Accepted Manuscript

Large-scale Multiview 3D Hand Pose Dataset

Francisco Gomez-Donoso, Sergio Orts-Escalano, Miguel Cazorla

PII: S0262-8856(18)30180-X

DOI: <https://doi.org/10.1016/j.imavis.2018.12.001>

Reference: IMAVIS 3730

To appear in: *Image and Vision Computing*

Received date: 20 July 2017

Revised date: 1 October 2018

Accepted date: 3 December 2018

Please cite this article as: Francisco Gomez-Donoso, Sergio Orts-Escalano, Miguel Cazorla , Large-scale Multiview 3D Hand Pose Dataset. Imavis (2018), <https://doi.org/10.1016/j.imavis.2018.12.001>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Large-scale Multiview 3D Hand Pose Dataset

Francisco Gomez-Donoso and Sergio Orts-Escolano and Miguel Cazorla

University Institute for Computer Research. University of Alicante

Abstract

Accurate visual hand pose estimation at joint level has several applications for human-robot interaction, natural user interfaces and virtual/augmented reality applications. However, it is still an open problem being addressed by the computer vision community. Recent novel deep learning techniques may help circumvent the limitations of standard approaches. However, they require large amounts of accurate annotated data.

Hand pose datasets that have been released so far present issues such as limited number of samples, inaccurate data or high-level annotations. Moreover, most of them are focused on depth-based approaches, providing only depth information (missing RGB data).

In this work, we present a novel multiview hand pose dataset in which we provide hand color images and different kind of annotations for each sample, i.e. the bounding box and the 2D and 3D location on the joints in the hand. Furthermore, we introduce a simple yet accurate deep learning architecture for real-time robust 2D hand pose estimation. Then, we conduct experiments that show how the use of the proposed dataset in the training stage produces accurate results for 2D hand pose estimation using a single color camera.

1. Introduction

Hand pose estimation in 2D images is a burgeoning research line that is booming in the field of computer vision. This problem is not only about the location of the hand itself, but also the localization of the hand joints, namely, the pose of the hand. The usefulness of this identification is wide ranging, from its use for sign language recognition [1, 2] to the identification of more complex behaviors such as hand gestures. It is also worth highlighting its importance in virtual reality and augmented reality applications. In order to interact in a natural way in these virtual worlds, the hand pose detection problem must first be tackled. Furthermore, to achieve a seamlessly integrated experience in virtual worlds, non-intrusive methods must be utilized for the hand pose detection. A way to accomplish this task is to use computer vision algorithms that work on images to retrieve the hand pose.

It is well known that deep learning systems require a large amount of labeled data to train robust models. In most cases, manual annotation is not feasible and researchers have come up with different procedures that can automatically generate ground truth annotations. Commonly, two different approaches are used for generating large amount of annotated data. First, there are tools for generating synthetic data. For example, in computer vision it is very common to use a graphics engine to simulate a real-world camera and so acquire synthetic images. However, this approach suffers from realistic-synthetic discrepancies and the trained network is therefore unable to correctly classify real-life data. The other approach is to use ad-hoc tools/sensors that can provide us with ground truth annotations. This is usually the approach that yields best results since it is able to automatize

the labeling procedure and the data that is acquired is more similar to real data.

As the review of the state of the art reveals, earliest hand pose datasets were composed of color images, but they were intended to perform static or dynamic gesture recognition, so they did not provide hand joint level annotations, but a class for each sample. Then, with the emergence of low cost range sensors, the datasets began to set aside the color information to provide depth maps instead. Despite most deep learning approaches utilize color information to tackle classification and regression problems, there also exist some approaches that take advantage of 3D information, such as depth maps or point clouds.

This paper presents a novel pose dataset. It provides images of hands from multiple points of view. The ground truth is composed of different kinds of data. For each image the hand location is provided: its bounding box, the X and Y image coordinates of each joint position of the hand, and the X, Y and Z real-world coordinates of the position of each joint of the hand. The dataset contains images of several individuals at different instants in time. There are even samples captured under special conditions which will challenge the generalization capabilities of the algorithms that use this dataset. Training and test splits generation methodology is also proposed.

