

# A Mixed Classification-Regression Framework for 3D Pose Estimation from 2D Images

Siddharth Mahendran<sup>1</sup>  
siddharthm@jhu.edu

Haider Ali<sup>2</sup>  
hali@jhu.edu

René Vidal<sup>1</sup>  
rvidal@cis.jhu.edu

<sup>1</sup> Center for Imaging Science,  
Mathematical Institute for Data Science,  
Johns Hopkins University,  
Baltimore, MD, USA

<sup>2</sup> Department of Computer Science,  
Johns Hopkins University,  
Baltimore, MD, USA

---

## Abstract

3D pose estimation from a single 2D image is an important and challenging task in computer vision with applications in autonomous driving, robot manipulation and augmented reality. Since 3D pose is a continuous quantity, a natural formulation for this task is to solve a *pose regression* problem. However, since pose regression methods return a single estimate of the pose, they have difficulties handling multimodal pose distributions (e.g. in the case of symmetric objects). An alternative formulation, which can capture multimodal pose distributions, is to discretize the pose space into bins and solve a *pose classification* problem. However, pose classification methods can give large pose estimation errors depending on the coarseness of the discretization. In this paper, we propose a mixed classification-regression framework that uses a classification network to produce a discrete multimodal pose estimate and a regression network to produce a continuous refinement of the discrete estimate. The proposed framework can accommodate different architectures and loss functions, leading to multiple classification-regression models, some of which achieve state-of-the-art performance on the challenging Pascal3D+ dataset.

## 1 Introduction

A fundamental problem in computer vision is to understand the underlying 3D geometry of the scene captured in a 2D image. This involves describing the scene in terms of the objects present in it (*i.e.* object detection and classification) and predicting the relative rigid transformation between the camera and each object (*i.e.* object pose estimation). The problem of estimating the 3D pose of an object from a single 2D image is an old problem in computer vision, which has seen renewed interest due to its applications in autonomous driving, robot manipulation and augmented reality, where it is very important to reason about objects in 3D. While general 3D pose estimation includes both the 3D rotation and translation between the object and the camera, here we restrict our attention to estimating only the 3D rotation.

Since 3D pose is a continuous quantity, a natural way to estimate it is to setup a pose regression problem: Given a dataset of 2D images of objects and their corresponding 3D

pose annotations, the goal is to learn a regression function (*e.g.* a deep network) that predicts the 3D pose of an object in an image. This requires choosing, *e.g.* a network architecture, a representation for rotation matrices (Euler angles, axis-angles or quaternions) and a loss function (mean squared loss, Huber loss or geodesic loss). However, a disadvantage of regression-based approaches is that they are unimodal in nature (*i.e.* they return a single pose estimate), hence they are unable to properly model multimodal distributions in the pose space which occur for object categories like boat and dining-table that exhibit strong symmetries.

An alternative approach that is able to capture multimodal distributions in the pose space is to setup a pose classification problem. Specifically, we discretize the pose space into bins and, instead of predicting a single pose output, we return a probability vector on the pose labels associated with these bins. This formulation is better at handling cases where two or more competing hypotheses have a high probability of being correct, for example due to symmetry. However, the drawback of this formulation is that we will always have a non-zero pose-estimation error even with perfect pose-classification accuracy due to the discretization process. Moreover, this error might be large if the binning is very coarse.

In this paper, we propose a mixed classification-regression framework that combines the best of both worlds. The proposed framework consists of two main components. The first one is a classification network that predicts a pose label corresponding to a “key pose” obtained by discretizing the pose space. The second one is a regression network that predicts the deviation between the (continuous) object pose and the (discretized) key pose. The outputs of both components are then combined to predict the object pose. The proposed framework is fairly general and can accommodate multiple choices, for the pose classification and regression networks, for the way in which their outputs are combined, for the loss functions used to train them, etc. Experiments on the Pascal3D+ dataset show that our framework gives more accurate estimates of 3D pose than using either pure classification or regression.

**Related work.** There are many non-deep learning methods for 3D pose estimation given 2D images. Due to space constraints, we restrict our review to only methods based on deep networks. The current literature on 3D pose estimation using deep networks can be divided in two groups: (i) methods that predict 2D keypoints from images and then recover the 3D pose from these keypoints, and (ii) methods that directly predict 3D pose from an image.

The first group of methods includes the works of [2, 5, 18, 19, 24]. The works of [18, 24] train on 2D keypoints that correspond to semantic keypoints defined on 3D object models. Given a new image, they predict a probabilistic map of 2D keypoints and recover 3D pose by comparing with some pre-defined object models. In [5], [19] and [2], instead of semantic keypoints, the 3D points correspond to the 8 corners of a 3D bounding box encapsulating the object. The network is trained by comparing the predicted 2D keypoint locations with the projections of the 3D keypoints on the image under ground-truth pose annotations. [5] uses a Huber loss on the projection error to be robust to inaccurate ground-truth annotations and is the current state-of-the-art on the Pascal3D+ dataset [25] to the best of our knowledge.

The second group of methods includes the works of [4, 13, 14, 15, 21, 22, 23]. All these methods, except [13], use the Euler angle representation of rotation matrices to estimate the azimuth, elevation and camera-tilt angles separately. [22], [21] and [4] divide the angles into non-overlapping bins and solve a classification problem, while [23] tries to regress the angles directly with a mean squared loss. [15] proposes multiple loss functions based on regression and classification and concludes that classification methods work better than regression ones. On the other hand, [13] uses the axis-angle and quaternion representations of 3D rotations and optimizes a geodesic loss on the space of rotation matrices.




