

Person Recognition in Personal Photo Collections

Seong Joon Oh

Rodrigo Benenson

Mario Fritz

Bernt Schiele

Max Planck Institute for Informatics
Saarbrücken, Germany

{joon, benenson, mfritz, schiele}@mpi-inf.mpg.de

Abstract

Recognising persons in everyday photos presents major challenges (occluded faces, different clothing, locations, etc.) for machine vision. We propose a convnet based person recognition system on which we provide an in-depth analysis of informativeness of different body cues, impact of training data, and the common failure modes of the system. In addition, we discuss the limitations of existing benchmarks and propose more challenging ones. Our method is simple and is built on open source and open data, yet it improves the state of the art results on a large dataset of social media photos (PIPA).

1. Introduction

Person recognition in private photo collections is challenging: people can be shown in all kinds of poses and activities, from arbitrary viewpoints including back views, and with diverse clothing (e.g. on the beach, at parties, etc., see Figure 1). This paper presents an in-depth analysis of the problem of person recognition in photo albums: given a few annotated training images of a person (possibly from different albums), and a single image at test time, can we tell if the image contains the same person?

Intuitively, the ability to recognize faces in the wild [22] is an important ingredient. However, when persons are engaged in an activity (i.e. not posing) their face becomes only partially visible (non-frontal, occlusion) or simply fully non-visible (back-view). Therefore, additional information is required to reliably recognize people. We explore three other sources: first, body of a person contains information about their shape and appearance; second, human attributes such as gender and age help to reduce the search space; and third, scene context further reduces ambiguities.

The main contributions of the paper are the following. First, we provide a detailed analysis of performance of different cues (§3). Second, we propose a more realistic and challenging experimental protocols over PIPA (§5.1) on which a deeper understanding of robustness of different cues can be attained (§5.2). Third, in the process, we obtain



| | | | | |
|------------|---|---|---|---|
| Head | ✓ | ✗ | ✗ | ✗ |
| Body | ✓ | ✓ | ✗ | ✗ |
| Attributes | ✓ | ✓ | ✓ | ✗ |
| All cues | ✓ | ✓ | ✓ | ✓ |

Figure 1. Person recognition in photo albums is hard. To handle the diverse scenarios we need to exploit multiple cues from different body regions and information sources. Photos show test cases successfully recognised by our system, ticks indicate which ingredient could handle it. For example, the surfer is not recognised when using only head or head+body cues. However, it is successfully recognised when the additional attribute cues are provided.

best results on the recently proposed PIPA dataset and show that previous performance can be matched without specialized face recognition or pose estimation (§4). Fourth, we analyse remaining failure modes (§5). Additionally, our top-performing method is based only on open source code and data, and we will make the auxiliary data produced by us public, including the new attribute annotations over PIPA (§3.7) and the new experimental setups (§5.1).

1.1. Related work

Data type The bulk of previous work on person recognition focus either on facial features [22] (only the head/face is visible), or on the surveillance scenario [3, 2] (full body is visible, usually in low resolution). Both settings have seen a recent shift from sophisticated classifiers based on hand-crafted features and metric learning approaches [20, 7, 5, 29, 26, 41, 1], towards methods based on deep learning [37, 36, 43, 33, 27, 38, 21].

In this paper we tackle a different scenario, where persons may appear at different zoom levels (e.g. only head, upper torso, full body visible), and in any pose (e.g. sitting, running, posing), and from any point of view (e.g. front, side, back view), see Figures 1 and 7. The “Gallagher collection person dataset” [15] was the first dataset covering

this scenario; however, it is quite small (~600 images, 32 identities) and only frontal faces are annotated. We build our paper upon the recently introduced PIPA dataset [40] which is two orders of magnitude larger (~40k images, ~2k identities), more diverse, and also provides identity annotations when the face is not visible. We describe PIPA in more detail in §2.

Recognition tasks There exist multiple tasks related to person recognition [19] differing mainly in the amount of training and testing data. Face and surveillance re-identification is most commonly done via “verification” (one reference image, one test image; do they show the same person?) [22, 2]. The scenario of our interest is ~20 training images and one test image.

Other related tasks are, for instance, face clustering [9, 33], finding important people [30], or associating names in text to faces in images [13, 14].

Recognition cues The base cue for person recognition is the appearance of the face itself. Face normalization (“frontalisation”) [44, 37, 12] improves robustness to pose, view-point and illumination. Similarly, pose-independent descriptors can be built for the body [8, 17, 40].

Multiple other cues have been explored, for example: attributes classification [24, 25], explicit cloth modelling [15], relative camera positions [18], social context [16, 35], space-time priors [28], and photo-album priors [34].