The rest of the paper is organized as follows. Section 2 reviews state of the art works related to the existent hand pose datasets. Next, in Section 3 the details of the capture device we use for composing the dataset are given. Then, Section 4 explains relevant details of the dataset itself. Section 5

describes a system for hand pose regression trained on this dataset to be used by other hand pose estimation based proposals as a baseline. Section 6 presents the main conclusions of the present work and some lines of future work are suggested.

2. Related Works

Hand pose estimation is one of the most common areas of research in human-computer interaction. Since early in the 2000's, efforts have been made to achieve highly accurate hand gesture and pose recognition through different methods. First, traditional computer vision algorithms were applied in static color images to segment the hand from the background. Then, machine learning methods were used to tackle the problem and lastly, the most sophisticated and also accurate system have been achieved through deep learning. Alongside these methods, different kinds of datasets emerged, intended to feed these systems with proper data. In this section, we review these datasets, and highlight their weaknesses and strengths.

As previously mentioned, early in the 2000's a set of two datasets for static pose detection [3] and dynamic gesture classification [4] appeared. The static pose version dataset consists of over 3,000 images of hands in static poses distributed in 6 different classes. The dynamic gestures dataset is composed of 60 sequences of hands performing a certain gesture. This dataset is distributed in 4 classes.

In both versions, only the class of each sample is provided, with no further annotations, and the samples are highly unbalanced, which may harm the accuracy of the learning algorithms. The resolution and quantity of the

samples are also clearly insufficient by today's standards.

Much later, in 2014, Pisharady et al. [5] released another hand pose classification dataset. The NUS hand posture dataset consists of 10 classes of postures, 24 sample images per class, which are captured by varying the position and size of the hand within the image frame. Both greyscale and color images are available at 160×120 pixels. The hand postures are selected in such a way that the inter class variation in the appearance of the postures is minimized, which makes the recognition task challenging, but the background of the images in this dataset is uniform.

Again, this dataset does not provide annotation of the joints in the hand, but a class for each sample, so it is not intended to be used in continuous hand pose estimation at joint level. Moreover, they state the background is uniform which is an undesirable feature for learning algorithms.

The dataset for hand gesture recognition introduced in [6] contains gestures from Polish Sign Language and American Sign Language. In addition, some special signs were included. The database consists of three series of gestures which include the subsequent data: original RGB images at different resolutions, but with great image quality; ground truth binary skin presence masks and hand feature points location. In total, this dataset contains over 1600 samples.

Despite its high resolution and great quality ground truth, this dataset contains insufficient number of samples. There is a further and bigger drawback: the background across the samples is always the same. This might harm the generalization capabilities of the learning system, as in a real world scenario the background is constantly changing, and clearly not always grey

with a great deal of contrast between the hand and the background.

Another dataset was released by Molina et al. [7]. This dataset is composed of natural and synthetic depth maps. Each depth map represents a segmented hand in a certain pose. The annotations include the 2D position of each joint in the depth map coordinate frame, and the depth value in that point. They included 65 natural samples divided in 6 different dictionaries split by their occlusion level, pose based, motion based, or compound content. The synthetic data was generated taking a natural seed and performing up to 200 random rotations.

This is a high representative hand pose dataset but it has a major drawback: it only provides depth maps, which forces whoever wants to use it to acquire a 3D sensor. The number of samples is also insufficient for deep learning algorithms and the synthetic samples do not represent the essence of the real world.

Following the depth trend spurred by the emergence of low cost range sensors, the NYU Hand Pose dataset [8] was created. This dataset contains 8252 test-set and 72757 training-set frames of captured RGBD data with ground-truth hand pose information. For each frame, the RGBD data from 3 Kinects is provided: a frontal view and 2 side views. The training set contains samples from a single user only, whilst the test set contains samples from two users. A synthetic re-creation (rendering) of the hand pose is also provided for each view.

The main drawback of this dataset is that the RGB images provided are projections of the point clouds obtained by a 3D sensor. This causes the 2D color images to be rendered in low quality, with a lot of unknown regions

out of range of the 3D sensor. Thus, for using this dataset a 3D sensor is required as well.