Figure 1: A high level overview of our problem statement and proposed network architecture

In this work we also use the axis-angle representation but within a mixed classification-regression framework (which we also call *Bin & Delta model*) instead of the pure regression approach of [13]. The proposed framework can be seen as a generalization of [6, 7, 12, 16], which also use a Bin & Delta model to combine classification and regression networks. Specifically, [16] is a variation of the *Geodesic Bin & Delta model* we propose in Eqn. (12) with a  $\cos - \sin$  representation of Euler angles, while [7] is a particular case of the *Simple Bin & Delta model* we propose in Eqn. (11) with a quaternion representation of 3D pose. On the other hand, the quantized regression model of [6, 7] uses the Bin & Delta model to generate dense correspondences between a 3D model and an image for face landmark and human pose estimation. [6] learns a modification of our *Simple Bin & Delta model* in Eqn. (5) with a separate delta network for every facial region while [7] is a particular case of the *Probabilistic Bin & Delta model* we propose in Eqn. (10). [7] also makes a connection between a Bin & Delta model and a mixture of regression experts proposed in [10], where the classification output probability vector acts as a gating function on regression experts.

Broadly speaking, there has been recent interest in trying to design networks and representations that combine classification and regression to model 3D pose but the authors of these different works have treated this as a one-off representation problem. In contrast, we propose a general framework that encapsulates prior models as particular cases.

**Paper outline.** The remainder of the paper is organized as follows. In §2 we define the 3D pose estimation problem along with our network architecture, 3D pose representation and some common-sense baselines that use only regression or classification. In §3 we describe our mixed classification-regression framework and discuss a variety of different models and loss functions that arise from the framework. Finally, in §4 we demonstrate the effectiveness of these models on the challenging Pascal3D+ dataset for the task of 3D pose estimation.

## 2 3D Pose Estimation

Given an image and a bounding box around an object in the image with known object category label, in this paper we consider the problem of estimating its 3D pose. We assume the bounding box around the object is given by an oracle, but the output of an object detection system like [20] can also be used instead. An overview of our problem is shown in Fig. 1(a).

**Network architecture.** We use a standard network architecture shown in Fig. 1(b) where we have a feature network shared across all object categories and a pose network per object category. The final pose output is selected as the output of the pose network corresponding to the input object category label. We use the ResNet-50 network [8, 9] minus the last classification layer as our feature network. Each pose network is essentially a multi-layer perceptron, with fully connected layers containing ReLU nonlinearities and Batch-Normalization layers in between, whose output depends on whether it is a classification, regression or mixed pose network. We describe the proposed Bin & Delta pose network in more detail in §3.

**Representations of 3D pose.** Let  $SO(3) = \{R \in \mathbb{R}^{3 \times 3} : R^T R = I, \det(R) = +1\}$  denote the Special Orthogonal group of rotation matrices of dimension 3. We use the axis-angle representation  $R = \exp(\theta[v]_{\times})$  of a rotation matrix  $R \in SO(3)$  with the corresponding axis-angle vector defined as  $y = \theta v$ , where  $\theta$  is the angle of rotation,  $v = [v_1, v_2, v_3]^T$  ( $\|v\|_2 = 1$ ) is the axis of rotation,  $[v]_{\times} = [[0, -v_3, v_2], [v_3, 0, -v_1], [-v_2, v_1, 0]]$ , and  $\exp$  is the matrix exponential. Assuming  $\theta \in [0, \pi)$  and defining  $R = I_3$  iff  $y = 0$  sets up a bijective mapping between axis-angle vector  $y$  and rotation matrix  $R$ . With the above notation, we can write the input-output map of the network architecture in Fig. 1(b) as  $y = \Phi(x; W)$ , where  $x$  denotes the input image,  $W$  denotes the network weights, and  $y$  denotes the axis angle representation. Technically, the output  $y$  is also a function of the object category label  $c$ , *i.e.*  $y = \Phi(x, c; W)$ . For the sake of simplicity, we will omit this detail in further analysis.

**Regression baseline.** A natural baseline is a regression formulation with a squared Euclidean loss between the ground-truth 3D pose  $y_n^*$  and the predicted pose  $y_n = \Phi_R(x_n; W)$  for image  $x_n$ . This involves solving the following optimization problem  $\mathcal{R}_E$  during training:

$$\mathcal{R}_E : \quad \min_W \frac{1}{N} \sum_n \|y_n^* - \Phi_R(x_n; W)\|_2^2. \quad (1)$$

However, as recommended by [13], it makes more sense to minimize a geodesic loss on the space of rotation matrices instead of the Euclidean loss as it better captures the geometry of the problem. This involves solving the following optimization problem  $\mathcal{R}_G$  during training:

$$\mathcal{R}_G : \quad \min_W \frac{1}{N} \sum_n \mathcal{L}_p(y_n^*, \Phi_R(x_n; W)), \quad (2)$$

where  $\mathcal{L}_p(y_1, y_2) \equiv \mathcal{L}(R_1, R_2) = \frac{1}{\sqrt{2}} \|\log(R_1 R_2^T)\|_F$  is the geodesic distance between two axis-angle vectors  $y_1$  and  $y_2$  with corresponding rotation matrices  $R_1$  and  $R_2$  respectively.

