

---

# TACKLING FACE VERIFICATION EDGE CASES: IN-DEPTH ANALYSIS AND HUMAN-MACHINE FUSION APPROACH

---

A PREPRINT

✉ **Martin Knoche**    ✉ **Gerhard Rigoll**  
Technical University of Munich  
Arcisstrasse 23, 80333 München, Germany  
Martin.Knoche@tum.de

April 19, 2023

## ABSTRACT

Nowadays, face recognition systems surpass human performance on several datasets. However, there are still edge cases that the machine can't correctly classify. This paper investigates the effect of a combination of machine and human operators in the face verification task. First, we look closer at the edge cases for several state-of-the-art models to discover common datasets' challenging settings. Then, we conduct a study with 60 participants on these selected tasks with humans and provide an extensive analysis. Finally, we demonstrate that combining machine and human decisions can further improve the performance of state-of-the-art face verification systems on various benchmark datasets. Code and data are publicly available on GitHub<sup>1</sup>.

## 1 Introduction

In recent years, face recognition technology has become increasingly prevalent in various fields, such as security, surveillance, and healthcare. The continuous performance improvement can be primarily attributed to the development of deep learning models. These models have reached near saturation levels on standard benchmark datasets like LFW [1]. However, several challenges remain to be addressed. For instance, the performance of face recognition systems is highly dependent on the quality of input images, as demonstrated by [2], and is susceptible to pose [3] or age [4] variations. More recently, the COVID-19 pandemic highlighted that face masks drastically affect the accuracy of traditional face recognition systems [5]. Despite ongoing advancements, current face recognition systems are still need to be robust enough to achieve 100% accuracy and thus require supervision by human operators in real-world applications. In this context, the human expert is a so-called "super-recognizer", making decisions based on their knowledge and the system's output.

Inspired by the recent work of Zheng *et al.* [6], which proposes a fusion of a neural network with pathologists for predicting EBVaGC from histopathology, we aim to propose and analyze a similar fusion of state-of-the-art deep learning networks with human decisions for face verification. In contrast to Carragher's *et al.* [7] previous work, our focus is on prioritizing machine decisions over human decisions.

In this work, we explore whether there are simple problems humans can solve but remain challenging for machines. To this end, we examine common face recognition benchmark datasets, focusing on cases where state-of-the-art models struggle. We hypothesize that humans can accurately and confidently solve the selected face verification problems that are difficult for machines. To test this hypothesis, we first conduct a baseline study with 60 participants to assess the consistency of human performance in face verification tasks without prior experience. Next, we carried out an investigation involving the same participants, asking them to verify selected image pairs that state-of-the-art face recognition models find challenging and to rate their confidence in their decisions.

---

<sup>1</sup><https://github.com/anonymous/anonym>

Our main contributions are as follows:

- We analyze the performance of state-of-the-art face verification models on several standard benchmark datasets by examining specific image pairs where the models fail or exhibit low confidence in their decisions.
- We conducted an in-depth analysis of a survey involving 60 participants, focusing on selected image pairs and highlighting the differences between machine and human performance, taking into account the participants' confidence levels.
- We demonstrate that a clever combination of machine and human decisions can enhance the overall performance of face verification systems on various benchmark datasets.

## 2 Related Work

Human performance in face verification or face matching has been extensively studied. However, only a few works have investigated human accuracy on popular datasets commonly used to evaluate automatic face recognition systems. In [8], the authors conducted a study using Amazon Mechanical Turk workers and evaluated human accuracy on the LFW [1] database, achieving 99.2%. To the best of our knowledge, no studies have yet been conducted on human accuracy for other large benchmark datasets.

The following studies delve more into the psychological aspects of face verification:

Howard et al. [9] investigated the influence of prior face identity decisions on subsequent human judgments about face similarity. Their study revealed that prior identity decision labels alter volunteers' internal criteria for judging face pairs.

Fysh and Bindemann [10] examined the impact of onscreen trial labels with consistent, inconsistent, or unresolved information on face-matching accuracy. Their experiments demonstrated that human face-matching decisions are influenced by onscreen identifications, leading to increased accuracy when information is correct and increased errors when information is misleading.

In [11], the authors explored the potential improvement in accuracy resulting from using multiple-image arrays in unfamiliar face-matching tasks. The study found that simultaneous viewing of multiple images of the same person improved matching accuracy.

Other works [12–14] have focused on developing standardized tests for human analysis in face-matching tasks.