The PIPA dataset was introduced together with the reference PIPER method [40]. PIPER obtains promising results combining three ingredients: a convnet (AlexNet [23]) pre-trained on ImageNet [10], the DeepFace re-identification convnet (trained on a large private faces dataset) [37], and Poselets [4] (trained on H3D) to obtain robustness to pose variance. In contrast, this paper considers features based on open data and use the same AlexNet network for all the image regions considered, thus providing a direct comparison of contributions from different image regions.

2. PIPA dataset

The recently introduced PIPA dataset (“People In Photo Albums”) [40] is, to the best of our knowledge, the first dataset to annotate identities of people with back views. The annotators labelled many instances that can be considered hard even for humans (Figure 7). PIPA features 37107 Flickr personal photo album images (Creative Commons), with 63188 head bounding boxes of 2356 identities. The dataset is partitioned into train, validation, test, and leftover sets, with rough ratio 45:15:20:20. Up to annotation errors, neither identities nor photo albums by the same uploader are shared among these sets.

For valid comparisons, we follow the PIPA protocol in [40]. The training set is used for feature learning and the validation set for exploring and optimising options. The

test set is for evaluation of our methods (Table 4); it is itself split in two parts, $\text{test}_0/\text{test}_1$, with roughly the same number of instances per identity. Given test_0 a classifier is learnt for each identity (11 examples per identity on average), and these are evaluated on test_1 (and vice-versa). Later we consider more challenging splits than the PIPA default (§5.1).

At test time, the system is fed with the photo of the test instance and the ground truth head annotation (tight around the skull, face and hair included; not fully visible heads are hallucinated by the annotators). The task is to find the corresponding identity of the head.

In the next section, various image regions and the corresponding recognition cues are defined (§3.1), and their validation set performances are compared (§3.3 to §3.7). The performance of our final system and comparisons to other methods and baselines are provided in §4. §5 will present an in-depth analysis of the systems, including the performance on the more realistic and challenging PIPA splits.

3. Cues for recognition

Our person recognition system is performant yet simple. At test time, given a (ground truth) head bounding box, we estimate (based on the box size) five different regions depicted. Each region is fed into one (or more) convnet(s) to obtain a set of feature vectors. The vectors are concatenated and fed into a linear SVM, trained per identity as one versus the rest (on $\text{test}_{0/1}$). In our final system all features are computed using the seventh layer of an AlexNet [23] pre-trained for ImageNet classification (albeit we explore alternatives in the next sections). The cues only differ among each other by the image region considered, and by the fine-tuning used to alter the AlexNet model (type of data or surrogate task). [23]

Compared to PIPER [40], we merge cues with a simpler schema and do not use specialized face recognition or pose estimation. Instead, we explore different directions: how informative are fixed body regions (no pose estimation) (§3.3)? How much does scene context help (§3.4)? And how much do we gain by using extended data (§3.6 & §3.7)? This section is based exclusively on validation set.

3.1. Image regions used

We choose five different image regions based on the ground truth head annotation (given at test time, see

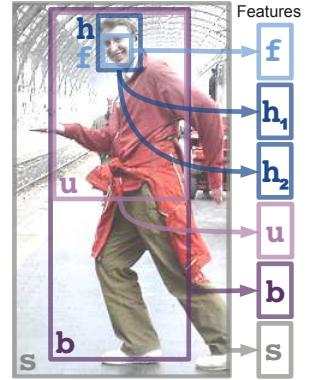


Figure 2. Regions considered for feature extraction: face f , head h , upper body u , full body b , and scene s . More than one feature vector can be extracted per region (e.g. h_1, h_2).

| Cue | | Accuracy |
|--------------|-------------|----------|
| Chance level | | 0.27 |
| Scene (§3.4) | s | 27.06 |
| Body | b | 80.81 |
| Upper body | u | 84.76 |
| Head | h | 83.88 |
| Face (§3.5) | f | 74.45 |
| Face+head | f+h | 84.80 |
| Full person | P = f+h+u+b | 91.14 |
| Full image | $P_s = P+s$ | 91.16 |

Table 1. Validation set accuracy of different cues. More detailed combinations in supplementary material.

§2). The head rectangle h corresponds to the ground truth annotation. The full body rectangle b is defined as $(3 \times \text{head width}, 6 \times \text{head height})$, with the head at the top centre of the full body. The upper body rectangle u is the upper-half of b . The scene region s is the whole image containing the head.

We use a face detector to find the face rectangle f inside each test head. We use the open source state of the art method of [31], which also provides a rough indication of the head yaw rotation (frontal, 45°, 90° side view). When no detection matches an annotation (e.g. back views), we regress the face area from the head bounding box. More details on the performance of this detector are in §3.5. Five respective image regions are illustrated in Figure 2.