Another depth-based dataset is proposed in [9]. This dataset contains sequences of 14 hand gestures performed in two ways: using one finger and the whole hand. Each gesture is performed between 1 and 10 times by 28 participants, resulting in 2800 sequences. Sequences are labelled following their gesture, the number of fingers used, the performer and the trial. Each frame of sequences contains a depth image, the coordinates of 22 joints both in the 2D depth image space and in the 3D world space creating a full hand skeleton. Once again, this dataset does not include the corresponding RGB images so a 3D sensor is required in order to use this dataset.

In [10], a dataset of hands from an egocentric perspective is proposed. This dataset is focused on hand segmentation and activity recognition. It provides 15000 color frames, the corresponding segmentation mask ground truth and semantic activity label. Despite its large number of samples and high quality, this dataset cannot be used for 2D or 3D fine grain hand pose estimation since it does not provide hand joints.

A different dataset is proposed in [11]. In this work, the authors propose a stereo dataset which provides 18000 stereo pairs of hand poses and their ground truth 3D joints and depth maps. Despite the quality of the annotations, this dataset lacks of variability. It only includes a hand in the scene which is limited to a counting with the hand motion, or random similar movements. In addition, it only features two different individuals.

Creating a 3D hand pose dataset is a challenging task, as authors remark in [12]. In this work, the authors propose a methodology to accurately label

hand poses from depth maps in a semi-automated fashion. First, the user records a sequence of hand poses and manually annotates some frames in the image plane. Then, the annotations are projected to 3D space and propagated to the remaining frames. A global optimization procedure is carried out in order to obtain accurate 3D ground truth. This methodology provides useful insight for a hand pose dataset creation. Nonetheless, the annotation of samples along the sequence, and the necessity of stating the z-index of each joint is a tedious process that would eventually provoke a lack of accuracy due to the user fatigue. In addition, this methodology assumes a depth map input, which is also a drawback.

Some of these datasets are used in [13] to perform tridimensional hand pose estimation from a single color image. This approach combines real and synthetic datasets with 2D and 3D annotations to create a robust 3D hand pose estimation pipeline.

In view of this review of the state of the art, we consider there is a need for a large scale hand pose dataset composed of color images and their corresponding 3D and 2D annotations.

3. Capture Setup

To generate accurate ground truth, a custom capture device was created, as shown in Figure 1. This device consists of an aluminum structure with three articulated arms that holds a total of four cameras. Three are Logitech C920 Pro cameras, and the last one, in a zenithal position, is a Microsoft LifeCam Studio. These cameras provide high quality images up to 1080p resolution, and although this model comes with auto focus function, this feature

was disabled as it would affect the intrinsic camera parameters estimation. As the cameras are calibrated, these parameters may not vary. The images were captured at 640×480 resolution.

The location of the cameras is as follows: one is located in the arm located in the middle arm, facing downwards; two more are located in the side arms with an orientation of 45 degrees; the last camera is located at the base of the structure. Finally, there is a Leap Motion Controller fixed to the aluminum frame, in front of the cameras.

Figure 1: We created a custom structure in order to generate the dataset. This device holds the color cameras and the Leap Motion Controller, and allows us to quickly generate large amounts of ground truth data.

The Leap Motion Controller is a device that accurately captures the pose of a person's hands. It provides sub-millimeter hand tracking at 200Hz. This device supplies, among other data, the 3D position of the joints of a hand. This information is obtained in order to automatically annotate the position

of the joints in the color images. The device is attached at the bottom of the structure, facing upwards. This configuration creates an overlapping space that is captured by both the cameras and the Leap Motion Controller.

This device can capture an enormous amount of accurate ground truth with no effort.

4. Dataset Description

The proposed dataset is intended to be used for 3D and 2D hand pose estimation and hand area location. The dataset is structured in various sequences and each sequence is composed of a set of frames. A frame is a collection of different kinds of ground truth data for a certain instant in time.

The ground truth data provided includes:

- A set of 3D points of the joints in the hand as provided by the Leap Motion, without further processing.
- Four sets of 2D points as the resultant projection of the 3D points to each color camera coordinate frame.
- Four sets of 3D points (hand joints), one for each color camera coordinates.
- Four bounding boxes computed from the projection of the 3D points to each camera coordinate frame.
- Four color images of a hand as captured for each color camera.
- A full alternative augmented set of frames (segmented hand poses and random backgrounds).