## 3 Method

### 3.1 Survey Construction

For constructing the surveys in our study, we utilize various modalities encompassed within commonly used face verification benchmark datasets, such as cross-age (CALFW)[4], cross-pose (CPLFW)[3], cross-quality (XQLFW)[15], and masks (MLFW)[5]. All datasets contain 1000 genuine (same identity) and 1000 imposter (different identity) color image pairs, with a physical image resolution of 112 × 112 pixels.

First, we obtain predictions from three state-of-the-art face recognition models: 1) ArcFace\*, a network based on ResNet50 [16] and trained with additive angular margin loss [17]. 2) FaceTrans\*, a vision transformer architecture trained on faces [18]. We use Knochel et al. [19] publicly available and more robust re-tuned variants of both models. 3) ProdPoly, a polynomial neural network architecture trained on faces proposed by Chakraborty [20]. We also calculate their certainty scores using the method proposed by Kechen [21].

Next, we divide the mean certainty of all models, ranging from 0 to 100%, into five equally sized bins. We then randomly select one imposter and one genuine image pair from each bin and each dataset, resulting in 40 image pairs. This collection serves as the pool of image pairs for our baseline survey, which we will derive Survey 1.

Second, to construct Survey 2–5 we divide the mean certainty of all models into ten equally sized bins and randomly select 30 samples from each bin to create a pool. Additionally, we identify all image pairs from the abovementioned databases where the minimum certainty of all three models falls below a threshold of 50% and include them in the pool. This procedure results in four pools containing 745 image pairs. Table 1 summarizes the image pair statistics for all constructed surveys.

Table 1: Detailed image pair statistics for all surveys.

| Survey           | Dataset |        |        |        |        |
|------------------|---------|--------|--------|--------|--------|
|                  | BASE    | CALFW  | CPLFW  | MLFW   | XQLFW  |
|                  | 1       | 2      | 3      | 4      | 5      |
| # Pairs in Pool  | 40      | 684    | 1385   | 3561   | 1815   |
| # Pairs per User | 40      | 12     | 24     | 60     | 31     |
| % of Total Pairs | -       | 5.37%  | 10.28% | 28.58% | 13.58% |
| % Genuine        | 50.00%  | 47.08% | 44.55% | 48.16% | 44.90% |

### 3.2 Study Design

Each of our 60 participants must register for our online study via a website, where they can provide their age, ethnicity, and gender. After registration, participants can complete the survey at their convenience. We have *Survey 1–5* available for participants to complete.

The baseline *Survey 1* is identical for all participants, *i.e.*, all 40 image pairs are displayed to every user in the same order.

All subsequent surveys are personalized for each participant. The image pairs for *Survey 2–5* are randomly selected from the corresponding pool of images (see Subsec. 3.1) and are unique, meaning no image pair is displayed twice. The ratio of genuine and imposter pairs varies for each user. Every user is asked to answer 167 image pairs, resulting in 9 845 questions. The last three participants received fewer image pairs as the respective image pair pools were depleted.

Participants are not informed of the ratio of genuine and imposter pairs or the labels of the image pairs. To prevent any negative influence on the participants’ performance, we do not display the models’ predictions to the participants as proposed in [10]. We also do not display the outcome of previous questions to the participants. This could influence their decision on subsequent questions, particularly by revealing whether the previous answer was correct or incorrect [9].

For each image pair in the survey datasets, participants need to decide whether the two images show the same or different identities and rate their certainty regarding the answer to the first question. Participants adjust a slider, which ranges from “uncertain” to “certain” without any numerical indicators or additional information. The default position of the slider is at the midpoint, corresponding to a value of 0.5, with the internal range of values extending from 0 (uncertain) to 1 (certain). The duration is also recorded.

After completing at least one question in the survey, participants can leave the survey and review the accuracies of their answers thus far, along with the corresponding FaceTrans\* model’s performance.

The survey platform’s code was written in Python using the Flask framework and hosted on Google Cloud.

