hand, we profit from the more reliable output of the coarse layers while still leveraging the more detailed output of the finer layers. On the other hand, due to the overlapping viewpoint bins, we mitigate the effect of misclassifications occurring between neighboring viewpoint classes.

Keypoint/Wireframe branch: This branch corresponds to a decoder network, upsampling the output of the input branch to the original input resolution. Inspired by [28] and [26], it is trained to produce one heatmap \mathcal{H}_c^c for every keypoint $c \in [1, C]$ in \mathcal{K}_A . Unlike previous work, we

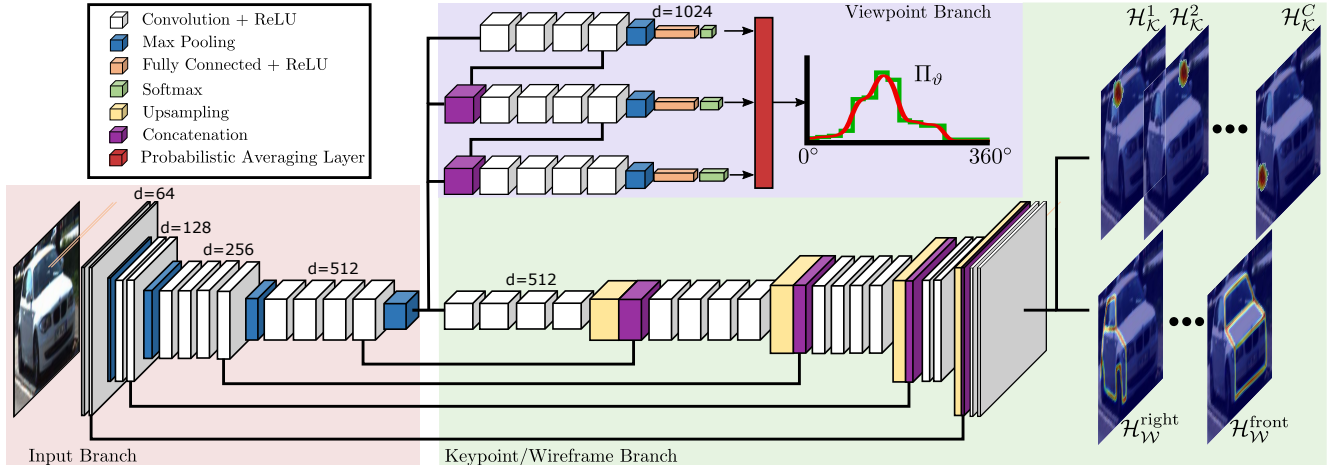


Figure 5: Architecture of our multi-task CNN. The input is a 3 channel image of size 224x224. The convolutional filters have size 3x3, max pooling and upsampling use filter size 2x2 and stride 2. The number of filters is denoted by d in the figure.

adapt the network to also output heatmaps $\mathcal{H}_{\mathcal{W}}^w$ for the vehicle wireframe edges. To this end, we subdivide the wireframe edges into four mutually non-exclusive subsets, each of which contains all edges belonging to the wireframe of one of the four vehicle sides $w \in \{\text{front, back, left, right}\}$. The values at each pixel position of the resulting heatmaps correspond to a probability for the presence of the respective keypoint/wireframe edge at that position. Together with the input branch, the keypoint/wireframe network follows a symmetrical UNet-like [34] architecture including skip connections between corresponding layers of the encoder and decoder blocks. The head of this branch consists of $C + 4$ binary classifiers using a sigmoid activation function to produce the C keypoint and the four wireframe heatmaps.

Training: The input branch is initialized from the corresponding layers of the VGG19 network [37], pre-trained on ImageNet [35] and frozen during training. The remaining convolutional layers are initialised using a normal distribution. We train the viewpoint branch using categorical cross-entropy loss at each classification head given the groundtruth bin containing the groundtruth viewpoint of the training vehicle images. We create training images for the keypoint/wireframe branch using annotated image keypoints. By automatically connecting keypoints corresponding to our wireframe definition we derive training data for the wireframe outputs. The training images are generated by placing 2D Gaussians at the corresponding keypoint positions, gaussian blurred edges along the reference wireframe, and zero everywhere else. The standard deviation of the Gaussians varies according to the distance of the vehicle to the camera and thus adapts to the size of the bounding box. Training of the keypoint/wireframe branch is done by comparing the predicted heatmaps to the groundtruth heatmaps. In [28] and [26], a mean squared error (MSE)

