

3D Hand Pose Estimation from RGB Using Privileged Learning with Depth Data

Shanxin Yuan*
Huawei Noah’s Ark Lab
shanxin.yuan@huawei.com

Björn Stenger
Rakuten Institute of Technology
bjorn@cantab.net

Tae-Kyun Kim
Imperial College London
tk.kim@imperial.ac.uk

Abstract

This paper proposes a method for 3D hand pose estimation given a large dataset of depth images with joint annotations, and a smaller dataset of depth and RGB image pairs with joint annotations. We explore different ways of using the depth data at the training stage to improve the pose estimation accuracy of a network that only takes RGB images as input. By using paired RGB and depth images, we are able to supervise the RGB-based network to learn middle layer features that mimic that of a network trained on large-scale, accurately annotated depth data. Further, depth data provides accurate foreground masks, which are employed to learn better feature activations in the RGB network. During testing, when only RGB images are available, our method produces accurate 3D hand pose predictions. The method is also shown to perform well on the 2D hand pose estimation task. We validate the approach on three public datasets, and compare it to other published methods.

1. Introduction

3D hand pose estimation accuracy has been greatly improving over the past few years, with new methods being introduced [8, 18, 37, 41, 9] and datasets being made publicly available [7, 30, 32, 34, 43]. There remains an accuracy gap between pose estimation from RGB and depth image input, which several recent works have aimed to narrow [17, 20, 27, 47]. One of the difficulties has been the lack of large-scale realistic RGB datasets with accurate annotations. This has been addressed by creating synthetic datasets [47], or employing GANs to generate training data [16]. In this paper we propose a different approach, using depth data as *privileged information* during training. Fully annotated depth datasets [7, 30, 32, 34, 43] are abundant in the literature, but so far no attempt has been made to use this data to support the task of 3D hand pose estimation

from RGB images. Some RGB-D datasets have been proposed recently to tackle the problem of 3D hand pose estimation from RGB images, however most of existing methods [16, 44, 47] use only RGB images for training.

Training with privileged information [36], also called training with hidden information [39], has been shown to improve performance in other domains, such as image classification [4], object detection [13], and action recognition [25]. In this paper we explore how to use depth data as privileged information to aid 3D hand pose estimation from RGB images. Existing methods for 3D hand pose estimation from RGB images pursue two main directions: (1) using only RGB images for 3D hand pose estimation [16, 44, 47], with different CNN models. Given the limited size of real RGB datasets, a large number of synthetic images [16, 47] are created to help the training, whether they are purely synthetic [47], or using CycleGAN to enforce a certain realism [16]. (2) Using RGB-D images for 3D hand pose tracking [17], where the input is the depth channel in addition to the RGB channels. This works well when the paired RGB and depth images are available at test time. Recently, a depth regularizer module was introduced in [2] to make use of only the paired depth images for weak supervision. The lack of large-scale annotated training data limits the success of this approach. Our study proposes a new framework for 3D hand pose estimation from RGB images, by using the existing large-scale annotated depth data during training. This helps improve 3D hand pose estimation using a single RGB image input at test time.

Our method transfers supervision from depth images to RGB images. We use two networks, an RGB-based network and a depth-based network, see Figure 1. We explore different ways to use depth data: (1) initially, we treat a large amount of independent external depth training data as privileged information to train the depth-based network. (2) After the initial training is completed, paired RGB and depth images are used to tune the RGB-based network and the depth-based network. The idea is to let activations in the middle layer of the RGB network mimic that of the depth network. (3) We use foreground masks from depth

*The work was conducted while Shanxin Yuan was a PhD student at Imperial College London