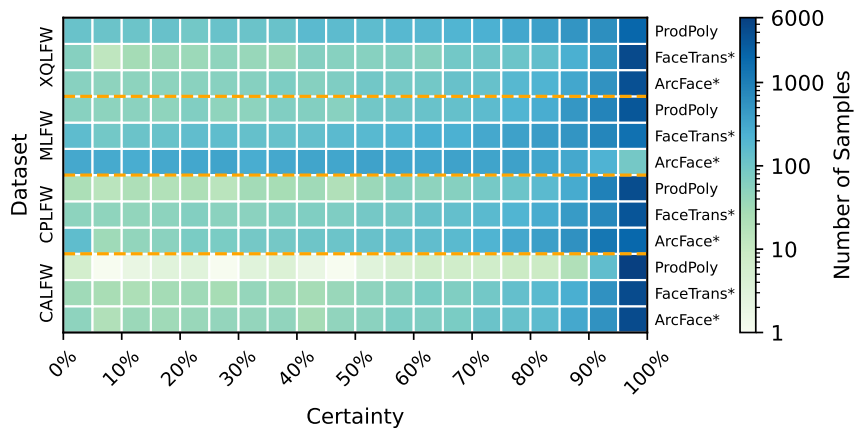


Figure 1: Certainty distributions on standard datasets for face recognition networks.




Figure 2: Example image pairs with correct human decisions and higher certainty than machine.

### 3.3 Human-Machine Fusion

In this work, we examine human-machine decision fusion in face verification. Our fusion algorithm can be outlined: We set a minimum certainty threshold  $t_{\min}$  for machine decisions to be directly accepted. Any machine decisions falling below this threshold are considered for comparison with the calibrated human certainty. If the human certainty for a particular question exceeds a specified threshold  $t_{\max}$  and surpasses the machine certainty by at least  $d_{\min}$ , we accept the human decision. Otherwise, we maintain the machine’s decision. After collecting human feedback on the pool of image pairs for all datasets, the proposed fusion algorithm is applied to all 6 000 pairs in each dataset.

## 4 Results

### 4.1 Dataset Insights

In Fig. 1, we present the distribution of certainty scores for our three models on the four datasets (see Subsec. 3.1). For all datasets, it is evident that the machine predicts with very high certainty ( $\geq 95\%$ ) most of the time. The masked dataset (MLFW) appears to be the most challenging for the models, which is also reflected in the accuracy (see Table 2). Additionally, the number of image pairs with certainty  $\leq 50\%$  is the highest in this dataset (refer to Table 1).

Furthermore, we conduct a closer visual inspection of the image pairs in the datasets, selecting pairs where the machine prediction is incorrect and has low confidence. In contrast, the human prediction is correct with very high confidence. In Fig. 2, we showcase three imposter and three genuine image pair examples that meet these criteria. Human operators made correct decisions for all these samples and rated them with 100% certainty, while the machine was incorrect and had much lower confidence.

### 4.2 Participants Characteristics

The data for this study was collected online over 14 days in March/April 2023. Fig. 3 displays the age, gender, and ethnicity distribution of the participants. The mean age is 37.22 (std=17.41, min=20, max=72), with 60% of the volunteers being male, 33.33% female, and four participants preferring not to state their gender. We primarily recruited

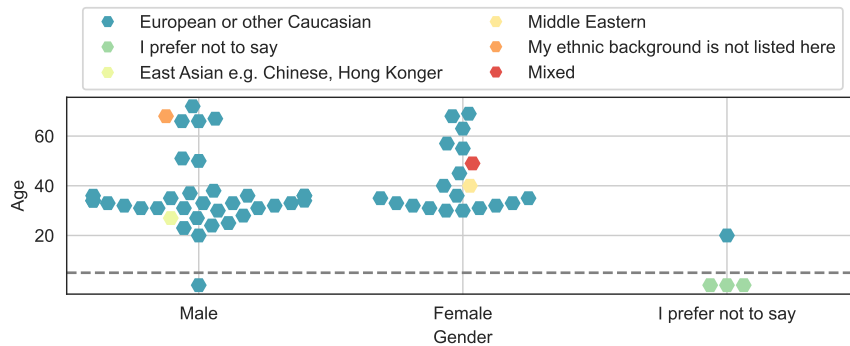


Figure 3: Age, gender and ethnicity distribution of the participants.