loss is used for training, in which each pixel contributes equally to the loss. In our use case of detecting keypoints and edges, this leads to an unfavorably broad shape of the loss function due to the extremely small proportion of keypoint/wireframe pixels w.r.t. non-keypoint/non-wireframe pixels in the groundtruth. To overcome this problem, we apply a new custom keypoint/wireframe loss $\mathcal{L}_{\mathcal{K}\mathcal{W}}$ with

$$\mathcal{L}_{\mathcal{K}\mathcal{W}} = \text{MSE}_{\text{true}} + \text{MSE}_{\text{false}} + \text{MSE}_{\text{pred}}. \quad (2)$$

Here, MSE_{true} , $\text{MSE}_{\text{false}}$ and MSE_{pred} correspond to individual MSE, each computed for a different subset of pixels. In MSE_{true} only pixels with groundtruth values larger than zero, whereas in $\text{MSE}_{\text{false}}$ only pixels with groundtruth values equal to zero are considered. In this way, the keypoint/wireframe and non-keypoint/non-wireframe pixels contribute to the loss with equal weight. Furthermore, we add an additional regularizing term MSE_{pred} to the loss function. In this term, only pixels whose prediction exceeds a predefined threshold are considered. Thus, this term puts emphasis on keypoint/wireframe detections and acts as additional penalty of false positive outputs. The network is trained using Keras [5], Adam optimizer for optimisation [18], a batch-size of 50, and a learning rate of 10^{-4} . To improve training, we drop the learning rate by a factor of 10 after 5 validation accuracy plateaus, use Batch normalisation [16], and apply Dropout [38] to the fully-connected layers with a rate of 0.5.

3.3. Probabilistic Model

Given the vehicle detections \mathbf{v}_k , we fit a vehicle model $M(\mathbf{s}_k)$ to each detection by finding the optimal state variables $\hat{\mathbf{s}}_k = (\hat{\mathbf{t}}_k, \hat{\theta}_k, \hat{\gamma}_k)$. Neglecting the index k (where it is possible) to simplify our notation in the following sections,

$\hat{\mathbf{s}}$ can be derived by maximising the posterior

$$p(\mathbf{s}|\mathbf{v}) = \frac{p(\mathbf{v}|\mathbf{s}) \cdot p(\mathbf{s})}{p(\mathbf{v})} \rightarrow \max. \quad (3)$$

Adapting the ideas from [7], we further factorize the likelihood and the prior according to

$$p(\mathbf{s}|\mathbf{v}) \propto \underbrace{p(X|\mathbf{s}) \cdot p(\mathcal{H}_{\mathcal{K}}|\mathbf{s}) \cdot p(\mathcal{H}_{\mathcal{W}}|\mathbf{s})}_{\text{Observation likelihood}} \cdot \underbrace{p(\mathbf{t}) \cdot p(\theta)}_{\text{State prior}}. \quad (4)$$

We minimise the negative logarithm of the posterior of Eq. 4. The individual likelihood and prior terms are explained in the following paragraphs.

3D likelihood: Based on the distances of the 3D points X_k to the surface M_{Δ} of the model $M(\mathbf{s})$ we calculate the 3D likelihood as

$$\log p(X_k|\mathbf{s}) = -\frac{1}{P} \sum_{x \in X_k} \frac{d_{\sigma_x}(x, M_{\Delta})}{2\sigma_x^2}. \quad (5)$$

Here, P is the overall number of 3D points in X_k and σ_x is the depth uncertainty of the individual 3D point x . We apply the Huber norm to calculate $d_{\sigma_x}(x, M_{\Delta})$, as it is more robust against outliers. This likelihood fits the 3D ASM to the 3D point cloud.

Keypoint likelihood: To calculate this term we backproject the visible model keypoints of \mathcal{K}_A to the stereo images, resulting in $c = [1, C_v]$ image points $\mathbf{u}_c^{l/r}$ for both, the left (l) and right (r) stereo images, respectively. The keypoint likelihood is calculated using the keypoint heatmaps $\mathcal{H}_{\mathcal{K}}$ by

$$\log p(\mathcal{H}_{\mathcal{K}}|\mathbf{s}) = -\frac{1}{2C_v} \sum_{i \in \{l, r\}} \sum_{c=1}^{C_v} \log(1 - {}^i\mathcal{H}_{\mathcal{K}}^c(\mathbf{u}_c^i)) \quad (6)$$

Here, $\mathcal{H}_{\mathcal{K}}^c(\mathbf{u}_c)$ denotes the output of the heatmap for the keypoint c at the location \mathbf{u}_c . This term fits the 3D ASM to the predicted keypoints.