Please note that regions overlap with each other, and that those pose agnostic crops may not match the actual regions.

3.2. Fine-tuning and parameters

Unless specified otherwise AlexNet is fine-tuned using PIPA’s person recognition training set ($\sim 30k$ instances, $\sim 1.5k$ identities), cropped at different regions, with 300k mini-batch iterations (batch size 50). We refer to the base cue thus obtained as f , h , u , b , or s , depending on the crop. On the validation set we found fine-tuning to provide a systematic ~ 10 percent points (pp) gain over not fine-tuned AlexNet. Since we use seventh layer of AlexNet, each cue adds 4 096 dimensions to our concatenated feature vector.

We train for each identity linear classifier using SVM regularization parameter $C = 1$. On the validation set the SVM classifier consistently outperforms by a ~ 10 pp margin the naive nearest neighbour (NN) classifier. Additional details can be found in supplementary material.

3.3. How informative is each image region ?

Table 1 shows the validation set results of each region individually and in combination. Head and upper body are the strongest individual cues. We discuss head and face in §3.5. Upper body is more reliable than the full body, because we observe that legs are commonly occluded (or out of frame) and thus become a distractor. Scene is, unsur-

| | Method | Accuracy |
|----------------------|-------------------------|----------|
| Gist | s_{gist} | 21.56 |
| PlacesNet scores | $s_{\text{places 205}}$ | 21.44 |
| raw PlacesNet | s_0 places | 27.37 |
| PlacesNet fine-tuned | s_3 places | 25.62 |
| raw AlexNet | s_0 | 26.54 |
| AlexNet fine-tuned | $s = s_3$ | 27.06 |

Table 2. Validation set accuracy of different feature vectors for the scene region s . See descriptions in §3.4.

prisingly, the weakest individual cue, but it still contains useful information for person recognition (far above chance level). Importantly, we see that all cues complement each other (despite having overlapping pixels).

Conclusion On the validation set at least, our features and combination strategy seems quite effective.

3.4. Scene (s)

Other than a fine-tuned AlexNet we considered multiple feature types to encode the scene information. s_{gist} : using the Gist descriptor [32] (512 dimensions). s_0 places: instead of using AlexNet pre-trained on ImageNet, we consider an AlexNet (PlacesNet) pre-trained on 205 scene categories of the “Places Database” [42] (~ 2.5 million images). $s_{\text{places 205}}$: Instead of the 4 096 dimensions PlacesNet feature vector, we also consider using the score vector for each scene category (205 dimensions). s_0, s_3 : we consider using AlexNet in the same way as for body or head (with zero or 300k iterations of fine-tuning on the PIPA person recognition training set). s_3 places: s_0 places fine-tuned for person recognition.

Results Table 2 compares the different alternatives on the validation set. The Gist descriptor s_{gist} performs only slightly below the convnet options (4 608 dimensional version of Gist gives worse results). Using the raw (and longer) feature vector of s_0 places is better than the class scores of $s_{\text{places 205}}$. Interestingly, in this context pre-training for places classification is better than pre-training for objects classification (s_0 places versus s_0). After fine-tuning s_3 reaches a similar performance as s_0 places. Experiments trying different combinations indicate that there is little complementarity between these features. Since there is not a large difference between s_0 places and s_3 , for sake of simplicity we use s_3 as our scene cue s in all other experiments.

Conclusion Scene by itself, albeit weak, can obtain results far above chance level. After fine-tuning, scene recognition as pre-training surrogate task [42] does not provide a clear gain over (ImageNet) object recognition.

| | Method | Accuracy |
|------------------------------------|--|----------|
| More data (§3.6) (head region) | h | 83.88 |
| | h + h _{cacd} | 84.88 |
| | h + h _{casia} | 86.08 |
| | h + h _{casia} + h _{cacd} | 86.26 |
| Attributes (§3.7) (head region) | h _{pipallm} | 74.63 |
| | h _{pipall} | 81.74 |
| | h + h _{pipall} | 85.00 |
| | u _{peta5} | 77.50 |
| (upper body region) | u + u _{peta5} | 85.18 |
| | A = h _{pipall} + u _{peta5} | 86.17 |
| | h + u | 85.77 |
| | h + u + A | 90.12 |

Table 3. Validation set accuracy of different cues based on extended data. See §3.6 and §3.7 for details.

3.5. Head (h) or face (f) ?

A large portion of work on face recognition focuses on the face region specifically. In the context of photo albums, we aim to quantify how much information is available in the head versus the face region.