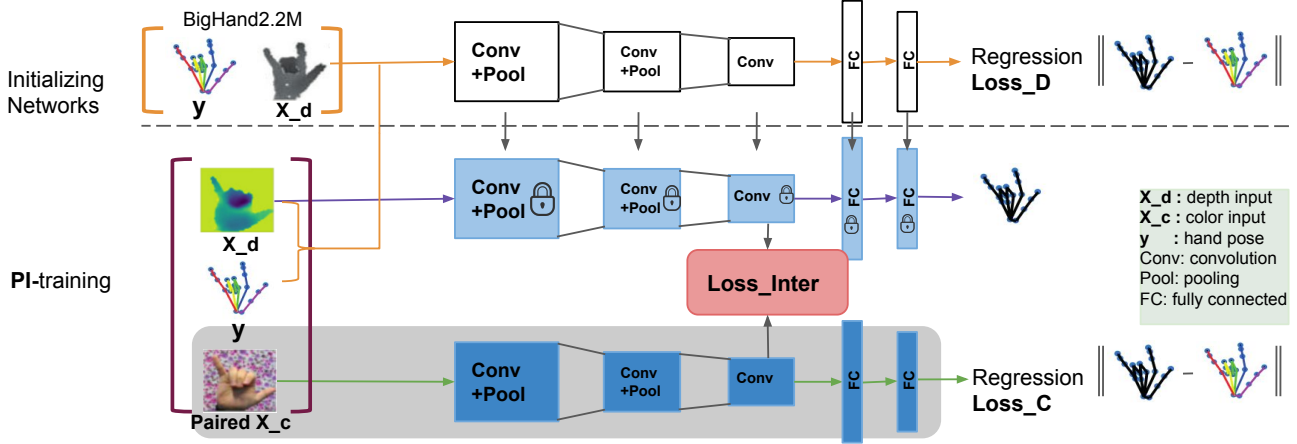


Figure 1. **Proposed framework for 3D hand pose estimation from an RGB image using privileged depth data.** Training proceeds in two stages, a pre-training stage and privileged information (PI)-training stage. In the first stage, a depth-based network (top) and an RGB-based (bottom) network are trained independently to minimize 3D pose loss L_D and L_C . In the second stage, we freeze the parameters of the depth-based network and continue training with paired RGB and depth images, by minimizing a joint loss, which includes L_C and a mid-level feature regression loss L_{inter} .

data to reduce spurious background activations in the RGB network.

Our training strategy can be easily embedded into existing pose estimation methods. We demonstrate this in experiments of 2D hand pose estimation with an RGB image input by a different CNN model. Results on 2D hand pose estimation, using our training strategy show improvement over state-of-the-art methods for 2D hand pose estimation with RGB input.

Comprehensive experiments are conducted on three datasets: the Stereo dataset [44], the RHD dataset [47], and the Dexter-Object dataset [29]. The Stereo dataset and RHD dataset are used for evaluating 3D pose estimation from an RGB input. All three datasets are used for evaluating 2D hand pose estimation from a single RGB image.

2. Related Work

3D hand pose estimation. Hand pose estimation from depth data has made rapid progress in the past years [5, 8, 18, 24, 37], where comprehensive studies [6, 31, 42] have been instrumental in advancing the field. Random forests [32, 38] and CNNs [8, 34, 37, 41] trained on large-scale public depth image datasets [7, 30, 32, 34, 43] have shown good performance. A recent benchmark evaluation [42] showed that modern methods achieve mean 3D joint position errors of less than 10mm. Hand pose estimation from RGB images is significantly more challenging [20, 27, 47]. Due to the difficulty in capturing real RGB datasets with accurate 3D annotations, recent methods employ synthetic data [47], or more realistic *GAN*-erated images [16]. Mueller *et al.* [16] use an image-to-image

translation network to create a large amount of RGB training images and combine a CNN with a kinematic 3D hand model for pose estimation. The method requires a predefined hand model, adapted for each user. Simon *et al.*'s *OpenPose* [27] system generates an annotated RGB dataset using a panoptic studio setup, using multiple views to bootstrap 2D hand pose estimation. Zimmermann and Brox [47] proposed combining hand segmentation and 2D hand pose estimation (using *CPM* [40]), followed by estimating 3D hand pose relative to a canonical pose. Panteleris and Argyros [20] estimate absolute 3D hand pose by first estimating 2D hand pose and then optimizing a 3D hand model with inverse kinematics. Note that there also exists a large body of work on the related task of recovering full 3D human body pose from images. One line of work aims to directly estimate the 3D pose from images [35, 45]. A second approach is to first estimate 2D pose, often in terms of joint locations, and then lift this to 3D pose. 2D key points can be reliably estimated using CNNs and 3D pose is estimated using structured learning or a kinematic model [26, 33, 46].

Learning with privileged information and transfer learning. Privileged information denotes training data that is available only during training but not at test time. The concept to provide teacher-like supervision at training time was introduced by Vapnik and Vashist [36]. The idea has proven useful in other domains [4, 13, 25]. Shi *et al.* [25] treated skeleton data as privileged information in CNN-RNNs for action recognition from depth sequences. Chen *et al.* [4] manually annotated object masks in 10% of the training data and treated these as privileged information for image classification. The idea is related to network compression and mimic learning proposed by Ba and Caruana [1] as