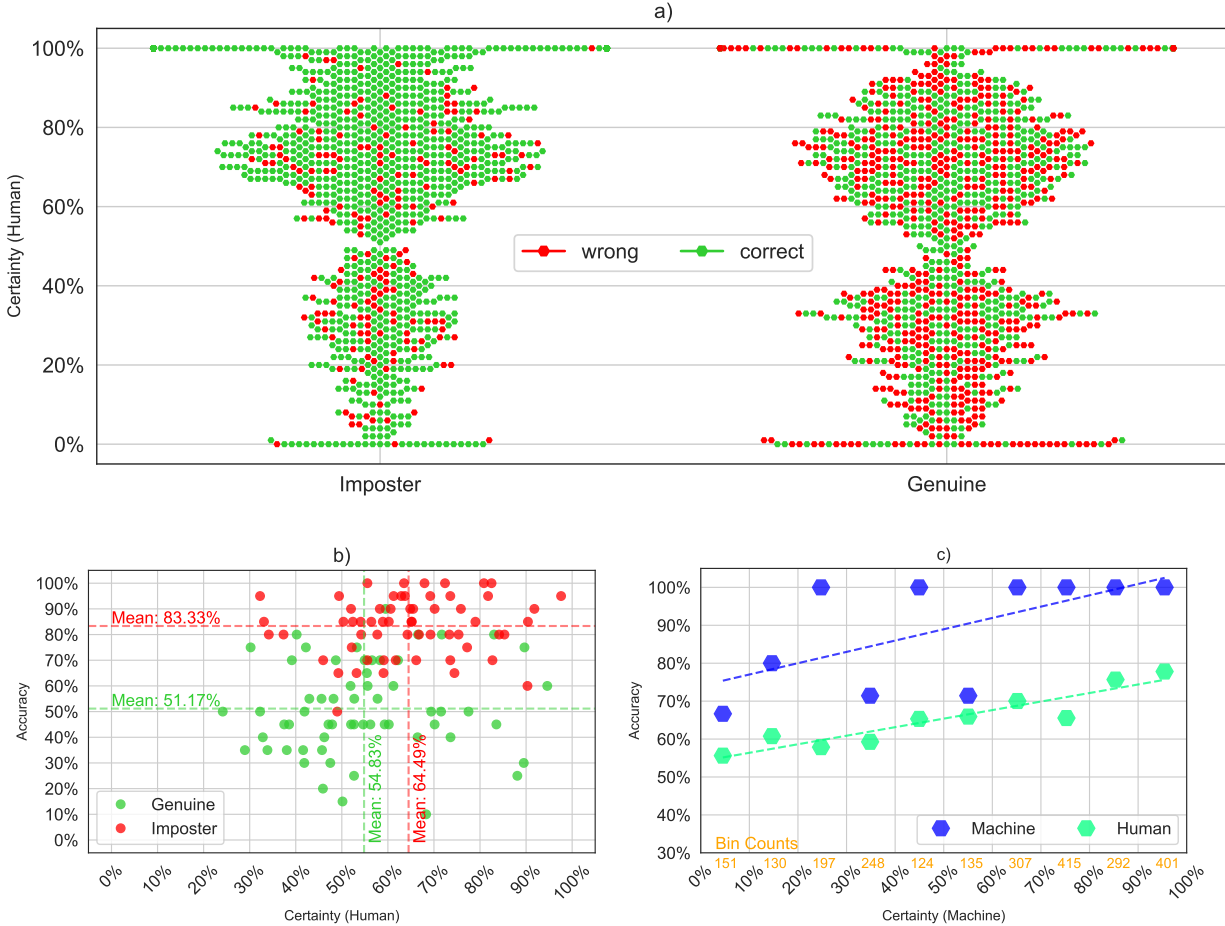


Figure 4: a) User prediction distribution b) Mean certainty and accuracy per user. c) Accuracy comparison for machine certainty histogram bins.

European/Caucasian users, which, although it may limit the generalizability of the study, more closely reflects the ethnicity distribution of the datasets used in our research, as reported by [22].

### 4.3 Experiment 1

The collected feedback from *Survey 1* is summarized in Fig. 4(a), which reveals that humans perform better on imposter pairs than on genuine pairs. This observation is also evident in Fig. 4(b), displayed using an alternative representation. Here, it can be seen that both accuracy and certainty are higher for imposter pairs than for genuine pairs. Additionally, a correlation factor of 0.44 demonstrates a moderate correlation between the accuracy and certainty of the participants’ answers. A notable dip at the 50% certainty level is visible in Fig. 4(a), which can be attributed to the rating slider design. The slider is initially centered for each question and must be moved by the user. We observe slightly more answers with certainties above 50%, particularly for imposter pairs, resulting in a total mean certainty of 59.66% (std=14.61) and a mean accuracy of 67.25% (std=8.53) for *Survey 1*.

We also assessed the average user certainty for true positives (56.73%, std=27.74), false positives (55.63%, std=25.73), true negatives (66.27%, std=28.94), and false negatives (52.83%, std=30.33) image pairs. The highest certainty is achieved for true negatives, consistent with the results in Fig. 4(b). Fig. 4(c) illustrates the mean accuracies for each machine-certainty bin. We compare human performance with the mean of three machine models (*cf.* Subsec. 3.1). As expected, the machine ensemble outperforms human operators, primarily because we selected image pairs covering the entire range of machine certainties, acknowledging that machines excel over humans in most cases. The number of image pairs assessed within each bin increases with certainty and accuracy.