Finally, it is worth noting that each sequence was captured in different conditions to assure high variability. Subject, moment of the day and speed of motion are some of the parameters we considered to achieve high variability. Further details on generalization features are given in 4.6.

4.1. 3D Points of the Joints in a Hand

For each frame a set of 3D point is provided. These 3D points correspond to the joints in the hand as provided by the Leap Motion Controller with no further preprocessing. Table 1 shows a list of each point and its semantic correspondence. In Figure 3, the different points are marked on a real hand.

We provide a 3D point for each knuckle in a finger, plus the fingertips. Additionally, we also provide the palm position and the normal of the palm.

4.2. 3D and 2D Points of the Joints of a Hand

Besides the 3D points of the joints in a hand as provided by the Leap Motion, we also provide the 2D correspondences and the 3D space reprojection for each color image. In order to be able to automatically annotate images from a color camera using the Leap Motion Controller, we had to calibrate both devices, estimating the extrinsic parameters between their coordinate frames.

The calibration process requires two steps. First, estimation of intrinsic camera parameters, K , of the color camera and then estimation of extrinsic parameters between the Leap Motion controller and the color camera, $[R|t]$. In this work, we use [14] for computing the intrinsic camera parameters. Then, we estimate the transformation that allows us to transform 3D points in the Leap Motion coordinate frame to the camera coordinate frame.

Joint ID	Finger	Bone	Joint ID	Finger	Bone
1	Index	Distal	11	Pinky	Intermediate
2	Index	Metacarpal	12	Pinky	Proximal
3	Index	Intermediate	13	Ring	Distal
4	Index	Proximal	14	Ring	Metacarpal
5	Middle	Distal	15	Ring	Intermediate
6	Middle	Metacarpal	16	Ring	Proximal
7	Middle	Intermediate	17	Thumb	Distal
8	Middle	Proximal	18	Thumb	Metacarpal
9	Pinky	Distal	19	Thumb	Intermediate
10	Pinky	Metacarpal	20	Thumb	Proximal

Table 1: Joints of the hand present in the dataset. The Joint ID columns reference the numerical label shown in Figure 3.

To estimate the transformation between these two sensors, we use a set of correspondences between Leap Motion data and color camera data. Given these correspondences between the 3D Leap Motion controller points and their projection on the 2D images, we captured 10 color frames and manually labeled corresponding 3D joints positions of each fingertip in the color frame. We computed the transformation matrix between the two coordinate frames by solving the Perspective-n-Point problem. Figure 2 illustrates this process. Basically, given a set of n 3D points in the Leap Motion reference frame and their corresponding 2D image projections, as well as the intrinsic camera parameters, K , we determine the six degree-of-freedom (6DoF) pose of the camera in the form of its rotation R and translation t with respect to the Leap

Motion coordinate frame. We use the Levenberg-Marquardt optimization algorithm to minimize the final reprojection error. Figure 3 illustrates the result of the reprojected 3D points after the calibration procedure.

Figure 2: Calibration setup used for extrinsic parameter estimation by solving the Perspective-n-Point problem.

As stated before, we use this transformation matrix to project the 3D points provided by the Leap Motion Controller on each camera.

As the calibration parameters are unique for each camera, we applied this method to the four cameras of the capture device. We thus, generate four sets of 2D and 3D points for each frame, as a result of projecting the 3D points provided by the Leap Motion Controller to each color camera.

4.3. 2D Bounding Box of a Hand

Finally, we also provide the exact location of the hand in the scene for each sample through its bounding box. The bounding box is computed by taking the projection of the 3D points to the camera coordinate frame and then extracting the maximum and minimum values for the X and Y coordinates.

Figure 3: Once the devices are calibrated and the transformation matrix computed, new examples can be annotated in an automatic fashion in order to generate reliable ground truth.

Then, the width and length of the bounding boxes are enlarged by a dynamic offset, which depends on the depth. **This offset is $\frac{f \times p}{C}$. f is equal to the focal length in pixels. C is the depth of the hand. $p = 12.68$ which is the result of empirically setting 20 pixels of offset for a hand placed at $390mm$ from the sensor in our camera. This way, the bounding box offset increases as the hand gets closer to the sensor.**

Finally it is worth noting that the dataset provides 4 different bounding boxes, one per color image.