well as network distillation by Hinton *et al.* [12], where intermediate layer outputs of one network are approximated by another, possibly smaller, network. These techniques can be used to significantly reduce the number of model parameters without a significant drop in accuracy. In our case, the application target is similar to transfer learning for domain adaptation. Information from one task, prediction from depth images, is shared with another, prediction from RGB images. In transfer learning and domain adaptation information is shared across different data modalities [3, 13, 23]. Chen *et al.* [3] proposed recognition in RGB images by learning from labeled RGB-D data. A common feature representation is learned across two feature modalities. Hoffman *et al.* [13] learned an additional *hallucination* representation, which is informed by the depth data in training. At testing, it used the softmax to select the final prediction between the predictions from the hallucination representation and the predictions from RGB representation. Luo *et al.* [15] recently proposed graph distillation for action detection with privileged modalities (RGB, depth, skeleton, and flow), where a novel graph distillation layer was used to dynamically learn to distill knowledge from the most effective modality, depending on the type of action. Cai *et al.* [2] proposed a weakly-supervised method for 3D hand pose estimation from RGB image by introducing an additional depth regularizer module, which rendered a depth image from the estimated 3D hand pose. Training was conducted by minimizing an additional loss term, which is the $L1$ distance between the rendered depth image and the ground truth depth image. In our case, we use paired depth and RGB images during training. Depth and RGB networks are first trained separately. Subsequently the RGB network are progressively updated, while the depth network parameters remain fixed.

Learning a latent space representation. Latent space representations shows promise for 3D hand pose estimation from RGB images [14, 28]. Spurr *et al.* [28] learned a cross-modal statistical hand model via a latent space representation that embeds sample points from multiple data sources such as 2D keypoints, images, and 3D hand poses. Multiple encoders were used to project different data modalities into a unified low-dimensional latent space, where a set of decoders reconstruct the hand configuration in either modality. Iqbal *et al.* [14] used latent 2.5D heatmaps, containing the latent 2D heatmaps and latent depth maps, to ensure the scale and translation invariance. Absolute 3D hand poses are reconstructed from the latent 2.5D heatmaps.

3. Methods

We propose a framework to train a hand pose estimation model from RGB images by using depth images as privileged information. The model learns a new RGB representation which is influenced by the paired depth representa-

tion through mimicking the mid-level features of a depth network. As shown in Figure 1, we mainly use depth images in two ways: (1) to train an initial depth-based network with the aim of regressing 3D hand poses. Depth data that is annotated with 3D full hand pose information is abundant in the literature, and we choose the largest real dataset BigHand2.2M [43] to train our depth-based model, see the top row of Figure 1. (2) Paired RGB and depth images are fed into the RGB-based and depth-based network with the parameters of the depth-based network being frozen. The training of the RGB-based network continues with the aim of minimizing a joint loss function. The joint loss function has two terms, the first term being the 3D hand pose regression loss, L_C , and the second term the mid-level regression loss, L_{inter} .

3.1. Architecture

Figure 1 shows the training architecture. There are two base models, one is a depth-based network and the other one is an RGB-based network. We use deep convolutional neural networks (CNNs), which have been widely used in hand pose estimation and have proven useful in transferring information from one network to another [12]. Prior work [17] has been shown useful in combining RGB and depth images as a four-dimensional RGB-D input to a single CNN model to estimate 3D hand pose. In our architecture, we share information in the middle layers of our two CNN models. Each CNN model produces a 3D hand pose estimation result given its input.

For clarity, we denote the depth-based network *Depth_Net*, the RGB-based network *RGB_Net* when this is trained before privileged information is used. When privileged information is introduced in the training, we denote the RGB-based network *RGB_PI_Net*. *RGB_Net* and *RGB_PI_Net* are the same CNN model trained before and after the paired RGB and images are used to train the RGB channel. To let the *Depth_Net* channel share information with *RGB_PI_Net*, we introduce an intermediate regression loss between the paired layers in the two models. This intermediate regression loss is inspired by prior work [12, 13], where similar techniques are used for model distillation [12], supervision transfer from well labeled RGB images to depth images with limited annotation [11], and hallucination of different modalities [13]. We therefore introduce an intermediate loss, which helps *RGB_PI_Net* to extract middle level features that mimic the responses of the corresponding layer of the *Depth_Net* using the paired depth image. The intermediate loss L_{inter} (see Figure 1) is defined as:

$$L_{inter}(k) = \|A_k^{Depth} - A_k^{RGB}\|_2^2, \quad (1)$$

where A_k^{Depth} and A_k^{RGB} are the k th layer activations for the Depth Network and RGB Network, respectively. In our

experiments we select the final layer before the fully connected layers. During testing, when only an RGB image is available, we feed the RGB image into *RGB_PINet* to estimate the 3D hand pose.