The face region f is defined by a state of the art face detector [31] (see §3.1). Since no face annotations are available on PIPA, we validate the face detection location by learning a linear regressor from f to h (per DPM component). When using these heads estimates ($\sim 75\%$ of heads replaced) instead of the ground truth head (h in Table 1), results drop only 0.45% thus indirectly validating that faces are well localized.

Results When using the face region, there is a large gap of ~ 10 percent points performance between f and h in Table 1 highlighting the importance of including the head region around the face in the descriptor.

When evaluating only on the frontal faces of validation set (as indicated by the detector) f reaches 81% accuracy and 70% for non-frontal faces. The performance drop between frontal versus handling profile and back views is less dramatic than one could have suspected.

In comparison, on frontal faces in test set, DeepFace reaches $\sim 90\%$ [40], and returns the chance level (0.17%) otherwise. The test set contains about 50% of non-frontal faces. On test set f obtains 74% and 57% for frontal and non-frontal faces, respectively (18 pp drop), while h obtains 82% and 70%, respectively (12 pp drop).

Conclusion Using h is more effective than f, both due to improved recognition for frontal faces and robustness to head rotation. That being said, f results show fair performance to recognise non-frontal faces. As with other body cues, there is complementarity between h and f and we thus suggest to use them together.

3.6. Additional training data (h_{cacd}, h_{casia})

It is well known that deep learning architectures benefit from additional data. PIPER’s DeepFace is trained over $4.4 \cdot 10^6$ faces of $4 \cdot 10^3$ persons (the private SFC dataset [37]). In comparison our cues are trained over ImageNet and PIPA’s $29 \cdot 10^3$ faces over $1.4 \cdot 10^3$ persons. To measure the effect of training on larger data we consider fine-tuning using two open face recognition datasets: CASIA-WebFace (CASIA) [39] and the “Cross-Age Reference Coding Dataset” (CACD) [6].

CASIA contains $0.5 \cdot 10^6$ images of $10.5 \cdot 10^3$ persons (mainly actors and public figures), and is (to the best of our knowledge) the largest open dataset for face recognition. When fine-tuning AlexNet over these identities (using the head area h), we obtain the h_{casia} cue.

CACD contains $160 \cdot 10^3$ faces of $2 \cdot 10^3$ persons with varying ages. Although smaller than CASIA, CACD features greater number of face examples per subject ($\sim 2\times$). The h_{cacd} cue is built via the same procedure as h_{casia}.

Results The improvement of h + h_{cacd} and h + h_{casia} over h show that cues from outside training data are complementary to h (see top part of Table 3). h_{cacd} and h_{casia} on their own are about ~ 5 pp worse than h. h_{cacd} and h_{casia} exhibit slight complementarity.

Conclusion Adding more data, even from different type of photos, is an effective means to improve the performance.

3.7. Attributes (h_{pipall}, u_{peta5})

Albeit overall appearance might change day to day, one could expect that long term attributes provide means for recognition. We thus explore building feature vectors by fine-tuning AlexNet not on person recognition (like for all other cues), but rather for attributes classification as a surrogate task. We consider two sets of annotations.

We have annotated the PIPA train and validation sets (1409 + 366 identities) with five long term attributes: age, gender, glasses, hair colour, and hair length (11 binary bits total; see supplementary materials for details). We use the h crops to build h_{pipall}, as the attributes are head centric.

We also consider using the “PETA pedestrian attribute dataset” [11], which features 105 attributes annotations for $19 \cdot 10^3$ full-body pedestrian images. Out of 105 we chose the five binary attributes that are long term and are well represented in PETA: gender, age (young adult, adult), black hair, and short hair (details in supplementary material). Since upper-body u is less noisy than the full-body b (see Table 1), upper body crops of PETA are used to fine-tune AlexNet. The u_{peta5} cue is thus built.

Results For attribute fine-tuning we consider two approaches: training a single network for all attributes (multi-label classification with sigmoid cross entropy loss), or tuning one separate network per attribute (softmax loss) and

| | Method | Accuracy |
|------------------------|----------------------------|--------------|
| | Chance level | 0.17 |
| Body | GlobalModel [40] | 67.60 |
| | b | 69.63 |
| Head | DeepFace [40] | 46.66 |
| | h | 76.42 |
| Extended data | $h + h_{casia} + h_{cacd}$ | 79.63 |
| | PIPER [40] | 83.05 |
| Head+Body | $h+b$ | 83.36 |
| Full person | $P = f+h+u+b$ | 85.33 |
| Full image | $P_s = P+s$ | 85.71 |
| Extended data | naeil = P_s+E | 86.78 |
| Combining with [40] | PIPER [40]+P | 87.67 |
| | PIPER [40]+naeil | 88.37 |

Table 4. Test set accuracy of different cues and their combinations under the original PIPA evaluation protocol.

