

# Disentangling Pose from Appearance in Monochrome Hand Images

Yikang Li<sup>1\*</sup>, Chris Twigg<sup>2</sup>, Yuting Ye<sup>2</sup>, Lingling Tao<sup>2</sup>, Xiaogang Wang<sup>1</sup>

{ykli, xgwang}@ee.cuhk.edu.hk {chris.twigg, yuting.ye, lingling.tao}@oculus.com

<sup>1</sup>The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>2</sup>Facebook Reality Labs, USA

## Abstract

Hand pose estimation from monocular 2D image is challenging due to the variation in lighting, appearance, and background. While some success has been achieved using deep neural networks, they typically require collecting a large dataset that adequately samples all the axes of variation of hand images. It would therefore be useful to find a representation of hand pose which is independent of the image appearance (like hand texture, lighting, background), so that we can synthesize unseen images by mixing pose-appearance combinations. In this paper, we present a novel technique that disentangles the representation of pose from a complementary appearance factor in 2D monochrome images. We supervise this disentanglement process using a network that learns to generate images of hand using specified pose+appearance features. Unlike previous work, we do not require image pairs with a matching pose; instead, we use the pose annotations already available and introduce a novel use of cycle consistency to ensure orthogonality between the factors. Experimental results show that our self-disentanglement scheme successfully decomposes the hand image into pose and its complementary appearance features of comparable quality as the method using paired data. Additionally, training the model with extra synthesized images with unseen hand-appearance combinations by re-mixing pose and appearance factors from different images can improve the 2D pose estimation performance.

## 1. Introduction

Hand pose estimation is an important topic in computer vision with many practical applications including virtual/augmented reality (AR/VR) [22, 41] and human-computer interaction [46, 31]. A large body of work has shown robust hand pose estimation using RGB-D cameras [11, 47, 17, 50] or stereo cameras [56, 48] that provides 3D information about hands. With the recent advance of deep learning techniques [19, 27, 25, 24, 26], researchers

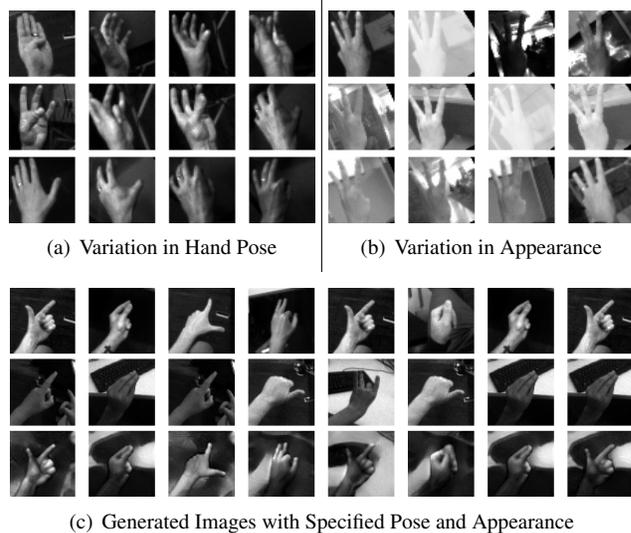


Figure 1. Robust hand pose detection requires handling the large variation in (a) hand pose, and, (b) Image appearance, e.g. different backgrounds, lighting conditions, camera exposures, hand textures, etc. (c) Image generated with pose from the first row and appearance from the second row.

have begun exploring the use of monocular 2D cameras [59, 34], which are cheap and ubiquitous thanks to their use in consumer devices like smart-phones and laptops.

Despite recent success of applying deep learning in hand pose estimation from monocular 2D images, there is still a substantial quality gap when comparing with depth-based approaches. We believe that the culprit is the variability in hand appearance caused by differences in lighting, backgrounds, and skin tones or textures. The same hand pose can appear quite differently in daylight than fluorescent lighting, and both harsh shadows and cluttered backgrounds tend to confuse neural networks. To ensure the robustness of neural networks, large amount of training data is typically required in order to adequately samples all the axes of variation.

In this work, we aim to improve the robustness of hand pose estimation from monocular 2D images by finding a representation of hand pose that is independent of its ap-

\*Work is done during an internship at Facebook Reality Labs.

pearance. We propose to train a neural network that learns to “disentangle” a hand image into two sets of features: the first captures the hand pose, while the second captures the hand’s appearance. Pose features refer to the informative factors used to reconstruct the actual hand pose (e.g. the locations of the hand joints), while the appearance feature denote the complementary “inessential” factors of the image, such as the background, lighting conditions, hand textures, *etc.* We refer to this decomposition as Factor Disentanglement.

Existing approaches to factor disentanglement generally require pairs of annotated data [33, 9], where the pairs share some features (e.g. object class) but vary in others (e.g. background). These pairs supervise the disentanglement process, demonstrating how different parts of the feature vector contribute to the composition of the image. While it is relatively easy to find multiple images that share the same object class, finding pairs of images with identical hand but different appearance is considerably more challenging. Instead, we propose to learn to disentangle the images using the supervision we do have: labeled poses for each training image.

To do this, we start with the following principles:

1. We should be able to predict the (labeled) hand pose using only the pose feature.
2. We should be able to reconstruct the original image using the combination of pose plus appearance features.

However, this is not sufficient, because there is nothing that ensures that the pose and appearance features are *orthogonal*, that is, they represent different factors of the image. The network can simply make the “appearance” feature effectively encode both pose and appearance, and reconstruct the image while ignoring the separate pose component. Therefore:

3. We should be able to combine the pose feature from one image with the appearance feature from another to get a novel image whose pose matches the first image and whose appearance matches the second.