3.2. Training with privileged information

This section explains the details of training the proposed architecture. We choose a base CNN for *DepthNet* and *RGB_PINet* for 3D hand pose estimation. For the base model, we build on Convolutional Pose Machine (CPM)’s [40] feature extraction layers with two fully connected layers to regress a 63 dimensional 3D hand pose with 21 joints. The *DepthNet* is initially trained on the Big-Hand2.2M [43] dataset, which has 2.2 million depth images with annotated 3D pose. After training, the model is further trained on the depth images of a smaller dataset (e.g., Stereo [44] and RHD [47] datasets) that has fully annotated paired RGB and depth images. The *RGBNet* is initially trained on the RGB images from the respective dataset.

When the initial training is completed for both CNN models, we freeze the parameters of the *DepthNet* and start training *RGB_PINet* with the paired depth images. During the privileged training stage, we want the *RGB_PINet*’s middle layer activations to match the activations of the corresponding layers of the *DepthNet*. We have two losses to optimize: (1) L_{inter} (Eqn. 1) to match the middle layer activations of the two CNN models. (2) L_C (see Figure 1) is the L_2 loss between the ground truth and the estimated 3D hand pose. We use a joint loss:

$$L_{joint}(k) = L_{inter}(k) + \lambda L_C, \quad (2)$$

where λ is used to balance the two losses, a larger value of λ means less supervision is required from the privileged information, a smaller value means that the model depends more on the supervision. We set λ to 100 for all experiments.

3.3. Foreground mask as privileged information

In addition to the supervision from depth images, we also explore the idea of extracting hand masks from depth images and embedding the hand masks into CNN layers of *RGB_PINet* to suppress the background features. We treat the hand mask M_h as privileged information. At test time, when the hand mask is not available, the CNN model is viewed as a standard CNN with convolutional layers, pooling layers and full-connected layers, where the L_{mask} is not used. In the training stage, the foreground hand mask is introduced in the last convolutional layer. Pixels of the mask M_h are zero on the hand region, and one otherwise. We suppress background features by minimizing the regression loss L_{mask} :

$$L_{mask} = \|A_k^{RGB} \odot M_h\|_2^2, \quad (3)$$

Dataset	No. Training	No. Test	No. Joints	Annotation	Type
Stereo [44]	15,000	3,000	21	2D, 3D	real
RHD [47]	41,258	2,728	21	2D, 3D	synthetic
Dexter-Object [29]	-	3,111	5 (tips)	2D, 3D	real

Table 1. Public datasets used in our experiments.

where \odot denotes element-wise multiplication. By minimizing the regression loss, where the response on the hand is multiplied by zero and the response outside the hand is multiplied by one, the response from outside the hand area is suppressed, focusing the response on the hand region.

4. Experiments

We carry out experiments on both 3D and 2D hand pose estimation from RGB images. Our experiments are conducted on three public RGB-D datasets: the RHD dataset [47], the Stereo dataset [44], and the Dexter-Object dataset [29], as shown in Table 1. The RHD dataset [47] is created synthetically and contains 41,258 training and 2,728 test images, with a resolution of 320×320 . The Stereo [44] dataset is a real RGB-D dataset, which has 18,000 pairs of RGB and depth images with a resolution of 640×480 pixels. The Dexter-Object (DO) [29] dataset contains 3,111 images of two subjects performing manipulations with a cuboid. The dataset provides RGB and depth images, but only fingertips are annotated. The RGB images have a resolution of 640×320 pixels. Due to the incomplete hand annotation, we use this dataset for cross-dataset generalization. During testing on a GTX 1080 Ti, the network forward steps take 6ms for 3D pose estimation and 8ms for the 2D case. The image cropping and normalization is the same as in [47]. To crop the hand region, we use ground truth annotations to obtain an axis-aligned crop, resized to 256×256 pixels by bilinear interpolation. Examples are shown in the first row of Figure 3. For 3D hand pose estimation, we use the root joint’s world coordinates and the hand scale to normalize the results.