Extended data E = $h_{casia} + h_{cacd} + h_{pipal1} + u_{peta5}$.

then concatenating their feature vectors. The results on validation data indicate the second choice (h_{pipal1}) performs better than the first ($h_{pipal1m}$).

Table 3 (bottom) shows that attribute classification as surrogate task does help person recognition. Both PIPA (h_{pipal1}) and PETA (u_{peta5}) annotations behave similarly (~ 1 pp gain over h and u), and show good complementary among themselves (~ 5 pp gain over h+u). Amongst the attributes considered, gender contributes the most to improve recognition accuracy (for both attributes datasets).

Conclusion Adding attributes classification as a surrogate task improves performance.

4. Test set results

All experiments in this paper are limited to a person recognition scenario where head boxes are provided by human annotations, and all test faces belong to a known finite set. Table 4 reports the performance on the test set of the different cues described in previous sections. We study their complementarity to each other, and compare them against the PIPER components [40]. A more detailed table and the corresponding validation set results are included in the supplementary material.

We also report computational times for some pipelines in our method. The feature training takes 2-3 days on a single GPU machine. The SVM training takes 42.20s for h (4096 dim) and 1303.30s for naeil (4096 \times 17 dim) on the Original split (581 classes, 6443 samples). Note that this corresponds to a realistic user scenario in a photo sharing service where ~ 500 identities are known to the user and the average number of photos per identity is ~ 10 .

Compared to PIPER, our framework is computationally efficient in two aspects. First, our system does not need to learn to assign weights for different cues. Second, the

PIPER feature has roughly 4096×108 dimensions, requiring far more memory and training time than our final system (4096×17 dim).

Body b Considered alone, our body cue b is a reimplementation of PIPER’s GlobalModel [40]. As expected, we obtain a similar accuracy.

Head h On the other hand, our head cue h is more effective than the corresponding PIPER’s DeepFace. As discussed in §3.5, we have observed that: a) for this task the head region is more informative than the face (focusing on the face region is detrimental); b) our approach is much more robust for non-frontal faces ($\sim 50\%$ of test cases), where h reaches 70% accuracy, while DeepFace becomes uninformative in this case. When extending the training data our head performance further improves (see also the discussion in §5.4).

Head+body h+b Our minimal system matching PIPER’s performance is h+b, with accuracy 83.36%. Note that the feature vector of h+b is about 50 times smaller than PIPER’s.

In principle PIPER captures the head region via one of its poselets. Thus, h+b extracts cues from a subset of PIPER’s “GlobalModel+Poselets” [40], which only reaches 78.79%.

Full person P Similar to the validation set results, having more cues further improves results. $P = f+h+u+b$ obtains a clear margin over PIPER, yet is a simpler system (neither specialised face recognition nor pose estimation used) built with less training data (only PIPA for fine-tuning, ImageNet for pre-training, and the face detector training set).

naeil Adding scene s (§3.4) and extended data E (§3.6 & §3.7) contributes to the last 1 percent point. We name our final method naeil¹. Its feature vector is 6 times smaller than PIPER’s, and it provides the best known results on the PIPA dataset.

Figures 1 and 7 show some example results of our system. §5.4 analyses the remaining hard test cases.

4.1. Complementarity between PIPER and naeil

Since PIPER uses different training data than naeil we can expect some complementarity between the two methods. For experiments, we use the PIPER scores provided by the authors of [40]. Note, however, that the PIPER features are unavailable. By averaging the output scores of the two methods (PIPER+naeil) gain ~ 1.5 percent points, reaching 88.37%. Using a more sophisticated strategy might provide more gain, but we already see that naeil covers most of the performance from PIPER.

¹“naeil”, 내일, means “tomorrow” and sounds like “nail”.

4.2. Towards an open world setting

All experiments in this paper are limited to a person recognition scenario where head boxes are provided by human annotations, and all test faces belong to a known finite set. Not providing ground truth heads at test time is an arguably more realistic and challenging scenario in which both person detection and recognition need to be solved jointly.

Using a face detector (§3.5) as our person detector over the test set, we reach $\sim 78\%$ recall at (average) ten detections per image ($\sim 70\%$ at 3 detections/image). If we use naeil to label these faces, we reach $\sim 65\%$ recall on the $\text{test}_{0/1}$ identities ($\sim 62\%$ at 3 detections/image).

The performance drops, but less dramatically than what one might expect. It remains as future work to implement a detailed evaluation in the open world setting.