**Classification baseline.** An alternative to the regression baselines  $\mathcal{R}_E$  and  $\mathcal{R}_G$  is a classification baseline  $\mathcal{C}$ . Here, we first run K-Means clustering on the ground-truth pose annotations  $\{y_n^*\}_{n=1}^N$  to obtain two things: (i) a pose label  $l_n^*$  associated with every image  $x_n$  and (ii) a K-Means dictionary  $\{z_k\}_{k=1}^K$  which also acts as the set of key poses in the discretization process. We then train a network to predict the pose labels by minimizing the cross-entropy loss  $\mathcal{L}_c$  between the ground-truth pose labels  $l_n^*$  and the predicted pose labels  $\Phi_C(x_n; W)$ :

$$\mathcal{C} : \quad \min_W \frac{1}{N} \sum_n \mathcal{L}_c(l_n^*, \Phi_C(x_n; W)). \quad (3)$$

## 3 Mixed Classification-Regression Framework

This section presents the proposed classification-regression framework. §3.1 presents an overview of the Bin & Delta model and its network architecture, while §3.2-§3.5 present various loss functions that arise from different modeling choices within our general framework.

### 3.1 Overview of the Bin & Delta model

Instead of a single multi-layer perceptron as the pose network like we have in our regression and classification baselines, the Bin & Delta model has two components: a bin network and a delta network as shown in Fig. 2. They both take as input the output of

the feature network,  $f_n = \Phi_F(x_n; W_F)$ , where  $\Phi_F$  is the feature network parameterized by weights  $W_F$  and  $x_n$  is the input image. Given feature input  $f_n$ , the bin network predicts a pose label,  $l_n = \Phi_B(f_n; W_B)$ , where  $\Phi_B$  is the bin network parameterized by weights  $W_B$ .

The pose label  $l_n$  references a key pose  $z_{l_n}$  in the discretization process. The delta network predicts a pose residual,  $\delta y_n = \Phi_D(f_n; W_D)$ , where  $\Phi_D$  is the delta network parameterized by weights  $W_D$ . The discrete pose  $z_{l_n}$  and the pose residual  $\delta y_n$  are combined to predict the continuous pose  $y_n$  as




Figure 2: Overview of the Bin&Delta model

$$y_n = g(z_{l_n}, \delta y_n), \quad (4)$$

where different choices for  $g$  lead to different ways of combining the outputs of the Bin & Delta models. We denote all the parameters of the Bin & Delta model as  $W = [W_F, W_B, W_D]$ .

The above framework for combining classification and regression is very general and there are many design choices that lead to different models and loss functions. For example:

1. *How to combine the classification and regression outputs?* Choosing the function  $g$  to be the addition operation, *i.e.*  $y_n = z_{l_n} + \delta y_n$ , leads to our models in §3.2, §3.3 and §3.5. Alternatively, taking the log of the product of the rotations associated to the outputs of the Bin & Delta models, *i.e.*  $y_n = \log(\exp(z_{l_n}) \exp(\delta y_n))$ , leads to our model in §3.4.
2. *Where to apply the regression loss?* We can choose to provide supervision during training at either the final pose output or at the intermediate output of the delta network. The former leads to our model in §3.3 and the later leads to our model in §3.2.
3. *Hard or soft assignment in the pose-binning step?* Instead of assigning a single pose label for every image (a hard assignment), we can assign a probability vector over pose-bins (a soft assignment). This leads to our model in §3.5.
4. *Single delta network for all pose-bins or one per pose-bin?* This is a decision choice we can make for all the models in §3.2-§3.5 and we discuss it in §3.6.

Also, note that even though we present everything in the context of axis-angle representation of 3D pose, all our proposed models can be generalized to any choice of pose representation.

## 3.2 Simple/Naive Bin & Delta

Given training data  $\{x_n, y_n^*\}_{n=1}^N$ , of images  $x_n$  and corresponding ground-truth pose-targets  $y_n^*$ , we run the K-Means discretization process outlined in the classification baseline to associate a pose label  $l_n^*$  with every image. Given this label and the key poses  $\{z_k\}_{k=1}^K$ , we can obtain a ground-truth delta  $\delta y_n^* = y_n^* - z_{l_n^*}$  for every image. Now, the Bin & Delta networks can be trained on modified training data  $\{x_n, l_n^*, \delta y_n^*\}_{n=1}^N$  with a cross-entropy loss  $\mathcal{L}_c$  for the bin network and a Euclidean loss  $\|\cdot\|_2^2$  for the delta network. More specifically, the parameters  $W$  are learned by solving the following optimization problem:

$$\mathcal{M}_S: \quad \min_W \frac{1}{N} \sum_n [\mathcal{L}_c(l_n^*, l_n) + \alpha \|\delta y_n^* - \delta y_n\|_2^2], \quad (5)$$

where  $\alpha \geq 0$  is a relative weighting parameter for balancing the two losses.