4.1. 3D hand pose estimation from RGB

In this section, we investigate the usefulness of depth images to improve the performance of 3D hand pose estimation from an RGB image. Our base CNN model is built upon the feature extraction layers of Convolutional Pose Machine (CPM) [40] with two fully connected layers. The final output is a 63-d vector denoting the 21 joint 3D locations. Specifically, our base CNN model contains 14 convolutional layers, 4 pooling layers, and 2 fully-connected layers. At training stage, we have access to paired RGB and depth images. Initially the *DepthNet* is trained on *Big-Hand2.2M* [43]. We continue to train the *DepthNet* using the depth images from the small dataset, e.g., Stereo dataset or RHD dataset. We train the *RGBNet* with the RGB images from the small dataset. When the initial training is

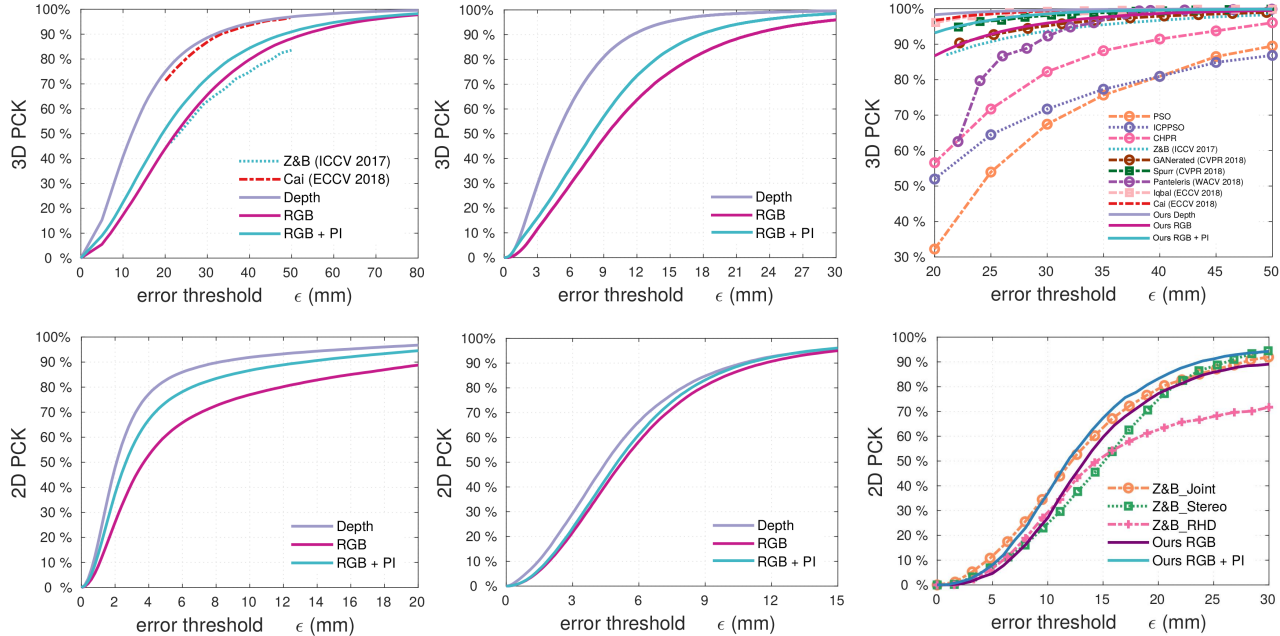


Figure 2. **Results on Stereo and RHD dataset for 3D hand pose and 2D pose accuracy.** Top row shows the comparisons of 3D hand pose accuracy, bottom row shows the comparisons of 2D hand pose accuracy. Top-left is self-comparison on RHD dataset, top-middle is self-comparison on Stereo dataset, top-right is comparison with state-of-the-art on Stereo dataset. Bottom-left is self comparison on RHD dataset, bottom-middle is self comparison on Stereo dataset, bottom-right is comparison with state-of-the-art on the Dexter-Object dataset.

completed, we start PI-training with the paired RGB and depth images. We freeze the weights of the *Depth_Net* and add the intermediate regression loss L_{inter} among the mid-level features of *Depth_Net* and *RGB_PI_Net*, then we continue the training of *RGB_PI_Net* by minimizing the joint loss L_{joint} . We apply the intermediate loss to the last convolutional layers of both branches, where the parameter k is set to 18 in Equation 1 and Equation 2.

Effect of PI-Learning: We conduct experiments with the two baseline CNNs and the CNN after PI training, see the accuracy curves in Figure 2 (top-left plot), and some examples in Figure 5. Our networks only estimate relative 3D pose from a cropped RGB image patch containing the hand, to yield 3D hand pose in world coordinates, we follow a similar procedure of [47], *i.e.*, by adding the absolute position of the root joint to our estimated results. For comparison we choose the percentage of correct keypoints (PCK) for varying threshold values. Training with depth data significantly improves the performance of the RGB-based network, narrowing the gap to the depth-based network.