Because we have no way to know *a priori* what this novel image should look like, we can not supervise it with an image reconstruction loss. Instead, we use *cycle consistency* [58]: if we disentangle this novel image, it should decompose back into the original pose and appearance features. This will ensure that the network does not learn to encode the pose into the appearance feature. We apply these three principles during our training process shown in Fig. 2.

The proposed self-disentanglement framework is applied on a dataset of monochrome hand images. We show that learning to disentangle hand pose and appearance features greatly improves the performance of hand pose estimation module in two ways: 1. the pose estimation module can

learn a better pose feature representation when the factor disentanglement is learned jointly as an auxiliary task. 2. the dataset can be augmented by generating new images with different pose-appearance combinations during the process. Both methods lead to improvement over baseline on our hand pose dataset. In addition, we also show comparable results to a factor disentanglement network trained with the supervision of paired images. Due to the challenge of capturing perfect paired data, we resort to a synthetic dataset for this comparison, where a pair of hand images are rendered using path tracing [42] with identical hand pose but different background, lighting, and hand textures. Although our experiments are done using monochrome images, this framework can be easily extended for the case of RGB images.

The main contribution of this paper is as follows:

- A self-distanglement network to disentangle the hand pose features from its complementary appearance features in monochrome images, without paired training data of identical poses.
- The proposed framework improves the robustness of supervised pose estimation to appearance variations, without use of additional data or labels.

## 2. Related Work

**Hand Tracking:** Due to the large pose space, occlusions, and appearance variation such as lighting and skin tone, hand tracking from images is a challenging problem. Methods that use multiple views [56, 38, 4] can deal with occlusions but are more difficult to set up. A large body of work [11, 47, 35, 50, 50] has demonstrated high-quality tracking use depth/RGBD cameras. While powerful, these sensors are still not largely available in consumer products. More recently, there has been work on tracking hands from monocular RGB camera [59, 34] using deep learning techniques. In this work, we will focus on monochrome cameras due to their increased sensitivity to low light, but the methods should generalize to RGB camera case as well.

**Encoder-Decoder Structure and Autoencoder:** Our base architecture is built on two encoder-decoder structures (also termed as contracting-expanding structures) [37, 1], which are neural networks with an encoder (contracting part) and a decoder (expanding part). The encoder is designed to extract a feature representation of the input (image) and the decoder translates the feature back to the input (autoencoder [7, 14]) or the desired output space. This structure is widely used in image restoration [6, 30, 55], image transformation [13, 9], pose estimation [36] or semantic segmentation [37, 1]. In addition, it has also been utilized for unsupervised feature learning [39] and factor disentangling [33, 49, 15]. In our framework, this architecture is adopted for both the image reconstruction module and the

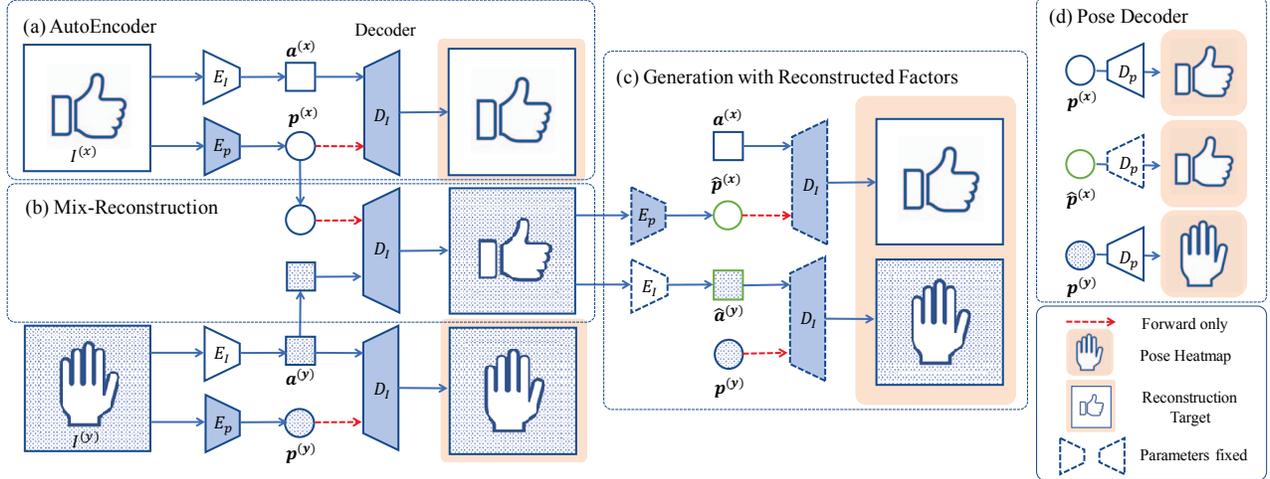


Figure 2. The overview of the self-disentanglement training process. (a) Input images  $I^{(x)}$  and  $I^{(y)}$  are encoded into Pose and Appearance factors, which contains the hand joint locations and its complementary image appearance information (e.g. background, lighting, hand texture, etc.) respectively.  $E_p$  and  $E_I$  are encoders for the pose and appearance factors respectively. Image decoder  $D_I$  is used to reconstruct the original images using the pose and appearance factors. (b) We combine the pose factor from  $I^{(x)}$  and appearance factor from  $I^{(y)}$  to construct a “mix-reconstructed” hand image with expected pose and appearance. (c) The mix-reconstructed image is decomposed back to the pose and appearance factors, and the resulting pose (appearance) feature is combined with the original appearance (pose) feature to generate a new decoded image, which should be similar to the original image. (d) The pose factors are trained to predict the pose heatmap with the pose decoder  $D_p$ . The dashed arrow indicates that we don’t allow the gradients from the image reconstruction loss to back-propagate through the pose factors. The dashed-outlined modules mean they just work as an estimator to provide the gradients to the early stage.

hand joint localization module, while we propose a novel unsupervised training method to ensure the separation of factors generated by the encoders.