4.4. Color Images of a Hand

Our proposed dataset is intended to be used, but not exclusively, by deep learning models. These methods would take images of hands as input hopefully producing accurate pose estimations. Having said that, the quality of the images is a key factor in order to make these systems work.

It is well known that images captured by a sensor reflect a slightly deformed version of the real scenario. Some of the more notorious deformations are caused by the camera’s lens curvature (i.e. radial and tangential lens distortion). To deal with this problem, and to provide the best ground truth possible, an undistortioning process was carried out for each image as described in [15]. No further preprocessing was applied to the images.

As stated before, the dataset provides four color images for each frame. These four images picture a hand in the same time instant but from different perspectives.

4.5. Augmentated set

In addition to the dataset, we artificially created an alternative set which comprises some of the the hand poses with different random backgrounds.

To create this augmented subset, first we took the color frames and segmented the hand from the background from all the dataset. To do so, we used the HandSegNet[13] approach. This architecture is able to perform pixel-wise labeling and was trained to state if a pixel is part of a hand or not. We forwarded all the images through this network, thus achieving the hand segmentation.

As expected, the segmentation was not perfect as it tends to remove some parts of the fingers. To improve the segmentation of the fingers, we used the ground truth 2D positions of the joints to create a segmentation mask by creating links between them. The thickness of these links is a factor of the depth, with a manually chosen seed for each individual. By adding both segmentation masks, the one achieved with the HandSegNet, and the one achieved by linking the 2D joints, we finally get a robust segmentation of the

hand.

The next step is to isolate the hands and blend them on random backgrounds. The background images are composed of 20 different indoor landscapes extracted from a free stock photo database.

Finally, the augmented samples were manually reviewed in order to remove those sequences which do not provide a high quality segmentation.

This augmented subset is provided in order to improve the generalization capabilities of the algorithms trained on this dataset.

Figure 4 depicts some of these augmented samples.

Figure 4: Some samples extracted from the additional augmented subset.

4.6. Variability and Generalization Capabilities

As previously mentioned, this dataset for pose estimation was composed to provide accurate ground truth with high variability data. To ensure this, the creation of the dataset involved 9 different individuals recorded at different moments. The motion of the hand was spontaneous and unrestricted, with no plan at all. Therefore, a large variety of combinations are covered, in contrast to gesture-based dataset in which some fixed poses are considered. In addition, some sequences were carried out under special conditions: there are two sequences in which the subject wore different gloves, and a sequence

Figure 5: This plot depicts the data distribution. Note that there are no distinguishable clusters in the distribution, thus ensuring the high variability of the samples. Colors are randomly applied for easy visualization.

in which the subject wore a mask. Finally, it is worth to remark that all the subjects were right-handed.

In total, we provide over 20500 different frames distributed in 21 sequences. For each frame, 4 color images and 13 different annotations were provided, so over 82000 color images and over 266500 annotations are provided. Figure 6 shows some examples present in the dataset. Figure 5 depicts the data distribution. Each sample is composed by concatenating the x , y and z coordinates of its keypoints. Then, a t-SNE dimensionality reduction algorithm is applied. **Each point of the plot portrays a sample of the dataset. The emergence of clusters would mean that there are several poses that are covered whilst some other poses are not.** As show, there are no distinguishable clusters in the distribution, thus ensuring

the high variability of the samples.

To generate the training and test splits, we propose a leaving-one-user-methodology. To benchmark a system with this dataset, a leaving-one-user-out cross-validation is proposed. Each sequence features a user, but a user could appear in many sequences. Since the dataset is structured by sequences, the mapping between user and sequences is **as shown in Table 2**.