### 3.3 Geodesic Bin & Delta

The Simple Bin & Delta model penalizes incorrect predictions in the individual Bin & Delta networks. It is not cognizant of the fact that what we care about is the final predicted pose. To address this issue, we propose a new Bin & Delta model that regresses the final output pose instead of the intermediate output of the delta network. We call this a Geodesic Bin & Delta model because we apply a geodesic regression loss  $\mathcal{L}_p$  between the ground-truth pose and predicted pose by solving the following optimization problem:

$$\mathcal{M}_G : \quad \min_W \frac{1}{N} \sum_n [\mathcal{L}_c(l_n^*, l_n) + \alpha \mathcal{L}_p(y_n^*, z_{l_n} + \delta y_n)]. \quad (6)$$

Notice that that this model has strong connections to the regression baseline  $\mathcal{R}_G$ , except that we now model multimodal 3D pose-distributions and have an additional classification loss. As noted in [13], the geodesic loss is non-convex with many local minima and a good initialization is required. We initialize the networks by training on problem  $\mathcal{M}_S$  for 1 epoch.

### 3.4 Log-Euclidean Bin & Delta

In the two models discussed so far, we have assumed that  $\delta y_n = y_n - z_{l_n}$ , which does not truly capture the geometry of the rotation group  $SO(3)$ . Technically, the delta model  $\delta y_n$  should capture the deviations between the continuous pose  $R_n$  and the discrete pose  $\tilde{R}_k = \exp(z_k)$  associated with key pose  $z_k$ . Specifically,  $\delta y_n$  should be a tangent vector at  $\tilde{R}_k$  in the direction of  $R_n$  whose norm is equal to the geodesic distance between  $\tilde{R}_k$  and  $R_n$ . Mathematically, this is expressed using exponential and logarithm maps as:

$$\delta y_n = \log_{\tilde{R}_k} (R_n) = \log(\tilde{R}_k^T R_n) \text{ or } R_n = \tilde{R}_k \exp(\delta y_n). \quad (7)$$

The geodesic loss between the ground-truth and predicted rotations  $\mathcal{L}(R_n^*, \tilde{R}_k \exp(\delta y_n)) = \mathcal{L}(\tilde{R}_k^T R_n^*, \exp(\delta y_n))$  can be approximated by the Euclidean distance on the tangent space at the identity as  $\|\log(\tilde{R}_k^T R_n^*) - \log(\exp(\delta y_n))\|_2$ . This new regression loss gives us the Log-Euclidean Bin & Delta model, which is based on solving the following optimization problem

$$\mathcal{M}_{LE} : \quad \min_W \frac{1}{N} \sum_n [\mathcal{L}_c(l_n^*, l_n) + \alpha \|\log(\tilde{R}_k^T R_n^*) - \delta y_n\|_2^2], \quad (8)$$

where the term  $\log(\tilde{R}_k^T R_n^*)$  can be precomputed for efficiency of training. Fig. 3 shows a toy example of our proposed model in the context of a circle. We show a circle with 5 tangent planes corresponding to key poses  $\tilde{R}_i, i = 1, \dots, 5$ . The rotation  $R$  (shown in red) is now a combination of the key pose  $\tilde{R}_1$ , with pose label  $l = 1$ , and the delta  $\delta y$  (shown in orange).




Figure 3: Toy example of the Log-Euclidean Bin & Delta model for a circle

### 3.5 Probabilistic Bin & Delta

In all the models discussed so far we have used a deterministic (hard) assignment obtained from K-Means in which we assign a single key-pose to an image. A more flexible and possibly more informative model would be to do a probabilistic (soft) assignment to all key-poses. Specifically, post K-Means we can generate a probabilistic assignment as:

$$p_{nk}^* = \frac{\exp(-\gamma \|y_n^* - z_k\|_2^2)}{\sum_k \exp(-\gamma \|y_n^* - z_k\|_2^2)}. \quad (9)$$

Now, the classification loss can be modified to be a Kullback-Leibler (KL) divergence between ground-truth and predicted probabilities. The regression loss is also updated to fully utilize this probabilistic output as shown in the following optimization problem:

$$\mathcal{M}_P : \quad \min_W \frac{1}{N} \sum_n \left[ \mathcal{L}_{KD}(p_n^*, p_n) + \alpha \sum_k p_{nk} \mathcal{L}_p(y_n^*, z_k + \delta y_n) \right]. \quad (10)$$

The predicted pose is now  $y_n = z_{l_n} + \delta y_n$ , where  $l_n = \operatorname{argmax}_k p_{nk}$ . Another variation that is of relevance here is that instead of using K-Means followed by the probabilistic assignment of Eqn. (9), one could learn a Gaussian Mixture Model (GMM) in the pose-space to do the soft assignment in a more natural way. We did not explore this variation but mention it to demonstrate that many more models can be described as particular cases of our framework.

### 3.6 One delta network per pose-bin

An implicit assumption in all the models discussed so far is that there is a single delta model for all pose-bins. An alternative modeling choice, as shown in Fig. 4, is to have a delta model for every single pose-bin. This modeling decision is equivalent to deciding whether to have a common covariance matrix across all clusters or have a different covariance matrix for every cluster in a GMM. If we choose to now have one delta network per pose-bin, we can update all the previous optimization problems as follows (the change is highlighted in red):