**Learning Disentangled Representations.** Disentangling the factors of variation is a desirable property of learned representations [5], which has been investigated for a long time [51, 52, 13, 15]. In [13], an autoencoder is trained to separate a translation invariant representation from a code that is used to recover the translation information. In [43], the learned disentangled representations is applied to the task of emotion recognition. Mathieu *et al.* combine a Variational Autoencoder (VAE) with a GAN to disentangle representations depending on what is specified (i.e. labeled in the dataset) and the remaining unspecified factors of variation [33].

Recently, factor disentanglement has also been used to improve visual quality of synthesized/reconstructed images and/or to improve recognition accuracy for research problems such as pose-invariant face recognition [54, 40], identity-preserving image editing [16, 21, 23], and hand/body pose estimation [29, 3]. However, these factor disentanglement methods usually either require paired data or explicit attribute supervision to encode the expected attribute. Two recent techniques,  $\beta$ -VAE [12] and DIP-VAE [20], build on variational autoencoders (VAEs) to disentangle interpretable factors in an unsupervised way. However, they learn it by matching to an isotropic Gaussian prior, while our method learns disentanglement using a

novel cycle-consistency loss. [2] improves the robustness of pose estimation methods by synthesizing more images from the augmented skeletons, which is achieved by obtaining more unseen skeletons instead of leveraging the unseen combinations of the specified factor (pose) and unspecified factors (background) in the existing dataset like ours. The most related work is [57], which proposes an disentangled VAE to learn the specified (pose) and additional (appearance) factors. However, our method explicitly makes the appearance factor orthogonal to the pose during training process, while [2] only guarantees that the pose factor does not contain information about the image contents.

### 3. Learning Self-Disentanglement

In this section, we present our self-disentanglement framework. An overview of the framework can be found in Fig. 2. Our framework encodes a monochrome hand image into two orthogonal latent features: the pose feature  $\mathbf{p}$  and the appearance feature  $\mathbf{a}$  using pose encoder  $E_p$  and appearance encoder  $E_I$ . Without explicit supervision on how these two features disentangle, we introduce the following consistency criteria for self supervision.

#### 3.1. Pose Estimation Loss

To encode the pose feature, we use a model similar to the contracting-expanding structure of UNet [44]. As shown in the top of Fig. 3, we use down-sampling layers (the pose encoder  $E_p$ ) to encode the image  $I^{(x)}$  into a latent pose fea-



image as  $\tilde{I}^{(x)}$  and  $\tilde{I}^{(y)}$ . Now we build a disentangle-mix-disentangle-reconstruct cycle to generate back to the original input (denoted as the self-disentanglement), and we use the following cycle consistency losses during training:

$$\begin{aligned} L_{cycle-img}^{(x)} &= \mathbb{E} \left[ \left\| \tilde{I}^{(x)} - I^{(x)} \right\|_1 \right] \\ L_{cycle-img}^{(y)} &= \mathbb{E} \left[ \left\| \tilde{I}^{(y)} - I^{(y)} \right\|_1 \right] \end{aligned} \quad (4)$$

The reconstructed pose factors  $\hat{\mathbf{p}}^{(x)}$  and  $\hat{\mathbf{a}}^{(y)}$  should also match the  $\mathbf{p}^{(x)}$  and  $\mathbf{a}^{(y)}$ . An additional dual feature loss is also added as an auxiliary supervision to enforce the feature-level consistency:

$$\begin{aligned} L_{dual-pose}^{(x)} &= \mathbb{E} \left[ \left\| \hat{\mathbf{p}}^{(x)} - \mathbf{p}^{(x)} \right\|_1 \right] \\ L_{dual-img}^{(y)} &= \mathbb{E} \left[ \left\| \hat{\mathbf{a}}^{(y)} - \mathbf{a}^{(y)} \right\|_1 \right] \end{aligned} \quad (5)$$

where  $\mathbf{p}^{(x)}$  and  $\mathbf{a}^{(y)}$  here only serve as fixed training targets, and the gradients are not back-propagated through to  $E_p$  and  $E_I$ .

In addition, the mix-reconstructed image  $\hat{I}^{(xy)}$  is also expected to output the pose from  $I^x$ . Therefore, as shown in Fig. 2 (d) the reconstructed pose code  $\hat{\mathbf{p}}^{(x)}$  is decoded with the pose decoder  $D_p$  to the hand joint heatmap  $\tilde{P}^{(x)}$ , which should match the original heatmap  $P^{(x)}$ :

$$L_{cycle-pose}^{(x)} = \mathbb{E} \left[ \left\| \tilde{P}^{(x)} - P^{(x)} \right\|_1 \right] \quad (6)$$

### 3.5. End-to-end Training

The model is trained end-to-end with randomly sampled pairs  $I^{(x)}$  and  $I^{(y)}$ :

$$\begin{aligned} Loss &= L_p + \alpha_1 L_I + \alpha_2 L_D + \alpha_3 L_G \\ &\quad + \alpha_4 L_{cycle-img} + \alpha_7 L_{cycle-pose} \\ &\quad + \alpha_6 L_{dual-img} + \alpha_5 L_{dual-pose} \end{aligned} \quad (7)$$

where  $L_{(\cdot)}$  denotes the sum of the corresponding losses for the pairs  $L_{(\cdot)}^{(x)}$  and  $L_{(\cdot)}^{(y)}$ .