User	Sequences	User	Sequences
1	1, 2, 4 (mask), 7, 12 (white glove), 13 (back glove)	6	14, 15
2	3	7	16, 17, 18
3	5, 6	8	19
4	8, 9	9	20, 21
5	10, 11		

Table 2: **This table shows the performers of each sequence and its additional features.**

4.7. Ground Truth Evaluation

In order to provide a ground truth evaluation metric, 436 random samples were manually annotated by two different individuals. These manual annotations include 2D positions of each joint. Then we calculate the mean distance per joint by computing the euclidean distance for each manual and automatic keypoint annotations. Then, the distances were averaged, producing a mean distance error or 7.71 pixels, **which is 0.01204 error with**

Figure 6: Some images extracted from the proposed dataset. Notice the high variability of the samples.

respect to the image width. Figure 7 depicts percentage of keypoints which euclidean distance is below a certain threshold. Furthermore, more than 90% of the keypoints yield a distance below 15 pixels.

5. Baseline - 2D Joints Estimation

In this section, we present a baseline system for 2D hand pose estimation in color images using a deep learning approach. To do so, the following pipeline is proposed: first, a regular camera is used in order to capture a color image. This image is forwarded through a Faster R-CNN network [16] trained to detect the hands present in the scene. This network outputs both the bounding box of a hand and a score that represents the detection confidence. If this score is above 0.85, the estimated bounding box is used to extract the hand from the image. The crop of the hand is resized to

Figure 7: The leftmost plot depicts the percentage of keypoints which euclidean distance is below a certain threshold. Note that more than 90% of the automatically annotated keypoints yield a distance below 15 pixels with respect to the manually annotated keypoints. The rightmost images shows random images with their automatic annotations (green) and the manual annotations (blue).

224 $\times$ 224 pixels and fed to a custom ResNet50 [17] network. By using this network in regression mode we are able to estimate X and Y coordinates in the image plane for each joint position. In Section 5.2, the process we followed to perform regression on a ResNet50 architecture is described. The overall pipeline is shown in Figure 8.

Figure 8: Overall pipeline of the proposed method. First, an image is captured using a regular web camera. This image is forwarded to a hand detector in order to obtain the bounding box of the hand. Then, the hand is cropped and passed to the hand pose regressor to finally estimate the position of the hand joints.

Notice that there are two critical stages in this pipeline: the hand detection (bounding box estimation) and the hand pose regression stage. These subsystems are detailed in the following subsections.

Finally, it is worth noting that this pipeline provides highly accurate hand pose estimation while maintaining a reasonable computation cost. These features allow the system to run under real-time constraints. Further results and experimentation are detailed in Section 5.3.

5.1. Detecting Hand Location

The proposed hand detection module is based on the work originally presented in [16]. We adopted the VGG-16 network [18] pre-trained on the ImageNet dataset [19] as the backbone network for feature extraction. In addition, we modified the number of outputs in final layers, since we are only interested in detecting hands.

Figure 9: The Faster R-CNN network is able to accurately estimate the bounding box of a certain object, in our case, the network was trained for detecting and localizing hands.

We trained a Faster R-CNN network using different hands datasets. First, we used the dataset presented in [20]. This dataset contains hand images collected from various different public image dataset sources as listed in [20]. We extended the number of training images (1000) creating two additional

sources of data. The first additional set of images was taken from the proposed dataset. The second subset of images was created using Google Images. We used the keywords 'hands' and 'close-up' to download images containing close-up hands. Finally, we also manually annotated these additional images.

The main reason to train this system on different datasets was to avoid overfitting to a specific environment that would produce lack of generalization.

Figure 9 shows estimated bounding boxes over some test images once the Faster R-CNN network was trained.

5.2. *Hand Pose Regression*

The hand pose is computed by a modified ResNet50. The ResNet50 deep learning architecture is currently the state of the art convolutional neural network on image recognition, achieving a top-1 error (22.85) on the ImageNet validation split. This architecture introduces the 'residual' term, which consists of the aggregation of the previous block feature map to the output feature map of a convolution block. **The resultant feature map yields higher values where the response for the filters are high, in contrast with the positions where there are no response. These positions yield much lower values rather than no activation at all as a result of the residual aggregation.** However, the output of a convolution layer in a non residual-based convolutional neural network is only the result of the neuron activation. If a neuron is not triggered on a certain region of the input image, the output remains with lower (**near zero**) activation values. When the network computes the weights update in the backpropagation stage, the values on non-activated regions lead to very low upgrades, eventually even

resulting in no upgrade at all, which causes the learning to stall. This issue is known as the vanishing gradient problem. The inclusion of the 'residual' term helps fight the vanishing gradient problem and allows the design of even deeper architectures. Currently, the best performer on several tasks of the ImageNet challenge is based on the 'residual' approach introduced by ResNet.