4.3. A naive baseline

Given the inherent difficulty of the PIPA person recognition task (see Figure 7) reaching a $\sim 85\%$ accuracy seems suspiciously high. Thus, we investigate the issue using a crude baseline h_{rgb} that takes the raw RGB pixel values of the head area as features (after downsizing to 40×40 pixels and blurring), and uses a nearest neighbour classifier. By design h_{rgb} is only able to recognize near identical heads across the $\text{test}_{0/1}$ split, yet it reaches a surprisingly high 33.77% (49.46%) accuracy on the test set (validation set).

Conclusion About $1/3$ of the original PIPA test splits is easy to solve. This motivates us to explore more realistic splits and protocols. In the next section we discuss the issue and propose solutions via new test splits.

5. Analysis of person recognition challenges

This section provides a detailed analysis of the obtained results and shares insights on addressing future challenges.

As we have seen in §4.3, the current setup includes many easy examples, limiting us from exploring more difficult dimensions of the problem. Accordingly, we propose three new $\text{test}_0/\text{test}_1$ splits of PIPA in §5.1. Based on the new splits, we analyse the robustness of different cues across appearance changes (§5.2). We then discuss the effect of the amount of person specific training data (§5.3), and provide a failure mode analysis in §5.4.

5.1. New PIPA splits with varying difficulty and challenges

We have seen a strong performance of our main system naeil (86.78% on test set, Table 4) and the baseline h_{rgb} (33.77% on test set, §4.3) despite the challenging task of person recognition in photo albums. This motivates us to investigate more difficult and realistic setups.



Figure 3. Visualisation of Original and Day splits for one identity. Greater appearance changes are observed across the Day split.

Limitations of original setup The main limitation of the original PIPA protocol is that the $\text{test}_0/\text{test}_1$ splits are even-odd instances from a sample list that largely preserves the photo orders in albums. When photos are taken in a short period of time, adjacent photos can be nearly identical. However, a main challenge in person recognition is to generalise across long-term appearance changes of a person; we thus introduce a range of new splits on PIPA in the order of increasing difficulty:

Original split \mathcal{O} : We keep the original split in our study for comparison. The split is on the odd vs even basis.

Album split \mathcal{A} : All samples are organised by albums. This split assigns for each person identity samples from separate albums, while keeping the number of samples equal for the splits. Since it is not always possible to satisfy both conditions, a few albums are shared between the splits. In this setup, training and test samples are split across different events and occasions.

Time split \mathcal{T} : This split investigates the temporal dimension of the photos. For each person identity, we sort all photos by their “photo-taken-date” metadata. We split them into newest versus oldest images. The instances without time metadata are distributed evenly. This split emphasises the temporal distance between training and test.

Day split \mathcal{D} : \mathcal{T} does not always make a time gap: many people appear only on one event, and the time metadata are often missing. We thus make the split manually according to the two conditions: either a firm evidence of date change such as {change of season, continent, event, co-occurring

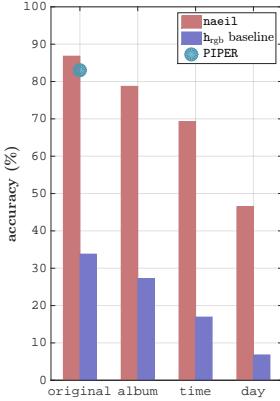


Figure 4. Recognition accuracy across different experimental setups on test set.

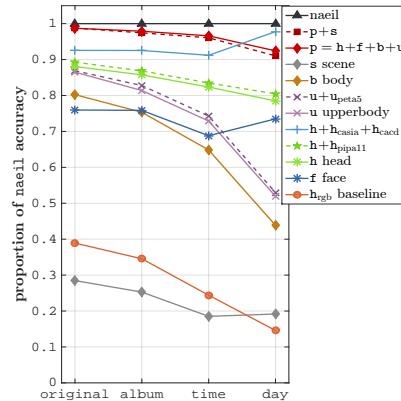


Figure 5. Test set accuracy of cues in different settings, relative to naeil.

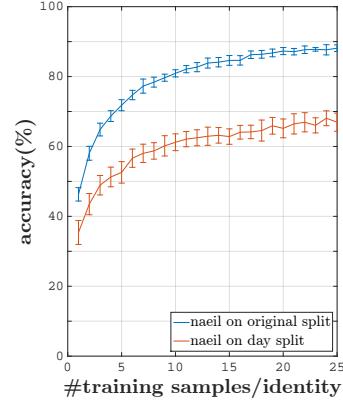


Figure 6. Recognition accuracy at different sizes of training examples.

people} between the splits, or visible changes in {hairstyle, make-up, head or body wear}. These rules enforce “appearance changes”. For each identity, we randomly discard instances from the larger test set until sizes match. If there are less than 5 instances in the split, we discard the identity altogether (Original split applies the same criterion). After pruning, 199 identities (out of 581) were left, with about 20 training samples per identity (similar range as all other splits).