Comparison with the state of the art: We compare our results with state-of-the-art methods, including PSO [19], ICPPSO [22], Zhang *et al.* [44], Z&B [47], GAN-erated [16], Cai *et al.* [2], Spurr *et al.* [28], Iqbal *et al.* [14], Panteleris *et al.* [21], see Figure 2 (top-right plot). We outperform (Z&B) [47] and [16]. While both [47] and [16] used extra training data, [47] used both Stereo (real) and RHD (synthetic) data to train their network. [16] used syn-

thetic (GAN-erated) data to train their network. The proposed method uses less RGB training data and achieved the best performance.

Feature activation maps: To give more intuition into training using additional privileged information, we visualize the activations of the mid-level feature maps for the three networks. Feeding an RGB image into each network, we aggregate all the mid-level feature maps into feature map by taking the maximum across all feature maps (similar to the maxout operation [10]). As shown in Figure 3, training with privileged information helps to select more representative features, where the visualized activations are close to the foreground.

Loss function evolution: We plot the loss during our training on the Stereo dataset, see Figure 4. The loss for 3D hand pose (left) of the RGB network on the test data converges at iteration 15,000, we continue training for another 5,000 iterations. From iteration 20,000, we fix the depth network parameters and connect mid-level features between the RGB and depth networks, and continue training by minimizing the joint loss (right) using RGB-D image pairs. The intermediate loss (middle) is used to suppress the difference between the mid-level feature between the RGB and depth networks. Loss for 3D hand pose of the RGB network, and the joint loss stop decreasing at around iteration 30,000.




Figure 3. **Feature activation maps.** (top row) input images, (row 2) activations of the RGB network trained on RGB only, (row 3) activations of the RGB network trained with additional depth data, (row 4) activations of the depth network.

4.2. 2D hand pose estimation from RGB

In this section, we choose the base CNN model as CPM [40], which has shown great performance for 2D human pose estimation [40], and 2D hand pose estimation [47]. Results are reported in Table 2, where ‘EPE’ is the average end point error in pixels, where an end point is a hand joint. Qualitative examples are shown in Figure 6. In this part of experiments, we treat the hand mask as privileged data, the CNN base model is CPM [27]. The baseline is obtained by the normal training procedure, *i.e.*, feeding the pre-processed hand image into CPM and obtaining the 2D hand pose by finding out the maximum location in each of the 21 heatmaps. For training with privileged information, we randomly select a certain proportion of RGB training data and use the hand masks, which are obtained by thresholding the depth images, in L_{mask} to suppress the background responses. As shown in Table 2, where 0.2 and 0.8 denotes the percentage of images when the L_{mask} is used during the training for 2D hand pose estimation.

Performance on hand-object interaction dataset: In Figure 2 (bottom-right plot), we show a comparison in terms of 2D PCK (in pixels) on the Dexter-Object(DO) [29] dataset. $Z\&B_Joint$ denotes the method of Z&B [47] trained on both RHD and Stereo datasets, which is better than $Z\&B_Stereo$ (trained on Stereo) and $Z\&B_RHD$ (trained on RHD). Our approach outperformed $Z\&B_Joint$ even though we used less RGB training data.

Method	Testing	Training	EPE median	EPE mean
Z&B [47]	RHD	RHD/Stereo	5.001	9.135
Baseline RGB	RHD	RHD	3.708	7.841
Baseline Depth	RHD	RHD	2.087	3.902
RGB + PI training	RHD	RHD	2.642	5.223
Z&B [47]	Stereo	RHD/Stereo	5.522	5.013
Baseline RGB	Stereo	Stereo	5.250	6.533
Baseline Depth	Stereo	Stereo	4.775	5.883
RGB + PI training (0.2)	Stereo	Stereo	5.068	6.280
RGB + PI training (0.8)	Stereo	Stereo	4.515	5.801
Z&B [47]	DO	RHD/Stereo	13.684	25.160
Baseline RGB	DO	RHD	13.360	18.278
RGB + PI training	DO	RHD	11.809	14.593

Table 2. **2D Hand Pose Accuracy.** Results when training on the RHD and Stereo datasets. The median and mean errors of RGB+PI training are consistently lower than the Baseline RGB error.

5. Conclusions

In this paper, we proposed a framework for 3D hand pose estimation from RGB images, with privileged information at the training stage in the form of depth data. We proposed three ways to use the privileged information: as training data for a depth-based network, as paired depth data to transfer supervision from the depth-based network to the RGB-based network, and to suppress background activations in the RGB-based network. This training strategy can be easily embedded into existing pose estimation methods. As an illustration, we showed how to trained a different CNN model to predict 2D hand pose from an RGB image. During testing, when only RGB images are available, our model significantly outperformed models trained on RGB images only.