Wireframe likelihood: This term is based on a measure of similarity between the backprojected edges of the model wireframe $M_{\mathcal{W}}(\mathbf{s})$ and the wireframe heatmaps $\mathcal{H}_{\mathcal{W}}$ resulting from our CNN. To this end, considering self-occlusion, we backproject the visible parts of the wireframe subsets $w \in \{\text{front, back, left, right}\}$ to the left and right images, resulting in binary wireframe images ${}^l I_{\mathcal{W}}^w$ and ${}^r I_{\mathcal{W}}^w$ with entries of 1 at pixels that are crossed by a wireframe edge and 0 everywhere else. We blur the wireframe images using a Gaussian filter to account for generalisation effects. The size of the filter is defined according to the backprojection uncertainty of the model keypoints given the generalisation

error of the ASM which is quantified to be 10 cm. The wireframe likelihood is calculated according to

$$\log p(\mathcal{H}_{\mathcal{W}}|\mathbf{s}) = -\frac{1}{2} \sum_{i \in \{l, r\}} \sum_w \log(1 - BC({}^i I_{\mathcal{W}}^w, {}^i \mathcal{H}_{\mathcal{W}}^w)) \quad (7)$$

where we use the Bhattacharyya coefficient $BC(\cdot, \cdot)$ as similarity measure between the wireframe images and the wireframe heatmaps. This likelihood will be large if the backprojected wireframes correspond well to the wireframes predicted by the CNN.

Position prior: The position prior is derived from the probabilistic free-space grid map Φ (cf. Sec 3.1.2). It is calculated based on the amount of overlap between the minimum enclosing 2D bounding box M_B of the model $M(\mathbf{s})$ on the ground plane and the free-space grid map cells Φ_g given their probability ρ_g of being free space:

$$\log p(\Phi|\mathbf{s}) = \frac{\lambda_{\Phi}}{A_B} \sum_{g=1}^G \log(1 - \rho_g) \cdot o(M_B, \Phi_g). \quad (8)$$

A_B is the area of the model bounding box. The function $o(\cdot, \cdot)$ calculates the overlap between the model bounding box and a cell Φ_g . The factor $\lambda_{\Phi} = \min(1, \frac{l_{\Phi}}{\sigma_M})$ is used to weight this likelihood term based on the grid cell size l_{Φ} and the depth uncertainty σ_M of a stereo-reconstructed point in the distance of the model $M(\mathbf{s})$. This likelihood penalises models that are partly or fully located in areas which are observed as not being occupied by 3D objects.

Orientation prior: To calculate the orientation prior for the model $M(\mathbf{s})$ we use the probability distribution Π_{ϑ} for the vehicle viewpoint inferred by our CNN. We compute the viewpoint ϑ_M from the model orientation θ . The image ray direction ρ is derived from the ray connecting the camera projection center and the center of the vehicle model $M(\mathbf{s})$. The orientation prior is calculated according to

$$\log p(\theta) = \log \Pi_{\vartheta}(\vartheta_M) + \log \left(\frac{1 + \cos(\vartheta_{\text{CNN}} - \vartheta_M)}{2} \right). \quad (9)$$

$\Pi_{\vartheta}(\vartheta_M)$ denotes the probability for the angle ϑ_M according to the output of the viewpoint classification branch of the CNN. As we assume incorrect viewpoint classifications to appear especially between neighboring viewpoints, this term alone is prone to cause small orientation biases. This is why we additionally consider the cosine distance of the most likely viewpoint ϑ_{CNN} predicted from the vehicle CNN and the model viewpoint ϑ_M in this prior.

Inference To find the optimal pose and shape parameters for each detected vehicle we minimize the negative

logarithm of Eq. 4. As this function is non-convex and discontinuous we apply the sequential Monte Carlo sampling approach described in [7] to approximate the parameter set for which the energy function becomes minimal. Starting from an initial state particle, we generate a number of particles in each iteration by jointly sampling the pose and shape parameters from a uniform distribution centered at the preceding parameter values. In contrast to [7], the initial particle orientation is derived from the viewpoint estimated by the CNN. For more details we refer the reader to [7].