As stated before, we used the ResNet50 version (50 layers). Our ResNet50 incarnation takes as input 224×224 pixels color images. The number of neurons from the last fully connected layer was adjusted to fit our problem. As there are 20 joints to locate composed of X and Y values for each joint, the number of output neurons was set to 40. We also modified the default softmax activation function was changed for a sigmoid one. By introducing these modifications, we are able to perform the regression of the keypoints in presence of an input hand image.

This architecture was trained from scratch using the training split of the proposed dataset. First, hands were extracted using the bounding boxes, and the joint annotations were remapped accordingly and normalized in a range between 0 and 1. The network was trained on the resultant dataset for 200 iterations reaching a training loss of 0.0000607571 and a test loss of 0.00031547. The test loss was used as an early stop criteria. The optimizer of choice was Adam with a learning rate of 0.001. The loss function we used was mean squared error.

5.3. Experimentation

In this section, we first describe the experimentation setup and then present results obtained for the proposed pipeline.

All timings and results for this pipeline were obtained by conducting the experiments in the following test setup: Intel Core i5-3570 with 8 GiB of Kingston HyperX 1600 MHz and CL10 DDR3 RAM on an Asus P8H77-M PRO motherboard (Intel H77 chipset). Principal storage was provided by a Seagate Desktop HDD. Additionally, the system included a NVIDIA GTX1080 GPU used for training and inference.

The execution of the full pipeline runs at 23 frames per second, but it is affordable to run at 30 – 35 frames per second with minor tweaks on the code. For example, the hand detector takes 15 ms per frame on a NVIDIA GTX 1080. In the current system the hand detector is running for every frame. However, with a naive motion estimation technique, bounding boxes could be tracked along the sequence and the hand position would only be recomputed when the system detects large displacements, thus reducing the computation time.

The frameworks of choice were Caffe 0.14-RC3 and Keras 1.2.0 with Tensor Flow 0.12 as the backbone, running on Ubuntu 14.04.02. Both were compiled using CMake 2.8.7, g++ 4.8.2, CUDA 8.0, and cuDNN v5.0.

As described earlier, there are two trainable subsystem in our pipeline. Both the hand detector and the hand pose estimator were trained using our dataset.

In order to validate our baseline approach, the full pipeline was processed by performing a leaving-one-user-out validation. In this case, User 2 was left for test. We achieved a validation mean distance error of 11.79 pixels per keypoint over the input hand image (224×224), **which is 0.01842 with respect to the width of the full input image**. Figure 10 shows the

Figure 10: This plot shows the percentage of keypoints which euclidean distance error is below a certain threshold, in pixels.

percentage of keypoints below a distance error threshold.

Figure 11: Our proposed system is able to accurately estimate hand joints position on a variety of hand poses. Nonetheless, it presents some failure cases as shown in the last image.

Figure 11 shows system performance for several hand poses during a live

test. These frames do not belong to the dataset. The last image shows a failure case due to the incorrect hand detection. As this approach is a dependent two-stage system, if the first step fails to detect the whole hand, the second system would try to regress the joints on wrong data, so it is very likely to fail.

6. Conclusions

Hand detection in both, 2D and 3D, is a challenging problem. In this work, we present a novel multiview hand pose dataset. It features over 20,500 annotated color images of hands. The annotations include, besides several images taken from different perspectives, the bounding box and the 2D and 3D position of each joint in the hand. Furthermore, we introduce a simple yet accurate architecture for real-time 2D hand pose estimation to be used as a baseline. This dataset could serve as a baseline for future works on hand detection. The dataset is available for download in ¹.

Following on this work we plan to design a new architecture for robust 3D hand pose estimation in which the system will provide not only x and y coordinates in the image plane but x, y and z, and the normal of the palm on real world coordinates.

Acknowledgements

This work has been funded by the Spanish Government TIN2016-76515-R grant for the COMBAHO project, supported with Feder funds. It has also

¹<http://www.rovit.ua.es/dataset/mhpdataset/>

been supported by the Generalitat Valenciana project GV/2018/22 and by a Spanish grant for PhD studies ACIF/2017/243. Experiments were made possible by a generous hardware donation from NVIDIA.