References

- [1] J. Ba and R. Caruana. Do deep nets really need to be deep? In *NIPS*, 2014. 2
- [2] Y. Cai, L. Ge, J. Cai, and J. Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *ECCV, 2018*. 1, 3, 5
- [3] L. Chen, W. Li, and D. Xu. Recognizing rgb images by learning from rgb-d data. In *CVPR 2014*. 3
- [4] Y. Chen, X. Jin, J. Feng, and S. Yan. Training group orthogonal neural networks with privileged information. *arXiv:1701.06772*, 2017. 1, 2
- [5] C. Choi, A. Sinha, J. Hee Choi, S. Jang, and K. Ramani. A collaborative filtering approach to real-time hand pose estimation. In *ICCV*, 2015. 2
- [6] A. Erol, G. Bebis, M. Nicolescu, R. D. Boyle, and X. Twombly. Vision-based hand pose estimation: A review. *CVIU*, 108(1):52–73, 2007. 2
- [7] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person hand action benchmark with RGB-




Figure 4. Loss function evolution on the stereo dataset [44].

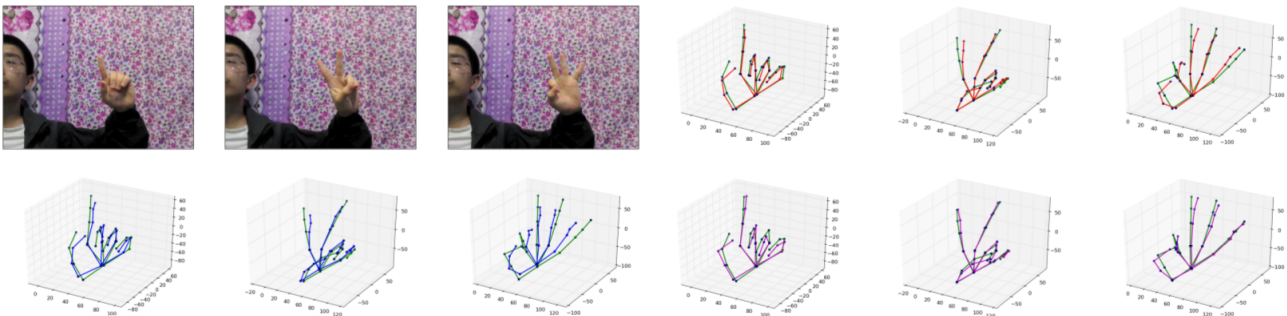


Figure 5. Qualitative 3D pose estimation results. Comparing the outputs of the RGB network (blue, left, second row), the RGB network with PI training (red, right, first row), and the Depth network (magenta, right, second row) with the ground truth 3D pose (green) on the Stereo dataset. Left, top row are the original images.

- D videos and 3D hand pose annotations. *CVPR*, 2018. 1, 2
- [8] L. Ge, H. Liang, J. Yuan, and D. Thalmann. 3D convolutional neural networks for efficient and robust hand pose estimation from single depth images. *CVPR*, 2017. 1, 2
- [9] L. Ge, Z. Ren, Y. Li, Z. Xue, Y. Wang, J. Cai, and J. Yuan. 3d hand shape and pose estimation from a single rgb image. In *CVPR*, 2019. 1
- [10] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. *ICML*, 2013. 5
- [11] S. Gupta, J. Hoffman, and J. Malik. Cross modal distillation for supervision transfer. In *CVPR 2016*. 3
- [12] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. In *arXiv:1503.02531*, 2015. 3
- [13] J. Hoffman, S. Gupta, and T. Darrell. Learning with side information through modality hallucination. In *CVPR*, 2016. 1, 2, 3
- [14] U. Iqbal, P. Molchanov, T. Breuel, J. Gall, and J. Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *ECCV, 2018*. 3, 5
- [15] Z. Luo, L. Jiang, J.-T. Hsieh, J. C. Niebles, and L. Fei-Fei. Graph distillation for action detection with privileged information. In *ECCV, 2018*. 3
- [16] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. GANerated hands for real-time 3D hand tracking from monocular RGB. In *CVPR, 2018*. 1, 2, 5
- [17] F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time hand tracking under occlusion from an egocentric RGB-D sensor. In *ICCV, 2017*. 1, 3
- [18] M. Oberweger, P. Wohlhart, and V. Lepetit. Training a feedback loop for hand pose estimation. In *ICCV*, 2015. 1, 2
- [19] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Efficient model-based 3D tracking of hand articulations using kinect. In *BMVC*, 2011. 5
- [20] P. Panteleris and A. Argyros. Back to RGB: 3D tracking of hands and hand-object interactions based on short-baseline stereo. In *ICCV Workshop*, 2017. 1, 2
- [21] P. Panteleris, I. Oikonomidis, and A. A. Argyros. Using a single rgb frame for real time 3d hand pose estimation in the wild. In *WACV 2018*. 5




Figure 6. **Qualitative 2D pose estimation results.** Comparing the outputs of (top) the RGB network and (bottom) the RGB network with PI training on the RHD dataset (left) and Stereo dataset (right).