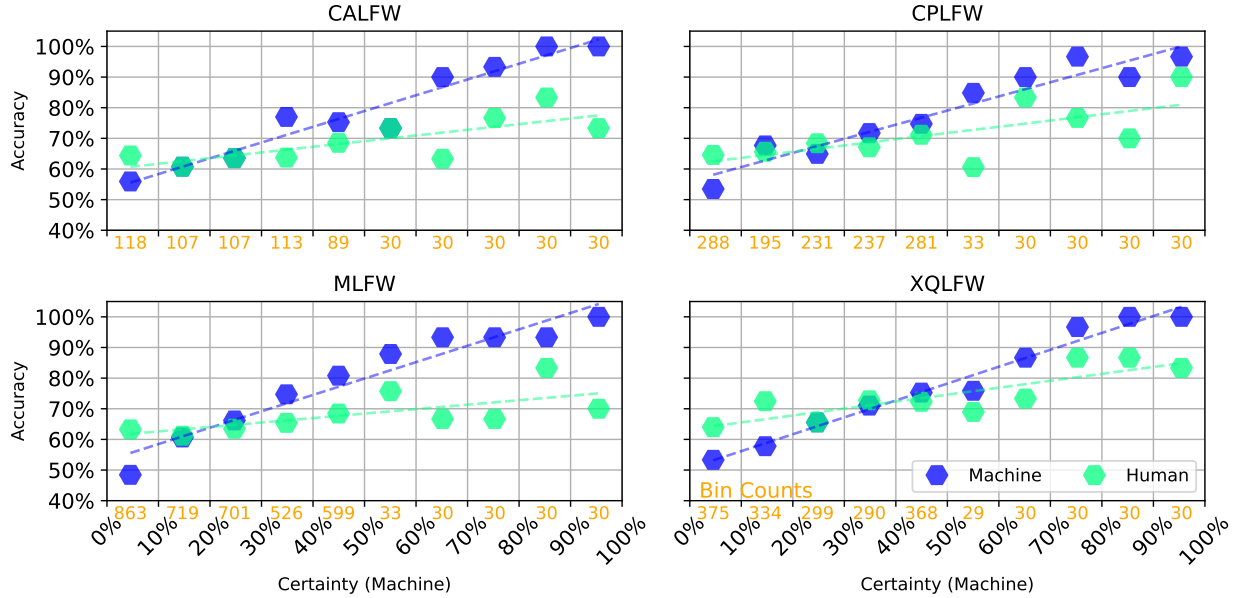


Figure 5: Accuracy comparison for machine certainty histogram bins across the datasets.

Interestingly, the decision duration (mean=10.71 sec, std=8.38) varied significantly ( $p=0.00003$ ) based on the user’s prediction. Users took an average of 10.22 (std=8.12) sec for the “different” answer and 11.66 (std=8.78) sec for the “same” answer. Investigating the precise reasons behind this finding would be an intriguing endeavor.

Lastly, accuracy analysis concerning participants’ ages revealed no significant differences.

#### 4.4 Experiment 2

Fig. 5 demonstrates that the mean accuracies for humans are higher than those for the machine ensemble when the machine ensemble certainties are very low (less than 10%) in *Survey 2–5*. This supports our fundamental assumption and motivation for our fusion approach (see Subsec. 3.3). The accuracy for both human and machine ensembles increases with rising certainty. The number of human decisions increases as the certainty score decreases.

#### 4.5 Human-Machine Fusion

In Table 2, we provide an overview of the performance of the three state-of-the-art models utilized in this study, combined with our fusion technique, on the four face verification benchmark datasets. The parameters  $t_{\min}$ ,  $t_{\max}$ , and  $d_{\min}$  were set to 0.37, 0.11, and 0.05, respectively. The results demonstrate that our fusion approach surpasses the

Table 2: Face verification accuracy [%] comparison and number of human decisions in fusion-algorithm.