4. Evaluation

4.1. Test data and test setup

We test and evaluate our proposed method on two datasets, the KITTI 3D object detection benchmark [11] and our own *StereoVehicle* benchmark. The official KITTI evaluation metrics are designed to assess the joint performance of both, detection and pose estimation. As our approach only focuses on the latter and because we want to obtain further insights in the performance of our algorithm, we only use the KITTI training set with known annotations to evaluate our approach using own evaluation metrics. It consists of 7481 stereo iamges and provides the 3D object location and the orientation for every vehicle. It distinguishes three levels of difficulty (*easy*, *moderate* and *hard*), which mainly depend on the level of object occlusion and truncation.

4.1.1 StereoVehicle dataset

For the acquisition of our dataset we equipped a vehicle with a calibrated and synchronised stereo camera rig using a baseline of 0.85 m and recorded in total more than one hour of data during different day times in urban environments. Further details can be found in [36]. Compared to the KITTI benchmark [11], our dataset has a larger image size (1936x1216), wider field of view using a focal length of 5 mm, a significantly larger baseline, and a higher frame rate of 25 fps. To evaluate our approach we labeled 2289 vehicles in 1000 image pairs. In contrast to the KITTI dataset, which only delivers 2D and oriented 3D bounding boxes as references, we manually fitted the most similar model out of a large set of vehicle CAD models to the individual vehicles, and thus deliver the reference shape and the reference vehicle type in addition to its 3D pose. We distinguish between *easy* vehicles, which are fully visible in the images, and *moderate* vehicles, which are occluded or truncated. We intend to make the data publicly available.

4.2. Parameter settings and training

We select the side length l_Φ of the free-space grid cells to be 25 cm. For the number of the eigenvalues and eigenvectors to be considered in the ASM we choose $j \in [1, 2]$,

which we found to be a proper tradeoff between the complexity of the model and the quality of the model approximation. For training our CNN we make use of the dataset provided by the authors of [43], who labeled 36 different vehicle keypoints in a subset of images of the KITTI dataset. (Note that this subset is not used for evaluation.) We crop these images by the provided reference bounding boxes to use them to train our network. The viewpoint classes needed for the viewpoint branch are derived from the groundtruth viewpoint angles according to our viewpoint class definition in Fig. 4. To train the keypoint/wireframe branch, we make use of the 36 labeled landmarks and create the 2D reference heatmaps for the keypoints and wireframe edges as described in Sec. 3.2. We horizontally flip the training images to double the amount of training data by adapting the viewpoint classes and keypoint/wireframe labels accordingly.

4.3. Vehicle reconstruction results

To evaluate the vehicle reconstruction, we compare the resulting pose parameters from each fitted 3D vehicle model and the reference data for location and orientation of the vehicles. We report results for position estimates whose euclidean distance from the reference position is smaller than 0.75 m. To evaluate the orientation, we show results in three stages (θ_5 , θ_{10} and $\theta_{22.5}$), in which we consider an orientation to be correct if its difference from the reference is less than 5° , 10° and 22.5° , respectively. Additionally, we calculate the average errors ε_t and ε_θ for position and orientation from all vehicle reconstructions.

4.3.1 Results on the KITTI benchmark

The percentage of vehicles in the KITTI dataset detected by the mRCNN is shown in the leftmost column of Tab. 1. Besides, the resulting values for the described pose evaluation metrics are shown in the second rightmost column, compared to state-of-the-art results [7] reported in the rightmost column. Furthermore, we report results depending on the vehicle’s distance from the camera. Additionally, we show the depth uncertainty σ_x of a stereo reconstructed 3D point assuming an uncertainty of disparity of 1 px. Throughout the three levels of difficulty, we achieve a total percentage between 77.5 and 80.6% of correct position estimates (cf. Tab. 1). However, it can be noticed that while the percentage of correct position estimates lies even between 96.3 and 97.1% for vehicles having a distance between 5 and 10 m, the amount decreases to 52.8%-56.7% for vehicles further away from the camera than 20 m. Accordingly, the average error of position estimates also increases drastically and more than doubles from around 30 cm for vehicles in a distance between 5 and 10 m to more than 76 cm for vehicles being more distant than 20 m. We suspect the

increasing depth uncertainty of distant 3D points to be responsible for this effect. Compared to the position estimates, the number of correct orientation estimates is higher with $\theta_{22.5} = 98.9\%$ and $\theta_5 = 86.7\%$ for the easy category (getting worse for the more challenging levels). However, the same effect of decreasing correct results for vehicles in increasing distance is visible for the orientation estimates, although less distinct compared to the position estimates. We also assume this effect to be caused by the increasing uncertainty with increasing distance to the camera. Comparing our results to [7], we obtain very similar values for the position estimates. However, we significantly outperform the orientation estimation results of [7], especially for the moderate and hard levels by up to 5.7% for the θ_5 metric and even up to 9.4% for the $\theta_{22.5}$ metric.