When evaluating cycle consistency, the pose decoder  $D_p$  and the image decoder  $D_I$  serve as an evaluator to estimate whether the mix-reconstructed image  $\hat{I}^{(xy)}$  can output the correct hand joint heatmap and can be encoded into expected features. We don't necessarily want to train them based on the mix-reconstructed image because it may be poor in quality, especially during the early stage of training. Therefore we fix the parameters in these two decoders in Fig. 2 (c-d) (shown with dash outline). They are simply a copy of the modules in Sec. 3.1 and Sec. 3.2, but do not accumulate gradients in back-propagation. This simple strategy greatly stabilizes training.

Dataset	Train (#frames)	Testing (#frames)
Real	123,034	13,416
Synthetic	$123,034 \times 2$	$13,416 \times 2$

Table 1. Statistics of the real and synthetic hand image datasets. The synthetic dataset is made up of pairs of images, which share the pose but differ in backgrounds, lighting conditions and the hand textures.

## 4. Experiments

### 4.1. Data Preparation

We collect a dataset of monochrome hand images captured by headset-mounted monochrome cameras in a variety of environments and lighting conditions. To obtain high quality ground truth labels of 3D joint locations, we rigidly attach the monochrome camera to a depth camera, then apply [53] on the depth image to obtain 3D key points on the hand. With careful camera calibration, we transform the 3D key points to the monochrome camera space as ground truth labels. The training images are then generated as a 64x64 crop around the hand.

In addition, we render a synthetic dataset of hand images using the same hand poses from the monochrome dataset. Each pose is rendered into a pair of images with different environment maps, lighting parameters and hand textures. This synthetic dataset offers perfectly paired data of the same pose with different appearances. Tab 1 shows statistics of the two datasets.

### 4.2. Implementation Details

We use an encoder-decoder architecture following UNet [44] without skip-connections as the base model. The encoder is shared between the pose feature  $E_I$  and the appearance feature  $E_p$  before the last downsampling layer. Two different decoders are used in the Hand Joint Localization branch and the Image Reconstruction branch respectively, where the image reconstruction branch decodes from both the pose feature and the appearance feature (Fig 3).

Both the encoder and the two decoders have a depth of 4. Each block in the encoder consists of repeated application of two 3x3 convolutions, each followed by a rectified linear unit (ReLU) and a 2x2 max pooling operation with stride 2 for down-sampling. The pose decoder  $D_p$  employs a 2x2 deconvolution layer in each block for up-sampling, while the image decoder  $D_I$  uses nearest-neighbor upsampling followed by a 3x3 convolution instead to avoid the checkerboard artifact [8]. Fig. 3 illustrates the detailed model structure. At training time, we initialize all parameters randomly, and use the Adam [18] optimizer with a fixed learning rate 0.001. A total of 75 epochs is run with a batch size of 128.

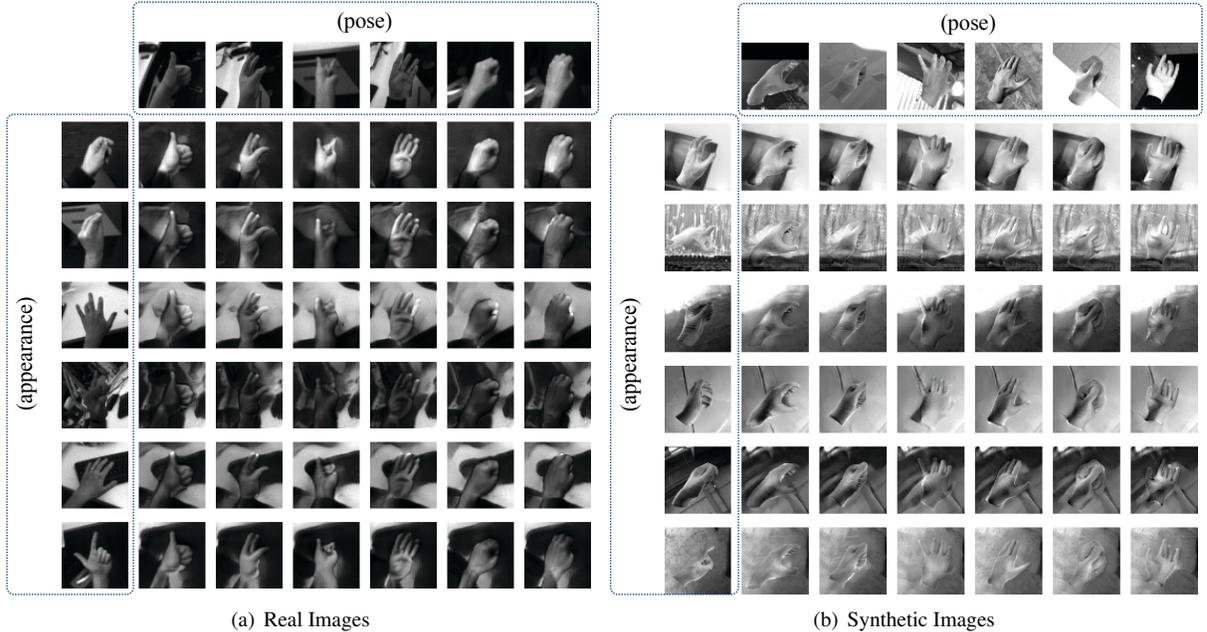


Figure 4. Self-disentanglement on Real (left) and Synthetic (right) data. The image on the top row provide the “pose” while the image on the left-most column offers the “appearance”. The images in the middle matrix are generated with our proposed method using the corresponding pose and the row-wise appearance.