Figure 4: One delta network per pose-bin

to now have one delta network per pose-bin, we can update all the previous optimization problems as follows (the change is highlighted in red):

$$\mathcal{M}_{S+} : \quad \min_W \frac{1}{N} \sum_n \left[ \mathcal{L}_c(l_n^*, l_n) + \alpha \|\delta y_n^* - \delta y_n^{l_n}\|_2^2 \right], \quad (11)$$

$$\mathcal{M}_{G+} : \quad \min_W \frac{1}{N} \sum_n \left[ \mathcal{L}_c(l_n^*, l_n) + \alpha \mathcal{L}_p(y_n^*, z_{l_n} + \delta y_n^{l_n}) \right], \quad (12)$$

$$\mathcal{M}_{LE+} : \quad \min_W \frac{1}{N} \sum_n \left[ \mathcal{L}_c(l_n^*, l_n) + \alpha \|\log(\tilde{R}_{l_n}^T R_n^*) - \delta y_n^{l_n}\|_2^2 \right], \quad (13)$$

$$\mathcal{M}_{P+} : \quad \min_W \frac{1}{N} \sum_n \left[ \mathcal{L}_{KD}(p_n^*, p_n) + \alpha \sum_k p_{nk} \mathcal{L}_p(y_n^*, z_k + \delta y_n^k) \right]. \quad (14)$$

## 4 Results

First, we describe the Pascal3D+ dataset [25], which is the benchmark dataset used for evaluating 3D pose estimation methods. Then, we demonstrate the effectiveness of our framework with state-of-the-art performance on this challenging task. Finally, we present an ablation study on two hyper-parameters: size of the K-Means dictionary  $K$  and relative weight  $\alpha$ .

**Dataset.** The Pascal3D+ consists of images of twelve object categories: aeroplane (aero), bicycle (bike), boat, bottle, bus, car, chair, diningtable (dtable), motorbike (mbike), sofa, train and tvmonitor (tv). These images were curated from the Pascal VOC [11] and ImageNet [8] datasets, and annotated with 3D pose in terms of the Euler angles  $(az, el, ct)$ . We use

the ImageNet-trainval and Pascal-train images as our training data and the Pascal-val images as our testing data. Following the protocol of [21, 22] and others, we use ground-truth bounding boxes of un-occluded and un-truncated objects. We use the 3D pose-jittering data augmentation strategy of [13] and the rendered images of [21] to augment our training data.

**Evaluation metrics.** We evaluate our models on the Pascal3D+ dataset under two standard metrics: (i) median angle error across all test images *MedErr* and (ii) percentage of images that have angle error less than  $\frac{\pi}{6}$ ,  $Acc_{\frac{\pi}{6}}$ . Here the angle error is the angle between predicted and ground-truth rotation matrices given by  $\frac{1}{\sqrt{2}} \|\log(R_1 R_2^T)\|_F$ . For all the results, we run each experiment three times and report the mean and standard deviation across three trials.

**3D pose estimation using Bin & Delta models.** As can be seen in Tables 1, 2 and Fig. 5, we achieve state-of-the-art performance with our Geodesic and Probabilistic Bin & Delta models with one delta network per pose-bin, namely models  $\mathcal{M}_{G+}$  and  $\mathcal{M}_{P+}$ , respectively. We improve upon the existing state-of-the-art [9], from  $10.88^\circ$  to  $10.10^\circ$  in the *MedErr* metric and from 0.8392 to 0.8588 in the  $Acc_{\frac{\pi}{6}}$  metric under model  $\mathcal{M}_{G+}$ . Due to space constraints Table 2 shows results for four of our models only. Please see the supplementary material for additional results. Other important observations include: (1) Models  $\mathcal{M}_{G+}$  and  $\mathcal{M}_{P+}$  that have one delta network per pose-bin perform better than the models  $\mathcal{M}_G$  and  $\mathcal{M}_P$  with a single delta network for all pose-bins, as they have more freedom in modeling the deviations from key pose even though the discretization of the pose space is coarser. (2) Models that use a geodesic loss on the final pose output ( $\mathcal{M}_G$ ,  $\mathcal{M}_P$ ,  $\mathcal{M}_{G+}$  and  $\mathcal{M}_{P+}$ ) perform better than the pure classification & regression baselines  $\mathcal{C}$ ,  $\mathcal{R}_E$ ,  $\mathcal{R}_G$  as well as models that put losses on the individual components ( $\mathcal{M}_S$ ,  $\mathcal{M}_R$ ,  $\mathcal{M}_{S+}$  and  $\mathcal{M}_{P+}$ ). (3) Pure regression with a geodesic loss  $\mathcal{R}_G$  worked better than pure classification  $\mathcal{C}$  under the *MedErr* metric but worse under the  $Acc_{\frac{\pi}{6}}$  metric. (4) Geodesic regression  $\mathcal{R}_G$  is better than Euclidean regression  $\mathcal{R}_E$ .

**Ablation analysis.** There are two main hyper-parameters in our model, (i) the size of the K-Means dictionary  $K$  and (ii) the relative weighting parameter  $\alpha$ . In Tables 3 and 4, we vary the size of the K-Means dictionary for models  $\mathcal{M}_G$  ( $K = \{24, 50, 100, 200\}$ ) and  $\mathcal{M}_{G+}$  ( $K = \{4, 8, 16, 24\}$ ). We see that usually, a larger K-Means dictionary is better. We also note that even with a very coarsely discretized pose space ( $K = 4$ ),  $\mathcal{M}_{G+}$  is better than the very highly discretized pose space ( $K = 200$ ) of  $\mathcal{M}_G$ . In Tables 5 and 6, we tried three different values of  $\alpha = \{0.1, 1, 10\}$  for models  $\mathcal{M}_G$  and  $\mathcal{M}_{G+}$  and observed that performance improves with higher  $\alpha$ . This is because higher  $\alpha$  gives more importance to the geodesic loss on the final predicted pose which is the quantity we care about.