Results Figure 4 provides an overview of how the raw colour baseline h_{rgb} and our naeil approach perform across different splits. We observe that the unreasonably good performance by the h_{rgb} baseline consistently degrades from the Original over Album and Time to the Day splits, indicating the increasing amount of non-trivial recognition tasks. Compared to the 1/5 drop by the h_{rgb} baseline (33.77% to 6.78%), naeil’s performance is less impaired (86.78% to 46.48%), indicating naeil’s ability to address more realistic scenarios characterised by changes in appearance, location and time.

5.2. Importance of features

To gain a deeper understanding of relative importance of different cues and their robustness across splits, we consider Figure 5 which shows the results normalised by the performance of naeil (100%). This allows us to analyse which features maintain, loose or gain discriminative power when moving from the easier to the more challenging settings.

We observe the strongest drops in relative performance for body and upper body features, due to the loss of discriminability of global features (e.g. clothing). We see consistent gains for using surrogate training tasks such as attributes (h_{piap11} , u_{petas}) and, more prominently, external data for head features (h_{asia} , h_{caca}). External data for head features particularly pays off for the most difficult day split.

Conclusion The usage of significantly larger databases improves the robustness of our features, enabling recognition in the most challenging scenarios.

5.3. Importance of training data

We also investigate how much collecting more data from each person identity can help to improve performance. In Figure 6 we compare the Original to the Day split and show performance for different sizes of training samples. While on the original split already after 10 training examples 80% performance is reached, the performance on the Day split sees a relatively slow improvement and stays below 70% with 25 samples (lagging 20% behind the Original split).

Conclusion From Figure 6 we see that only increasing the training data will struggle to solve the harder Day split. Better features and better methods are required.

5.4. Analysis of remaining failure modes

In the supplementary material we provide detailed statistics to study failure modes in the Original and Day splits. We discuss here the main findings.

As expected, non-frontal faces are common failure cases for naeil’s in both Original and Day splits (~50%). For frontal faces, we observe in the Day split a larger proportion of failures than in the Original split. Even more, the majority of failures correspond to large heads (height > 100 pixels), where good features can be extracted. To handle better more realistic scenarios it is thus important to improve the recognition of frontal faces across diverse settings and long time-spans.

Another interesting aspect is that while naeil on the Original split has only one identity (out of 581) which is never correctly predicted, on the Day split the proportion of never correct identities jumps to 20%. This suggests that there are inherently difficult identities that our simplistic system currently cannot handle.



Figure 7. Examples of success cases on the Original split. First column shows the test instances that our systems correctly predict. Columns 5-7 correspond to train instances of the correct identity. Columns 2-4 are the training examples of the identity that PIPER [40] wrongly predicts. From top to bottom, the shown test instances are: (1) success case of $f+h$ and failure case of PIPER; (2) success case of $p=f+h+u+b$ and failure case of PIPER and $f+h$; (3) success case of $p+s$ and failure case of PIPER and p ; and (4) success case of naeil, and failure case of PIPER and $p+s$.

6. Conclusion

We analysed the problem of person recognition in photo albums where people appear with various viewpoints, poses and occlusions. There are four major conclusions from our studies. First, head region, even when face is not visible, is a strong cue for person recognition, better than the face region itself (§3.5). Second, different cues, although from overlapping regions, are complementary (§4). Third, feature learning with massive database of faces improves robustness across time and appearance (§5.2). Fourth, simply increasing the number of training examples per person does not automatically solve the problem, and better recognition systems must be devised (§5.3).

One possible research direction is collecting a large database of personal photo albums on which better features can be trained. One can also exploit album context, which is a rich source of identity information [16, 34, 35]; however, it was not used in this work for fair comparison.

Our experimental data will be published, e.g. attribute annotations, new splits, result files, and the trained models.

References

- [1] S. Bak, R. Kumar, and F. Bremond. Brownian descriptor: A rich meta-feature for appearance matching. In *WACV*, 2014.
- [2] A. Bedagkar-Gala and S. K. Shah. A survey of approaches and trends in person re-identification. *IVC*, 2014.
- [3] B. Benfold and I. Reid. Guiding visual surveillance by tracking human attention. In *BMVC*, 2009.
- [4] L. Bourdev and J. Malik. Poselets: Body part detectors trained using 3d human pose annotations. In *ICCV*, 2009.
- [5] X. Cao, D. Wipf, F. Wen, and G. Duan. A practical transfer learning algorithm for face verification. In *ICCV*, 2013.
- [6] B.-C. Chen, C.-S. Chen, and W. H. Hsu. Cross-age reference coding for age-invariant face recognition and retrieval. In *ECCV*, 2014.
- [7] D. Chen, X. Cao, F. Wen, and J. Sun. Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification. In *CVPR*, 2013.