### 4.3. Orthogonal Feature Space From Self-Disentanglement

We visually validate the orthogonality of the two feature spaces by reconstructing novel images using the pose feature from one image and the appearance feature from another. Fig. 4 shows a matrix of generated results on both the captured dataset and the synthetic dataset. We can successfully reconstruct the desired hand pose under different lighting and background. For instance, the hands in the first two rows of Fig. 4 (a) are lit by light source from the left, consistent in appearance with the source images. Even though the network cannot reproduce all the details in the background, it generates similar statistics. We refer readers to the supplementary video for more results.

There are still noticeable artifacts in the generated images, especially when the pose estimator does a poor job either in the appearance image (row 2 in Fig. 4(b)) or in the pose image (column 3 in Fig. 4(a)). Interestingly, because we don’t have any key points on the arm, it is encoded into the appearance feature by our network (row 6 in Fig. 4(a)).

### 4.4. Comparison to Supervised Disentanglement

To prove the effectiveness of our proposed self-disentanglement, we compare our method with two baselines: (1) the Auto-Encoder [32] with the structure shown in Sec. 3.2; (2) the factor disentanglement [9] using the paired data that have identical pose but different appearance. Detailed experimental results are shown in Tab. 2.

We can see that the images from the *Appearance Factor*

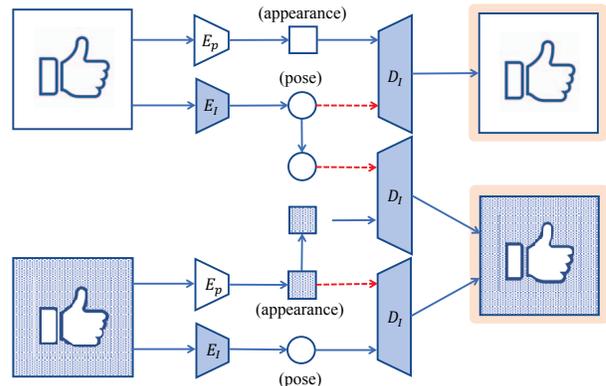


Figure 5. Factor Disentanglement with paired data [9]. The two inputs share the pose but differ in image appearance.

row and the Autoencoder results row in Tab. 2, are nearly the same. It shows that, without supervision on orthogonality, the Auto-Encoder model encodes the entire input image into the appearance feature, and discards the pose feature in decoding. Therefore, the pose and appearance factors are not fully disentangled. Checking the results of disentanglement with paired data [9] and our self-disentanglement, both methods are able to combine pose feature and appearance feature from two different source images to construct a novel image with specified pose and appearance. Our model generates visually similar images to the model trained with paired data.

Furthermore, we randomly swap the hand and appearance factors of the held-out set to generate a new set of im-