**Implementation details.** Please see the supplementary material for more details.

## 5 Conclusion

3D pose estimation is an important but challenging task and current deep learning solutions solve this problem using pure regression or classification approaches. We designed a framework that combines regression and classification loss functions to predict fine-pose estimates while modeling multi-modal pose distributions. This was implemented using a flexible Bin & Delta network architecture where different modeling choices led to different models and loss functions. We analyzed these models on the Pascal3D+ dataset and demonstrated state-of-the-art performance using two of them. We also tested various modeling choices and provided some guidelines for future work using these models.

**Acknowledgement.** This research was supported by NSF grant 1527340.



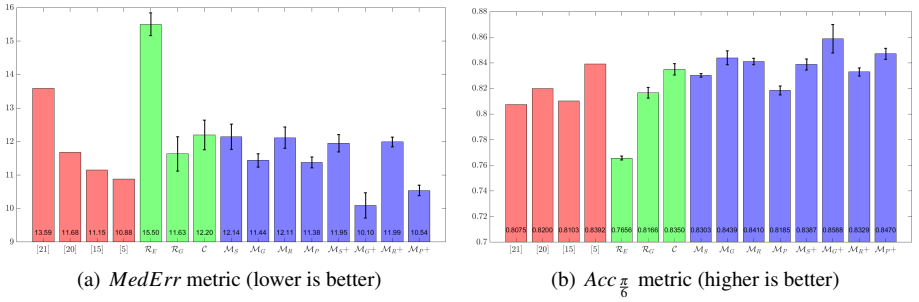


Figure 5: Performance on the Pascal3D+ dataset under two metrics (averaged over all object categories). Red bars are current state-of-the-art methods, green bars are our baselines and blue bars are our proposed models (best seen in color)

Method	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
$\square$	13.8	17.7	21.3	12.9	5.8	9.1	14.8	15.2	14.7	13.7	8.7	15.4	13.59
$\square$	15.4	<b>14.8</b>	25.6	9.3	3.6	6.0	<b>9.7</b>	<b>10.8</b>	16.7	<b>9.5</b>	6.1	12.6	11.68
$\square$	13.6	<b>12.5</b>	22.8	8.3	3.1	5.8	11.9	12.5	<b>12.3</b>	12.8	6.3	11.9	11.15
$\square$	<b>10.0</b>	15.6	<b>19.1</b>	8.6	3.3	<b>5.1</b>	13.7	11.8	<b>12.2</b>	13.5	6.7	<b>11.0</b>	10.88
$\mathcal{R}_E$	14.5	17.7	39.3	7.4	4.0	7.8	15.2	26.6	17.5	10.5	11.5	14.1	15.50
$\mathcal{R}_G$	11.8	15.9	27.2	7.2	2.9	5.2	11.6	15.0	14.3	10.8	<b>5.4</b>	12.4	11.63
$\mathcal{C}$	11.7	15.3	21.5	9.3	4.1	7.4	11.2	17.8	17.0	11.0	7.0	13.1	12.20
$\mathcal{M}_S$	11.0	15.5	21.0	8.8	3.8	7.0	10.8	21.0	16.6	10.7	6.5	13.1	12.14
$\mathcal{M}_G$	10.6	16.4	21.6	8.1	3.2	6.0	9.9	14.6	16.0	11.1	6.3	13.4	11.44
$\mathcal{M}_R$	12.8	15.2	23.4	9.0	4.0	7.4	11.1	16.8	16.1	10.7	6.6	12.3	12.11
$\mathcal{M}_P$	11.4	16.3	25.6	<b>7.0</b>	<b>2.6</b>	5.1	11.3	16.0	13.6	10.2	5.5	12.0	11.38
$\mathcal{M}_{S+}$	12.2	15.7	24.4	9.9	3.6	6.5	12.0	14.8	14.4	11.9	6.4	<b>11.6</b>	11.95
$\mathcal{M}_{G+}$	<b>8.5</b>	<b>14.8</b>	<b>20.5</b>	<b>7.0</b>	3.1	<b>5.1</b>	<b>9.3</b>	<b>11.3</b>	14.2	10.2	5.6	11.7	<b>10.10</b>
$\mathcal{M}_{R+}$	12.3	16.7	24.7	7.5	3.6	6.5	11.5	15.5	15.1	11.1	7.3	12.1	11.99
$\mathcal{M}_{P+}$	10.6	15.0	23.9	<b>6.7</b>	<b>2.7</b>	<b>4.7</b>	9.8	12.6	13.9	<b>9.7</b>	<b>5.3</b>	11.7	<b>10.54</b>

Table 1: Performance of our models under the *MedErr* metric (lower is better). Best results are highlighted in bold and the second best are shown in red (best seen in color).