		Vehicle distance				total	[7]
		5-10m	10-15m	15-20m	>20m		
σ_x [cm]		6-25	25-58	58-103	>103		
easy: 98.6%	t [%]	96.8	94.8	76.0	52.8	79.4	80.8
	ϵ_t [m]	0.31	0.34	0.54	0.76	0.49	-
	θ_5 [%]	91.6	90.4	86.8	78.7	86.7	84.8
	θ_{10} [%]	99.4	98.4	96.9	94.4	97.1	93.2
	$\theta_{22.5}$ [%]	99.7	99.6	98.8	97.6	98.9	94.8
	ϵ_θ [°]	2.75	2.87	4.32	5.94	4.02	-
moderate: 95.9%	t [%]	97.1	93.4	76.6	56.7	80.6	80.6
	ϵ_t [m]	0.30	0.37	0.55	0.76	0.50	-
	θ_5 [%]	89.4	85.7	81.6	73.3	82.3	77.8
	θ_{10} [%]	96.6	94.9	91.8	88.8	92.9	86.4
	$\theta_{22.5}$ [%]	98.0	97.5	95.2	93.8	96.1	89.0
	ϵ_θ [°]	4.49	4.88	7.35	9.43	6.6	-
hard: 85.5%	t [%]	96.3	89.7	72.8	53.7	77.5	75.9
	ϵ_t [m]	0.31	0.42	0.63	0.85	0.56	-
	θ_5 [%]	85.7	78.6	74.9	67.4	76.1	70.4
	θ_{10} [%]	93.0	88.3	85.1	81.9	86.7	78.4
	$\theta_{22.5}$ [%]	95.9	92.8	89.3	87.1	91.0	81.6
	ϵ_θ [°]	7.58	9.00	12.30	15.15	11.2	-

Table 1: Overall and distance dependent results of our vehicle reconstruction approach on the KITTI dataset.

4.3.2 Results on the StereoVehicle dataset

The reconstruction results on our *StereoVehicle* dataset are shown in Tab. 2. On the one hand, it is apparent that the overall results for the position estimates are distinctly better compared to the overall results on the KITTI dataset with 97.2% for the easy and 92.8% for the moderate level. Accordingly, the average position error achieved on our dataset is lower compared to the KITTI dataset. A reason for this might be the larger base length and consequently the lower depth uncertainty. However, the orientation estimates and average orientation errors are significantly worse compared to results achieved for the KITTI data, especially for the fine estimations reported in the θ_5 metric. This effect can be caused by the domain gap of our CNN which is trained on KITTI data and therefore performs better on data from the same domain. Also, it has to be noted that the defini-

tion of our *moderate* level covers the definitions of both, the moderate and hard level of the KITTI benchmark.

	t [%]	ϵ_t [m]	θ_5 [%]	θ_{10} [%]	$\theta_{22.5}$ [%]	ϵ_θ [°]
easy	97.2	0.28	74.4	90.4	95.4	10.4
moderate	92.8	0.35	70.1	85.8	90.8	16.8

Table 2: Results of our vehicle reconstruction approach on our *StereoVehicle* dataset.

To obtain a closer look on the distribution of orientation estimation errors, we also show a cumulative histogram for the orientation errors resulting from our own dataset in Fig. 6. According to this, the majority of the incorrectly recovered orientations have an error between 170° - 180° , i.e. they exhibit the opposed viewing direction and thus strongly influence the average orientation error. The reason for this effect might be ambiguities caused by the symmetric shape and appearance of vehicle front and back sides, which in some cases cannot be resolved by our approach.