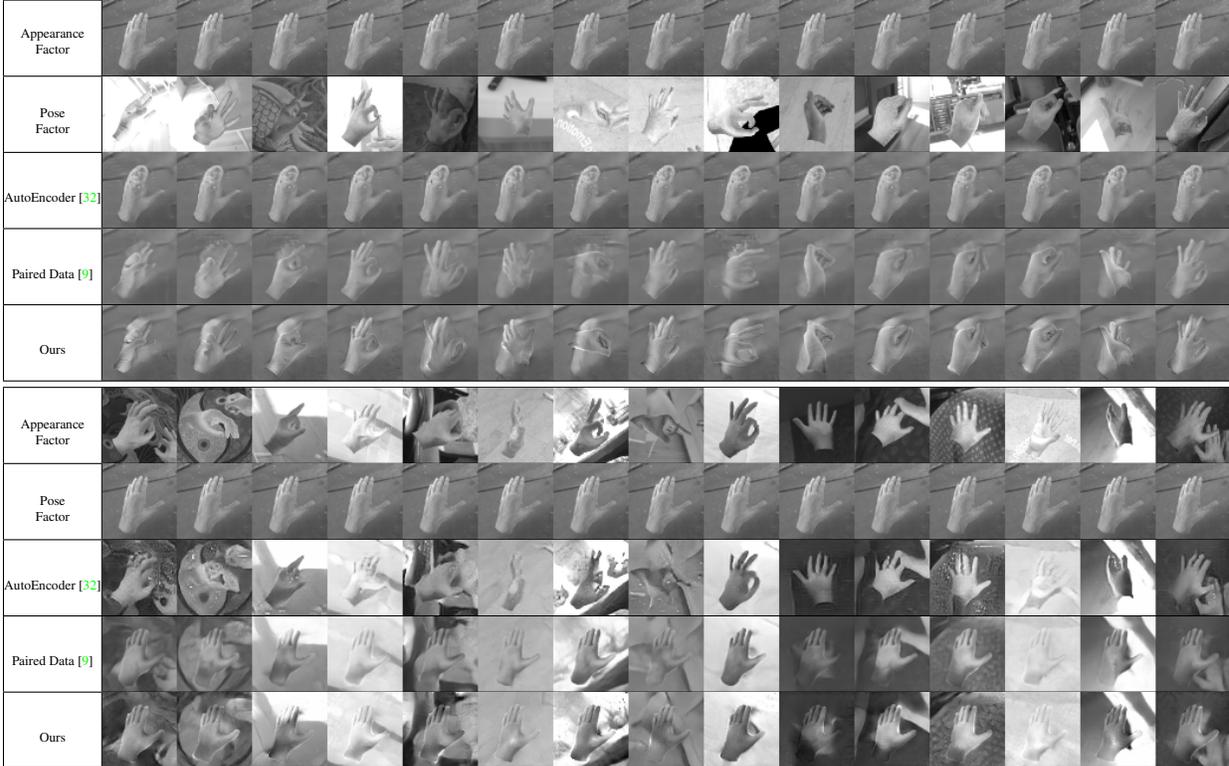


Table 2. Comparison with the existing methods on the paired synthetic data. Top part: fixed appearance factors with varying the pose ones. Bottom Part: varying pose factors with fixed appearance ones. **Appearance Factor** shows the images providing the appearance factors. **Pose Factor** shows the images providing the appearance factors. **AutoEncoder** denotes the image reconstruction along with the pose module shown in Sec. 3.2. **Paired Data** denotes the factor disentanglement using paired data [9]. **Ours** is our proposed self-disentanglement without leveraging the paired data.

Dataset	Paired Data [9]	Ours
I.S. [45]	$4.96 \pm 0.11$	$5.10 \pm 0.10$
Preference	51.66% (529 / 1024)	48.33% (495 / 1024)

Table 3. Quantitative Comparison of the factor disentanglement using Paired Data [9] and our proposed self-disentanglement, including Inception Scores [45] (I.S.) and User Study.

ages, and then calculate the inception scores [45] and perform a user study on the preference of between our method and [9] in Tab. 3. The comparable results validate our claims.

#### 4.5. Improve Pose Estimation with Disentanglement

An important application of our disentanglement model is to improve the robustness of the pose estimation module. We examine how each criterion in the disentanglement process affect the pose estimator step by step.

We fit the predicted heatmap of every joint to a gaussian distribution, and use the mean value as the predicted locations of the joints. Tab. 4 shows quantitative results, where the MSE denotes the mean square error of the predictions in pixels. The baseline pose estimator is trained with supervised learning (Sec 3.1). When we add the image recon-

struction loss (Sec 3.2, the accuracy is already improved by 4.60%. It suggests that the image reconstruction task encourages the shared base layers (Fig. 3) to extract more meaningful low level features for pose estimation. Adding the cycle consistency loss (Sec 3.4) further boosts the performance by 6.02%.

In Sec 3.5, we employ a strategy to stabilize training by disabling back-propagation to the pose feature (*Pose Detach*) as well as the back-propagation to the pose estimator parameters (*Pose Estimator Detach*). This is useful because the most reliable supervision is from the joint location labels, and we don't want to distract the pose estimator by auxiliary tasks that are more ambiguous in the early stage. However, once we have a reasonable disentanglement network, the additional supervision from image reconstruction and cycle consistency may help the pose estimator to better differentiate a pose from its appearance. We conduct two additional experiments using warm start from Model 3 to test this hypothesis. The new baseline trains our network as described in Sec 3.5 for another 75 epochs (Model 4). The first experiment allows back-propagation to the pose feature (Model 5), and the second experiment allows back-propagation to both the pose feature and the pose estimator

ID	Model	Epochs	MSE (in pixels)	Improvements
1	Baseline Pose Estimator	75	4.174	-
2	Pose Estimator + Image Reconstruction	75	3.982	4.60%
3	Our proposed Self-Disentanglement	75	3.923	6.02%
4	Our proposed Self-Disentanglement + Resume (**)	150	3.864	7.44%
5	(**) + No Pose Estimator Detach	150	3.756	10.02%
6	(**) + No Pose Estimator Detach + no Pose Detach	150	<b>3.735</b>	<b>10.53%</b>

Table 4. Ablation Study of the influence brought by Self-Disentanglement training on Hand Joints Localization. Mean Square Error (MSE) between the predicted location and the ground-truth is used to evaluate the accuracy, which is the lower the better. All models use the same model structure. **Resume** denotes resuming the training for another 75 epochs. **No Pose Estimator Detach** means when resuming training, the pose estimator will get trained on the mix-reconstructed images. **No Pose Detach** means when resuming training, the loss back-propagated from the image generation branch will go backward to the pose estimator through the pose factor.

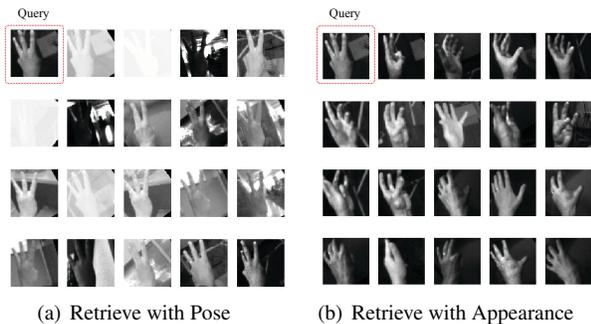


Figure 6. Image retrieval with disentangled factors.

parameters (Model 6). Both models are trained from Model 3 for 75 epochs. While the pose estimator benefits from warm start and the additional epochs, we can observe even greater improvements in accuracy when back-propagation is enabled. These two experiments demonstrate the effectiveness of self-disentanglement in improving the robustness of pose estimation to make it more resilient to environment variations.