- [8] D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, and V. Murino. Custom pictorial structures for re-identification. In *BMVC*, 2011.
- [9] J. Cui, F. Wen, R. Xiao, Y. Tian, and X. Tang. Easyalbum: an interactive photo annotation system based on face clustering and re-ranking. In *SIGCHI*, 2007.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [11] Y. Deng, P. Luo, C. C. Loy, and X. Tang. Pedestrian attribute recognition at far distance. In *ACMMM*, 2014.
- [12] C. Ding and D. Tao. A comprehensive survey on pose-invariant face recognition. *arXiv*, 2015.
- [13] M. Everingham, J. Sivic, and A. Zisserman. Hello! my name is... buffy—automatic naming of characters in tv video. In *BMVC*, 2006.
- [14] M. Everingham, J. Sivic, and A. Zisserman. Taking the bite out of automated naming of characters in tv video. *IVC*, 2009.
- [15] A. Gallagher and T. Chen. Clothing cosegmentation for recognizing people. In *CVPR*, 2008.
- [16] A. C. Gallagher and T. Chen. Using group prior to identify people in consumer images. In *CVPR*, 2007.
- [17] V. Gandhi and R. Ronfard. Detecting and naming actors in movies using generative appearance models. In *CVPR*, 2013.
- [18] R. Garg, S. M. Seitz, D. Ramanan, and N. Snavely. Where's waldo: matching people in images of crowds. In *CVPR*, 2011.
- [19] S. Gong, M. Cristani, S. Yan, and C. C. Loy. *Person re-identification*. Springer, 2014.
- [20] M. Guillaumin, J. Verbeek, and C. Schmid. Is that you? metric learning approaches for face identification. In *ICCV*, 2009.
- [21] Y. Hu, D. Yi, S. Liao, Z. Lei, and S. Li. Cross dataset person re-identification. In *ACCV workshop*, 2014.
- [22] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, UMass, 2007.
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [24] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *CVPR*, 2009.
- [25] R. Layne, T. M. Hospedales, S. Gong, and Q. Mary. Person re-identification by attributes. In *BMVC*, 2012.
- [26] W. Li and X. Wang. Locally aligned feature transforms across views. In *CVPR*, 2013.
- [27] W. Li, R. Zhao, T. Xiao, and X. Wang. Deep-reid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014.
- [28] D. Lin, A. Kapoor, G. Hua, and S. Baker. Joint people, event, and location recognition in personal photo collections using cross-domain context. In *ECCV*, 2010.
- [29] C. Lu and X. Tang. Surpassing human-level face verification performance on lfw with gaussianface. *arXiv*, 2014.
- [30] C. S. Mathialagan, A. C. Gallagher, and D. Batra. Vip: Finding important people in images. *arXiv*, 2015.
- [31] M. Mathias, R. Benenson, M. Pedersoli, and L. Van Gool. Face detection without bells and whistles. In *ECCV*, 2014.
- [32] A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV*, 2001.
- [33] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. *arXiv*, 2015.
- [34] J. Shi, R. Liao, and J. Jia. Codel: A human co-detection and labeling framework. In *ICCV*, 2013.
- [35] Z. Stone, T. Zickler, and T. Darrell. Autotagging facebook: Social network context improves photo annotation. In *CVPR workshops*, 2008.
- [36] Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. *arXiv*, 2014.
- [37] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, 2014.
- [38] D. Yi, Z. Lei, and S. Z. Li. Deep metric learning for practical person re-identification. *arXiv*, 2014.
- [39] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning face representation from scratch. *arXiv*, 2014.
- [40] N. Zhang, M. Paluri, Y. Taigman, R. Fergus, and L. Bourdev. Beyond frontal faces: Improving person recognition using multiple cues. In *CVPR*, 2015.
- [41] R. Zhao, W. Ouyang, and X. Wang. Person re-identification by salience matching. In *ICCV*, 2013.
- [42] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning Deep Features for Scene Recognition using Places Database. *NIPS*, 2014.
- [43] E. Zhou, Z. Cao, and Q. Yin. Naive-deep face recognition: Touching the limit of lfw benchmark or not? *arXiv*, 2015.
- [44] Z. Zhu, P. Luo, X. Wang, and X. Tang. Deep learning identity-preserving face space. In *ICCV*, 2013.