Method	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
$\square$	0.81	0.77	0.59	0.93	<b>0.98</b>	0.89	0.80	0.62	<b>0.88</b>	0.82	0.80	0.80	0.8075
$\square$	0.74	<b>0.83</b>	0.52	0.91	0.91	0.88	0.86	<b>0.73</b>	0.78	0.90	<b>0.86</b>	<b>0.92</b>	0.8200
$\square$	0.78	<b>0.83</b>	0.57	0.93	0.94	0.90	0.80	0.68	0.86	0.82	0.85	0.85	0.8103
$\square$	0.83	0.82	<b>0.64</b>	0.95	0.97	0.94	0.80	0.71	<b>0.88</b>	0.87	0.80	0.86	0.8392
$\mathcal{R}_E$	0.77	<b>0.75</b>	0.41	0.96	0.91	0.83	0.72	0.56	0.75	0.90	0.75	0.87	0.7656
$\mathcal{R}_G$	0.80	0.78	0.54	<b>0.97</b>	0.95	0.93	0.83	0.59	0.82	0.91	0.81	0.86	0.8166
$\mathcal{C}$	0.84	0.77	0.60	0.95	0.97	<b>0.95</b>	0.90	0.63	0.78	0.94	0.81	0.87	0.8350
$\mathcal{M}_S$	0.83	0.78	0.61	0.96	0.96	0.94	0.90	0.56	0.79	0.95	0.82	0.87	0.8303
$\mathcal{M}_G$	0.84	0.76	0.62	0.96	<b>0.98</b>	0.94	0.92	0.65	0.80	0.96	0.82	0.87	0.8439
$\mathcal{M}_{LE}$	0.83	0.77	0.58	0.96	0.96	0.94	0.91	0.71	0.81	0.93	0.81	0.87	0.8410
$\mathcal{M}_P$	0.80	0.77	0.56	<b>0.97</b>	0.97	0.93	0.82	0.57	0.81	0.92	0.82	0.88	0.8185
$\mathcal{M}_{S+}$	0.82	0.80	0.59	0.94	0.97	0.94	0.91	0.63	0.81	<b>0.97</b>	0.83	0.87	0.8387
$\mathcal{M}_{G+}$	<b>0.87</b>	0.81	<b>0.64</b>	0.96	0.97	<b>0.95</b>	<b>0.92</b>	0.67	0.85	<b>0.97</b>	0.82	0.88	<b>0.8588</b>
$\mathcal{M}_{LE+}$	0.81	0.77	0.56	0.96	0.97	0.92	0.86	<b>0.73</b>	0.79	0.93	0.80	0.89	0.8329
$\mathcal{M}_{P+}$	0.84	0.82	0.59	<b>0.97</b>	0.97	<b>0.95</b>	0.88	0.68	0.84	0.93	0.81	0.89	0.8470

Table 2: Performance of our models under the  $Acc_{\frac{\pi}{6}}$  metric (higher is better). Best results are highlighted in bold.

Metric	K	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
<i>MedErr</i>	24	12.4	16.3	23.5	8.7	2.7	5.4	11.6	17.4	16.3	13.7	6.2	15.6	12.48
	50	12.9	16.0	21.1	8.3	3.3	6.1	11.2	22.2	17.7	11.8	5.8	13.9	12.53
	100	12.1	16.0	19.9	8.9	3.4	6.5	10.8	15.2	16.4	9.6	5.9	13.0	11.48
	200	10.6	16.4	21.6	8.1	3.2	6.0	9.9	14.6	16.0	11.1	6.3	13.4	11.44
$Acc_{\frac{\pi}{6}}$	24	0.84	0.82	0.58	0.95	0.97	0.94	0.89	0.57	0.77	0.91	0.81	0.86	0.8266
	50	0.82	0.80	0.59	0.95	0.97	0.94	0.92	0.54	0.78	0.91	0.83	0.89	0.8281
	100	0.83	0.76	0.63	0.96	0.97	0.93	0.91	0.57	0.78	0.95	0.82	0.88	0.8335
	200	0.84	0.76	0.62	0.96	0.98	0.94	0.92	0.65	0.80	0.96	0.82	0.87	0.8439

Table 3: Ablation analysis of the size of K-Means dictionary in model  $\mathcal{M}_G$  under the *MedErr* (first row) and  $Acc_{\frac{\pi}{6}}$  (second row) metrics.

Metric	K	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
<i>MedErr</i>	4	10.4	13.3	21.9	7.2	2.9	5.3	9.9	16.3	14.1	10.4	5.0	12.5	10.78
	8	10.5	14.8	21.5	6.8	2.7	4.9	9.7	16.1	14.9	10.2	5.6	12.3	10.85
	16	9.9	14.3	21.3	7.3	2.7	4.9	9.6	13.0	14.7	10.8	5.2	11.7	10.46
	24	9.7	15.3	23.5	7.1	2.9	5.0	10.0	13.3	14.4	11.3	5.3	13.1	10.91
$Acc_{\frac{\pi}{6}}$	4	0.85	0.80	0.61	0.97	0.97	0.95	0.87	0.67	0.84	0.93	0.83	0.86	0.8453
	8	0.83	0.79	0.60	0.97	0.96	0.95	0.91	0.62	0.81	0.95	0.83	0.89	0.8427
	16	0.84	0.82	0.61	0.96	0.98	0.96	0.92	0.67	0.82	0.97	0.82	0.90	0.8553
	24	0.87	0.80	0.60	0.96	0.97	0.95	0.90	0.65	0.83	0.94	0.82	0.87	0.8467

Table 4: Ablation analysis of the size of K-Means dictionary in model  $\mathcal{M}_G+$  under the *MedErr* and  $Acc_{\frac{\pi}{6}}$  metrics.