References

- [1] F. Gomez-Donoso, M. Cazorla, A. Garcia-Garcia, and J. Garcia-Rodriguez, “Automatic schaeffer’s gestures recognition system,” *Expert Systems*, vol. 33, no. 5, pp. 480–488, 2016.
- [2] M. Cazorla, F. Gomez-Donoso, D. Viejo, and J. Martinez-Gomez, “Recognizing schaeffer’s gestures for robot interaction,” vol. 9422 of *Lecture Notes in Computer Science*, Springer, 2015.
- [3] S. Marcel and O. Bernier, “Hand posture recognition in a body-face centered space,” in *Proceedings of the International Gesture Workshop on Gesture-Based Communication in Human-Computer Interaction*, GW ’99, (London, UK, UK), pp. 97–100, Springer-Verlag, 1999.
- [4] S. Marcel, O. Bernier, J. E. Viallet, and D. Collobert, “Hand gesture recognition using input-output hidden markov models,” in *Proceedings Fourth IEEE International Conference on Automatic Face and Gesture Recognition (Cat. No. PR00580)*, pp. 456–461, 2000.
- [5] P. K. Pisharady, P. Vadakkepat, and L. A. Poh, *Hand Posture and Face Recognition Using Fuzzy-Rough Approach*, pp. 63–80. Singapore: Springer Singapore, 2014.

- [6] T. Grzejszczak, M. Kawulok, and A. Galuszka, “Hand landmarks detection and localization in color images,” *Multimedia Tools and Applications*, vol. 75, no. 23, pp. 16363–16387, 2016.
- [7] J. Molina, J. A. Pajuelo, M. Escudero-Viñolo, J. Bescós, and J. M. Martínez, “A natural and synthetic corpus for benchmarking of hand gesture recognition systems,” *Machine Vision and Applications*, vol. 25, no. 4, pp. 943–954, 2014.
- [8] J. Tompson, M. Stein, Y. Lecun, and K. Perlin, “Real-time continuous pose recovery of human hands using convolutional networks,” *ACM Transactions on Graphics*, vol. 33, August 2014.
- [9] Q. D. Smedt, H. Wannous, J.-P. Vandeborre, J. Guerry, B. L. Saux, and D. Filliat, “3D Hand Gesture Recognition Using a Depth and Skeletal Dataset,” in *Eurographics Workshop on 3D Object Retrieval* (I. Pratikakis, F. Dupont, and M. Ovsjanikov, eds.), The Eurographics Association, 2017.
- [10] S. Bambach, S. Lee, D. J. Crandall, and C. Yu, “Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions,” *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 1949–1957, 2015.
- [11] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang, “3D Hand Pose Tracking and Estimation Using Stereo Matching,” *ArXiv e-prints*, Oct. 2016.

- [12] M. Oberweger, G. Riegler, P. Wohlhart, and V. Lepetit, “Efficiently Creating 3D Training Data for Fine Hand Pose Estimation,” *ArXiv e-prints*, May 2016.
- [13] C. Zimmermann and T. Brox, “Learning to estimate 3d hand pose from single rgb images,” in *IEEE International Conference on Computer Vision (ICCV)*, 2017. <https://arxiv.org/abs/1705.01389>.
- [14] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, pp. 1330–1334, Nov. 2000.
- [15] J. Weng, P. Cohen, and M. Herniou, “Camera calibration with distortion models and accuracy evaluation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 965–980, Oct 1992.
- [16] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems 28* (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, eds.), pp. 91–99, Curran Associates, Inc., 2015.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [18] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2014.
- [19] 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)*, vol. 115, no. 3, pp. 211–252, 2015.
- [20] A. Mittal, A. Zisserman, and P. H. S. Torr, “Hand detection using multiple proposals,” in *British Machine Vision Conference*, 2011.

Highlights

- A large-scale hand pose dataset with more than 266500 annotations
- Color frames from 4 different points of view are provided
- Annotations include 3D joint positions, 2D joint positions and bounding boxes
- A baseline which can be used to benchmark other approaches