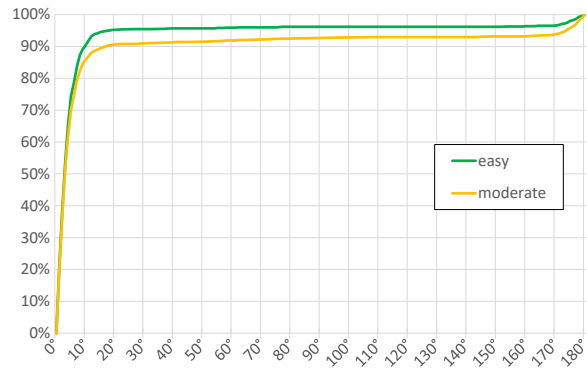


Figure 6: Cumulative histogram of absolute orientation errors on our *StereoVehicle* dataset.

5. Conclusion

In this paper, we proposed a probabilistic approach for vehicle reconstruction, jointly incorporating 3D data, scene knowledge and predictions for vehicle orientation, keypoints and wireframes inferred by our proposed CNN. For the prediction of vehicle orientation we presented a novel hierarchical classification structure, allowing the derivation of a probability distribution to be used as prior. The CNN based detection of vehicle wireframe edges attenuates the dependency on good gradients. These innovations lead to state-of-the-art results on the KITTI object detection benchmark, as well as on our presented *StereoVehicle* dataset, providing precisely fitted vehicle models as reference.

Acknowledgements

This work was supported by the German Research Foundation (DFG) as a part of the Research Training Group i.c.sens [GRK2159].

References

- [1] J. A. Ansari, S. Sharma, A. Majumdar, J. K. Murthy, and K. M. Krishna. The Earth ain't Flat: Monocular Reconstruction of Vehicles on Steep and Graded Roads from a Moving Camera. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 8404–8410, 2018.
- [2] S. Y. Bao, M. Chandraker, Y. Lin, and S. Savarese. Dense Object Reconstruction with Semantic Priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1264–1271, 2013.
- [3] F. Chabot, M. Chaouch, J. Rabarisoa, C. Teuliere, and T. Chateau. Deep MANTA: A Coarse-to-Fine Many-Task Network for Joint 2D and 3D Vehicle Analysis from Monocular Image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1827–1836, 2017.
- [4] X. Chen, K. Kundu, Y. Zhu, A. G. Berneshawi, H. Ma, S. Fidler, and R. Urtasun. 3d Object Proposals for accurate Object Class Detection. In *Advances in Neural Information Processing Systems (NIPS)*, volume 28, pages 424–432, 2015.
- [5] F. Chollet et al. Keras. <https://keras.io>, 2015.
- [6] M. Coenen, F. Rottensteiner, and C. Heipke. Detection and 3D Modelling of Vehicles from terrestrial stereo Image Pairs. In *ISPRS International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XLII-1/W1, pages 505–512, 2017.
- [7] M. Coenen, F. Rottensteiner, and C. Heipke. Precise Vehicle Reconstruction for autonomous Driving Applications. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 21–28, 2019.
- [8] A. Dame, V. A. Prisacariu, C. Y. Ren, and I. Reid. Dense Reconstruction using 3D Object Shape Priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1288–1295, 2013.
- [9] W. Ding, S. Li, G. Zhang, X. Lei, and H. Qian. Vehicle Pose and Shape Estimation through Multiple Monocular Vision. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2018.
- [10] F. Engelmann, J. Stückler, and B. Leibe. Joint Object Pose Estimation and Shape Reconstruction in urban Street Scenes using 3D Shape Priors. In *German Conference on Pattern Recognition (GCPR)*, volume 9796 of *Lecture Notes in Computer Science*, pages 219–230, 2016.
- [11] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous Driving? The KITTI Vision Benchmark Suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012.
- [12] A. Geiger, M. Roser, and R. Urtasun. Efficient Large-Scale Stereo Matching. In *Computer Vision – ACCV 2010*, volume 6492 of *Lecture Notes in Computer Science*, pages 25–38, 2011.
- [13] F. Güney and A. Geiger. Displets: Resolving stereo Ambiguities using Object Knowledge. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4165–4175, 2015.
- [14] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
- [15] M. Hödlmoser, B. Micusik, M. Pollefeys, M. Liu, and M. Kampel. Model-Based Vehicle Pose Estimation and Tracking in Videos Using Random Forests. In *International Conference on 3D Vision (3DV)*, pages 430–437, 2013.
- [16] S. Ioffe and C. Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- [17] A. Kar, S. Tulsiani, J. Carreira, and J. Malik. Category-Specific Object Reconstruction From a Single Image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1966–1974, 2015.
- [18] D. Kingma and L. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [19] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3D Object Representations for Fine-Grained Categorization. In *IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 554–561, 2013.
- [20] A. Kundu, Y. Li, and J. M. Rehg. 3D-RCNN: Instance-level 3D Object Reconstruction via Render-and-Compare. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3559–3568, 2018.
- [21] M. J. Leotta and J. L. Mundy. Vehicle Surveillance with a Generic, Adaptive, 3D Vehicle Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(7):1457–1469, 2011.
- [22] Y. Li, L. Gu, and T. Kanade. Robustly Aligning a Shape Model and Its Application to Car Alignment of Unknown Pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(9):1860–1876, 2011.
- [23] Y.-L. Lin, V. I. Morariu, W. Hsu, and L. S. Davis. Jointly Optimizing 3D Model Fitting and Fine-Grained Classification. In *European Conference on Computer Vision (ECCV)*, pages 466–480, 2014.
- [24] F. Manhardt, W. Kehl, and A. Gaidon. ROI-10D: Monocular Lifting of 2D Detections to 6D Pose and Metric Shape. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [25] A. Mousavian, D. Anguelov, J. Flynn, and J. Koseck. 3D Bounding Box Estimation using Deep Learning and Geometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5632–5640, 2017.
- [26] J. K. Murthy, G. V. S. Krishna, F. Chhaya, and K. M. Krishna. Reconstructing Vehicles from a single Image: Shape Priors for Road Scene Understanding. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 724–731, 2017.
- [27] J. K. Murthy, S. Sharma, and K. M. Krishna. Shape Priors for real-time monocular Object Localization in dynamic Environments. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 1768–1774, 2017.
- [28] A. Newell, K. Yang, and J. Deng. Stacked Hourglass Networks for Human Pose Estimation. In *Computer Vision – ECCV 2016*, pages 483–499, 2016.
- [29] M. Ozuysal, V. Lepetit, and P. Fua. Pose Estimation for Category specific multiview Object Localization. In *2009 IEEE*

- Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pages 778–785, 2009.
- [30] G. Pavlakos, X. Zhou, A. Chan, K. G. Derpanis, and K. Daniilidis. 6-DoF object pose from semantic keypoints. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018, 2017.
- [31] N. Payet and S. Todorovic. From Contours to 3D Object Detection and Pose Estimation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 983–990, 2011.
- [32] V. A. Prisacariu, A. V. Segal, and I. Reid. Simultaneous Monocular 2D Segmentation, 3D Pose Recovery and 3D Reconstruction. In *Computer Vision – ACCV 2012*, pages 593–606, 2012.
- [33] K. Ramnath, S. N. Sinha, R. Szeliski, and E. Hsiao. Car make and model recognition using 3D curve alignment. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 285–292, 2014.
- [34] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015.
- [35] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [36] S. Schön, C. Brenner, H. Alkhatib, M. Coenen, H. Dbouk, N. Garcia-Fernandez, C. Fischer, C. Heipke, K. Lohmann, I. Neumann, U. Nguyen, J.-A. Paffenholz, T. Peters, F. Rotensteiner, J. Schachtschneider, M. Sester, L. Sun, S. Vogel, R. Voges, and B. Wagner. Integrity and Collaboration in Dynamic Sensor Networks. *Sensors*, 18(7), 2018.
- [37] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [38] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [39] R. Wang, N. Yang, J. Stückler, and D. Cremers. DirectShape: Photometric Alignment of Shape Priors for Visual Vehicle Pose and Shape Estimation. In *arxiv*, 2019.
- [40] W. Xiao, B. Vallet, K. Schindler, and N. Paparoditis. Street-Side Vehicle Detection, Classification and Change Detection using mobile Laser Scanning Data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 114:166–178, 2016.
- [41] Z. Yan, H. Zhang, R. Piramuthu, V. Jagadeesh, D. DeCoste, W. Di, and Y. Yu. HD-CNN: Hierarchical Deep Convolutional Neural Networks for Large Scale Visual Recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2740–2748, 2015.
- [42] M. Z. Zia, M. Stark, B. Schiele, and K. Schindler. Detailed 3D Representations for Object Recognition and Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2608–2623, 2013.
- [43] M. Z. Zia, M. Stark, and K. Schindler. Towards Scene Understanding with detailed 3D Object Representations. *International Journal of Computer Vision*, 112(2):188–203, 2015.