#### 4.6. Image Retrieval using Disentangled Factors

We can examine the feature spaces by looking at images with similar features. For instance, if we query images with similar pose features, we will get images of similar hand poses under different environment variations. Likewise, if we query images with similar appearance features, we will get images in a similar environment but with different hand poses. Fig. 6 shows the top-20 nearest images from the monochrome dataset of the same query image in the pose space and the appearance space respectively. The query results further confirm the success of our method to disentangle the two factors.

### 5. Discussion

While we believe our method successfully disentangles pose features, we can only indirectly validate the result by reconstructing novel images from random pose-appearance feature pairs using a GAN framework. The reconstruction

captures the desired hand pose with consistent shading and background with the environment, but not without artifacts. The most severe issues are usually around the wrist and arm region, where the pose key points are sparse. Since key points are the only direct supervision, the model needs to differentiate hand pixels from background pixels based on the key points, and will make mistake where the connection is weak. Incorporation of pixel label masks or dense annotations as supervision to pose estimation and image reconstruction can potentially improve the image quality. Another interesting failure case is when the pose estimation makes a mistake, and the reconstruction image shows the wrongly estimated pose rather than the original input pose. It shows that while we are successful in disentanglement, there are other factors contributing to the robustness of pose estimation. In the future, we would like to investigate a more direct and quantitative measure of the effectiveness of disentanglement, and to improve the quality of image reconstruction to enrich any existing training dataset with a wide range of appearance variations.

### 6. Conclusion

In this paper, we present a self-disentanglement method to decompose a monochrome hand image into representative features of the hand pose and its complementary features of the image appearance. Without the supervision of paired images, we show that our method with cycle consistency principle is sufficient to ensure orthogonality of the pose feature and the appearance feature. Such flexibility makes our method applicable to any existing deep learning based pose estimation framework without requiring additional data or labels. When tested with a captured dataset of monochrome images, we demonstrate significant improvement in the robustness of the pose estimator to environment variations, comparing to a conventional supervised pose estimation baseline. Additionally, compared to a disentanglement model learned from paired training data, our model also performs similarly in terms of synthesized image quality proving the success of self-disentanglement.

## References

- [1] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*, 2015. **2**
- [2] S. Baek, K. In Kim, and T.-K. Kim. Augmented skeleton space transfer for depth-based hand pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8330–8339, 2018. **3**
- [3] G. Balakrishnan, A. Zhao, A. V. Dalca, F. Durand, and J. Guttag. Synthesizing images of humans in unseen poses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8340–8348, 2018. **3**
- [4] L. Ballan, A. Taneja, J. Gall, L. Van Gool, and M. Pollefeys. Motion capture of hands in action using discriminative salient points. In *European Conference on Computer Vision*, pages 640–653. Springer, 2012. **2**
- [5] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013. **3**
- [6] S. A. Bigdeli, M. Zwicker, P. Favaro, and M. Jin. Deep mean-shift priors for image restoration. In *Advances in Neural Information Processing Systems*, pages 763–772, 2017. **2**
- [7] H. Bourlard and Y. Kamp. Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, 59(4-5):291–294, 1988. **2**
- [8] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2016. **5**
- [9] A. Gonzalez-Garcia, J. van de Weijer, and Y. Bengio. Image-to-image translation for cross-domain disentanglement. *arXiv preprint arXiv:1805.09730*, 2018. **2, 4, 6, 7**
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. **4**
- [11] H. Hamer, K. Schindler, E. Koller-Meier, and L. Van Gool. Tracking a hand manipulating an object. In *Computer Vision, 2009 IEEE 12th International Conference On*, pages 1475–1482. IEEE, 2009. **1, 2**
- [12] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2017. **3**
- [13] G. E. Hinton, A. Krizhevsky, and S. D. Wang. Transforming auto-encoders. In *International Conference on Artificial Neural Networks*, pages 44–51. Springer, 2011. **2, 3**
- [14] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006. **2**
- [15] Q. Hu, A. Szabó, T. Portenier, M. Zwicker, and P. Favaro. Disentangling factors of variation by mixing them. *arXiv preprint arXiv:1711.07410*, 2017. **2, 3**
- [16] R. Huang, S. Zhang, T. Li, R. He, et al. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. *arXiv preprint arXiv:1704.04086*, 2017. **3**
- [17] C. Keskin, F. Kırac, Y. E. Kara, and L. Akarun. Real time hand pose estimation using depth sensors. In *Consumer depth cameras for computer vision*, pages 119–137. Springer, 2013. **1**
- [18] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. **5**
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. **1**
- [20] A. Kumar, P. Sattigeri, and A. Balakrishnan. Variational inference of disentangled latent concepts from unlabeled observations. *arXiv preprint arXiv:1711.00848*, 2017. **3**
- [21] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, et al. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems*, pages 5967–5976, 2017. **3**
- [22] T. Lee and T. Hollerer. Multithreaded hybrid feature tracking for markerless augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 15(3):355–368, 2009. **1**
- [23] M. Li, W. Zuo, and D. Zhang. Deep identity-aware transfer of facial attributes. *arXiv preprint arXiv:1610.05586*, 2016. **3**
- [24] Y. Li, N. Duan, B. Zhou, X. Chu, W. Ouyang, X. Wang, and M. Zhou. Visual question generation as dual task of visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6116–6124, 2018. **1**
- [25] Y. Li, W. Ouyang, X. Wang, and X. Tang. Vip-cnn: Visual phrase guided convolutional neural network. *CVPR*, 2017. **1**
- [26] Y. Li, W. Ouyang, B. Zhou, J. Shi, C. Zhang, and X. Wang. Factorizable net: An efficient subgraph-based framework for scene graph generation. In *The European Conference on Computer Vision (ECCV)*, September 2018. **1**
- [27] Y. Li, W. Ouyang, B. Zhou, K. Wang, and X. Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. **1**
- [28] Y. Liu, F. Wei, J. Shao, L. Sheng, J. Yan, and X. Wang. Exploring disentangled feature representation beyond face identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2080–2089, 2018. **4**
- [29] L. Ma, Q. Sun, S. Georgoulis, L. Van Gool, B. Schiele, and M. Fritz. Disentangled person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2018. **3**
- [30] X. Mao, C. Shen, and Y.-B. Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in neural information processing systems*, pages 2802–2810, 2016. **2**
- [31] A. Markussen, M. R. Jakobsen, and K. Hornbæk. Vulture: a mid-air word-gesture keyboard. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 1073–1082. ACM, 2014. **1**
- [32] J. Masci, U. Meier, D. Cireşan, and J. Schmidhuber. Stacked convolutional auto-encoders for hierarchical feature extraction. In *International Conference on Artificial Neural Networks*, pages 52–59. Springer, 2011. **6, 7**