- [22] C. Qian, X. Sun, Y. Wei, X. Tang, and J. Sun. Realtime and robust hand tracking from depth. In *CVPR*, 2014. 5
- [23] M. Rad, M. Oberweger, and V. Lepetit. Domain transfer for 3d pose estimation from color images without manual annotations. In *ACCV 2018*. 3
- [24] T. Sharp, C. Keskin, D. Robertson, J. Taylor, J. Shotton, D. Kim, C. Rhemann, I. Leichter, A. Vinnikov, Y. Wei, D. Freedman, P. Kohli, E. Krupka, A. Fitzgibbon, and S. Izadi. Accurate, robust, and flexible real-time hand tracking. In *CHI*, 2015. 2
- [25] Z. Shi and T.-K. Kim. Learning and refining of privileged information-based RNNs for action recognition from depth sequences. In *CVPR*, 2017. 1, 2
- [26] E. Simo-Serra, A. Ramisa, G. Alenyà, C. Torras, and F. Moreno-Noguer. Single image 3D human pose estimation from noisy observations. In *CVPR*, 2012. 2
- [27] T. Simon, H. Joo, I. Matthews, and Y. Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *CVPR*, 2017. 1, 2, 6
- [28] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal deep variational hand pose estimation. In *CVPR*, 2018. 3, 5
- [29] S. Sridhar, F. Mueller, M. Zollhöfer, D. Casas, A. Oulasvirta, and C. Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d input. In *ECCV*, pages 294–310. Springer, 2016. 2, 4, 6
- [30] X. Sun, Y. Wei, S. Liang, X. Tang, and J. Sun. Cascaded hand pose regression. In *CVPR*, 2015. 1, 2
- [31] J. S. Supancic, G. Rogez, Y. Yang, J. Shotton, and D. Ramanan. Depth-based hand pose estimation: methods, data, and challenges. *ICCV*, 2015. 2
- [32] D. Tang, H. J. Chang, A. Tejani, and T.-K. Kim. Latent regression forest: Structured estimation of 3D articulated hand posture. In *CVPR*, 2014. 1, 2
- [33] D. Tome, C. Russell, and L. Agapito. Lifting from the deep: Convolutional 3D pose estimation from a single image. *CVPR*, 2017. 2
- [34] J. Tompson, M. Stein, Y. Lecun, and K. Perlin. Real-time continuous pose recovery of human hands using convolutional networks. In *TOG*, 2014. 1, 2
- [35] A. Toshev and C. Szegedy. DeepPose: Human pose estimation via deep neural networks. In *CVPR*, 2014. 2
- [36] V. Vapnik and A. Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009. 1, 2
- [37] C. Wan, T. Probst, L. Van Gool, and A. Yao. Crossing nets: Combining GANs and VAEs with a shared latent space for hand pose estimation. In *CVPR*, 2017. 1, 2
- [38] C. Wan, A. Yao, and L. Van Gool. Hand pose estimation from local surface normals. In *ECCV*, 2016. 2
- [39] Z. Wang and Q. Ji. Classifier learning with hidden information. In *CVPR*, 2015. 1
- [40] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016. 2, 4, 6
- [41] Q. Ye, S. Yuan, and T.-K. Kim. Spatial attention deep net with partial PSO for hierarchical hybrid hand pose estimation. In *ECCV*, 2016. 1, 2
- [42] S. Yuan et al. 3D hand pose estimation: From current achievements to future goals. In *CVPR*, 2018. 2
- [43] S. Yuan, Q. Ye, B. Stenger, S. Jain, and T.-K. Kim. BigHand2.2M benchmark: Hand pose dataset and state of the art analysis. *CVPR*, 2017. 1, 2, 3, 4
- [44] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang. 3D hand pose tracking and estimation using stereo matching. *arXiv:1610.07214*, 2016. 1, 2, 4, 5, 7
- [45] X. Zhou, X. Sun, W. Zhang, S. Liang, and Y. Wei. Deep kinematic pose regression. In *ECCV*, 2016. 2
- [46] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3D human pose estimation from monocular video. In *CVPR*, 2016. 2
- [47] C. Zimmermann and T. Brox. Learning to estimate 3D hand pose from single RGB images. In *ICCV*, 2017. 1, 2, 4, 5, 6