Metric	$\alpha$	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
<i>MedErr</i>	0.1	11.8	16.1	20.8	8.4	3.3	6.4	10.6	28.5	15.0	11.2	6.0	12.3	12.53
	1	12.1	16.0	19.9	8.9	3.4	6.5	10.8	15.2	16.4	9.6	5.9	13.0	11.48
	10	12.1	14.5	22.8	8.7	3.1	6.5	10.9	15.1	16.3	10.6	6.0	13.0	11.63
$Acc_{\frac{\pi}{6}}$	0.1	0.84	0.77	0.62	0.96	0.96	0.94	0.91	0.51	0.82	0.96	0.81	0.88	0.8306
	1	0.83	0.76	0.63	0.96	0.97	0.93	0.91	0.57	0.78	0.95	0.82	0.88	0.8335
	10	0.82	0.79	0.59	0.96	0.97	0.94	0.91	0.67	0.81	0.95	0.82	0.88	0.8424

Table 5: Ablation analysis of the weighting parameter  $\alpha$  in model  $\mathcal{M}_G$  under the *MedErr* and  $Acc_{\frac{\pi}{6}}$  metrics.

Metric	$\alpha$	aero	bike	boat	bottle	bus	car	chair	dtable	mbike	sofa	train	tv	Mean
<i>MedErr</i>	0.1	10.3	16.0	24.0	7.1	3.2	5.5	10.3	11.8	15.2	10.6	6.0	12.3	11.01
	1	9.9	14.3	21.3	7.3	2.7	4.9	9.6	13.0	14.7	10.8	5.2	11.7	10.46
	10	8.5	14.8	20.5	7.0	3.1	5.1	9.3	11.3	14.2	10.2	5.6	11.7	10.10
$Acc_{\frac{\pi}{6}}$	0.1	0.86	0.81	0.59	0.96	0.98	0.95	0.90	0.65	0.80	0.96	0.82	0.89	0.8473
	1	0.84	0.82	0.61	0.96	0.98	0.96	0.92	0.67	0.82	0.97	0.82	0.90	0.8553
	10	0.87	0.81	0.64	0.96	0.97	0.95	0.92	0.67	0.85	0.97	0.82	0.88	0.8588

Table 6: Ablation analysis of the weighting parameter  $\alpha$  in model  $\mathcal{M}_G+$  under the *MedErr* and  $Acc_{\frac{\pi}{6}}$  metrics.

## References

- [1] The PASCAL Object Recognition Database Collection. <http://www.pascal-network.org/challenges/VOC/databases.html>.
- [2] Alberto Crivellaro, Mahdi Rad, Yannick Verdie, Kwang Moo Yi, Pascal Fua, and Vincent Lepetit. A novel representation of parts for accurate 3d object detection and tracking in monocular images. In *IEEE International Conference on Computer Vision*, 2015.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [4] Mohamed Elhoseiny, Tarek El-Gaaly, Amr Bakry, and Ahmed Elgammal. A comparative analysis and study of multiview CNN models for joint object categorization and pose estimation. In *International Conference on Machine learning*, 2016.
- [5] Alexander Grabner, Peter M. Roth, and Vincent Lepetit. 3d pose estimation and 3d model retrieval for objects in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [6] Riza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [7] Riza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. *coRR abs/1803.02188*, 2018. URL <http://arxiv.org/abs/1803.02188>.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. *CoRR abs/1603.05027*, 2016. URL <http://arxiv.org/abs/1603.05027>.
- [10] Michael I. Jordan and Robert A. Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural Computation*, 6(2):181–214, 1994. doi: 10.1162/neco.1994.6.2.181. URL <https://doi.org/10.1162/neco.1994.6.2.181>.
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [12] Chi Li, Jin Bai, and Gregory D. Hager. A unified framework for multi-view multi-class object pose estimation. *coRR abs/1801.08103*, 2018. URL <http://arxiv.org/abs/1801.08103>.
- [13] Siddharth Mahendran, Haider Ali, and René Vidal. 3D pose regression using convolutional neural networks. In *IEEE International Conference on Computer Vision Workshop on Recovering 6D Object Pose*, 2017.

- [14] Francisco Massa, Mathieu Aubry, and Renaud Marlet. Convolutional neural networks for joint object detection and pose estimation: A comparative study. *CoRR abs/1412.7190*, 2014. URL <http://arxiv.org/abs/1412.7190>.
- [15] Francisco Massa, Renaud Marlet, and Mathieu Aubry. Crafting a multi-task CNN for viewpoint estimation. In *British Machine Vision Conference*, 2016.
- [16] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *CVPR*, 2017.
- [17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [18] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-DoF object pose from semantic keypoints. In *IEEE International Conference on Robotics and Automation*, 2017.
- [19] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *IEEE International Conference on Computer Vision*, 2017.
- [20] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015.
- [21] Hao Su, Charles R. Qi, Yangyan Li, and Leonidas J. Guibas. Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In *IEEE International Conference on Computer Vision*, 2015.
- [22] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [23] Yumeng Wang, Shuyang Li, Mengyao Jia, and Wei Liang. Viewpoint estimation for objects with convolutional neural network trained on synthetic images. In *Advances in Multimedia Information Processing - PCM 2016*, 2016.
- [24] Jiajun Wu, Tianfan Xue, Joseph J Lim, Yuandong Tian, Joshua B Tenenbaum, Antonio Torralba, and William T Freeman. Single Image 3D Interpreter Network. In *European Conference on Computer Vision*, pages 365–382, 2016.
- [25] Roozbeh Mottaghi Yu Xiang and Silvio Savarese. Beyond PASCAL: A benchmark for 3D object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision*, 2014.