- [33] M. F. Mathieu, J. Zhao, A. Ramesh, P. Sprechmann, and Y. LeCun. Disentangling factors of variation in deep representation using adversarial training. In *Advances in Neural Information Processing Systems*, pages 5040–5048, 2016. 2, 3, 4
- [34] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. Gnerated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–59, 2018. 1, 2
- [35] 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 *Proceedings of International Conference on Computer Vision (ICCV)*, volume 10, 2017. 2
- [36] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016. 2
- [37] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015. 2
- [38] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints. 2011. 2
- [39] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016. 2
- [40] X. Peng, X. Yu, K. Sohn, D. N. Metaxas, and M. Chandraker. Reconstruction-based disentanglement for pose-invariant face recognition. *intervals*, 20:12, 2017. 3, 4
- [41] T. Piumsomboon, A. Clark, M. Billingham, and A. Cockburn. User-defined gestures for augmented reality. In *IFIP Conference on Human-Computer Interaction*, pages 282–299. Springer, 2013. 1
- [42] T. J. Purcell, I. Buck, W. R. Mark, and P. Hanrahan. Ray tracing on programmable graphics hardware. In *ACM SIGGRAPH 2005 Courses*, page 268. ACM, 2005. 2
- [43] S. Rifai, Y. Bengio, A. Courville, P. Vincent, and M. Mirza. Disentangling factors of variation for facial expression recognition. In *Computer Vision–ECCV 2012*, pages 808–822. Springer, 2012. 3
- [44] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3, 4, 5
- [45] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *NIPS*, 2016. 7
- [46] S. Sridhar, A. M. Feit, C. Theobalt, and A. Oulasvirta. Investigating the dexterity of multi-finger input for mid-air text entry. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3643–3652. ACM, 2015. 1
- [47] 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 *European Conference on Computer Vision*, pages 294–310. Springer, 2016. 1, 2
- [48] S. Sridhar, A. Oulasvirta, and C. Theobalt. Interactive markerless articulated hand motion tracking using rgb and depth data. In *Proceedings of the IEEE international conference on computer vision*, pages 2456–2463, 2013. 1
- [49] A. Szabó, Q. Hu, T. Portenier, M. Zwicker, and P. Favaro. Challenges in disentangling independent factors of variation. *arXiv preprint arXiv:1711.02245*, 2017. 2
- [50] A. Tagliasacchi, M. Schröder, A. Tkach, S. Bouaziz, M. Botsch, and M. Pauly. Robust articulated-icp for real-time hand tracking. *Computer Graphics Forum*, 34(5):101–114, 2015. 1, 2
- [51] J. B. Tenenbaum and W. T. Freeman. Separating style and content. In *Advances in neural information processing systems*, pages 662–668, 1997. 3
- [52] J. B. Tenenbaum and W. T. Freeman. Separating style and content with bilinear models. *Neural computation*, 12(6):1247–1283, 2000. 3
- [53] J. Tompson, M. Stein, Y. Lecun, and K. Perlin. Real-time continuous pose recovery of human hands using convolutional networks. *ACM Transactions on Graphics (ToG)*, 33(5):169, 2014. 4, 5
- [54] L. Tran, X. Yin, and X. Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, volume 3, page 7, 2017. 3
- [55] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep image prior. 2018. 2
- [56] R. Wang, S. Paris, and J. Popović. 6d hands: markerless hand-tracking for computer aided design. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 549–558. ACM, 2011. 1, 2
- [57] L. Yang and A. Yao. Disentangling latent hands for image synthesis and pose estimation. *arXiv preprint arXiv:1812.01002*, 2018. 3
- [58] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *ICCV*, 2017. 2
- [59] C. Zimmermann and T. Brox. Learning to estimate 3d hand pose from single rgb images. In *International Conference on Computer Vision*, volume 1, page 3, 2017. 1, 2