| Model             | Dataset       |               |                |               |
|-------------------|---------------|---------------|----------------|---------------|
|                   | CALFW [4]     | CPLFW [3]     | MLFW [5]       | XQLFW [15]    |
| ArcFace* [19]     | 93.85         | 88.37         | 73.53          | 93.27         |
| + Human           | 94.03 (+0.18) | 88.83 (+0.47) | 75.88 (+2.35)  | 93.65 (+0.38) |
| # Human Decisions | 241 (4.02 %)  | 470 (7.83 %)  | 1871 (31.18 %) | 257 (4.28 %)  |
| FaceTrans* [19]   | 94.93         | 91.58         | 85.63          | 95.12         |
| + Human           | 94.93 (+0.00) | 91.73 (+0.15) | 85.70 (+0.07)  | 95.28 (+0.17) |
| # Human Decisions | 187 (3.12 %)  | 311 (5.18 %)  | 778 (12.97 %)  | 172 (2.87 %)  |
| ProdPoly [20]     | 96.03         | 92.75         | 91.30          | 86.90         |
| + Human           | 96.13 (+0.10) | 92.82 (+0.07) | 91.35 (+0.05)  | 88.05 (+1.15) |
| # Human Decisions | 17 (0.28 %)   | 125 (2.08 %)  | 325 (5.42 %)   | 605 (10.08 %) |

individual models on all datasets. Furthermore, it is apparent that the lower the model’s accuracy, the more human decisions are required to enhance the overall performance of the fusion approach.

#### 4.6 Limitations

As mentioned in Subsec. 4.2, we recruited only 60 participants, primarily of European/Caucasian ethnic background, which may limit the generalizability of our study [23]. Our study encompasses a combination of familiar and unfamiliar face matching, which could also influence generalizability. [24].

## 5 Conclusion

Having initially analyzed the distribution of certainty for machine models across various datasets, we constructed five surveys from these datasets and evaluated them with 60 participants. The key findings of the results are that humans perform better on imposter image pairs than on genuine pairs. Additionally, users have very high certainty and provide the correct answer for some pairs where the machine has low certainty. Ultimately, our fusion algorithm demonstrates that combining machine and human models can further improve machine accuracies.

In the future, we plan to extend our work by calibrating user certainties to further enhance human-machine fusion accuracy. We also intend to investigate the effect of presenting face verification explanation maps from machine models to users.

## References

- [1] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” in *Workshop on faces in ‘Real-Life’ Images: detection, alignment, and recognition*, 2008.
- [2] Martin Knoche, Stefan Hörmann, and Gerhard Rigoll, “Susceptibility to image resolution in face recognition and training strategies to enhance robustness,” *Leibniz Transactions on Embedded Systems*, vol. 8, no. 1, pp. 01–1, 2022.
- [3] T. Zheng and W. Deng, “Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments,” Tech. Rep. 18-01, Beijing University of Posts and Telecommunications, February 2018.
- [4] Tianyue Zheng, Weihong Deng, and Jiani Hu, “Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments,” *CoRR*, vol. abs/1708.08197, 2017.
- [5] Chengrui Wang, Han Fang, Yaoyao Zhong, and Weihong Deng, “Mlfw: A database for face recognition on masked faces,” in *Biometric Recognition: 16th Chinese Conference, CCBR 2022, Beijing, China, November 11–13, 2022, Proceedings*. Springer, 2022, pp. 180–188.
- [6] Xueyi Zheng, Ruixuan Wang, Xinke Zhang, Yan Sun, Haohuan Zhang, Zihan Zhao, Yuanhang Zheng, Jing Luo, Jiangyu Zhang, Hongmei Wu, et al., “A deep learning model and human-machine fusion for prediction of ebv-associated gastric cancer from histopathology,” *Nature communications*, vol. 13, no. 1, pp. 2790, 2022.
- [7] Daniel J Carragher and Peter JB Hancock, “Simulated automated facial recognition systems as decision-aids in forensic face matching tasks.,” *Journal of Experimental Psychology: General*, 2022.
- [8] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar, “Attribute and simile classifiers for face verification,” in *2009 IEEE 12th international conference on computer vision*. IEEE, 2009, pp. 365–372.
- [9] John J Howard, Laura R Rabbitt, and Yevgeniy B Sirotnin, “Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making,” *Plos one*, vol. 15, no. 8, pp. e0237855, 2020.
- [10] Matthew C Fysh and Markus Bindemann, “Human–computer interaction in face matching,” *Cognitive science*, vol. 42, no. 5, pp. 1714–1732, 2018.
- [11] Adam Sandford and Kay L Ritchie, “Unfamiliar face matching, within-person variability, and multiple-image arrays,” *Visual Cognition*, vol. 29, no. 3, pp. 143–157, 2021.
- [12] David White, Daniel Guilbert, Victor PL Varela, Rob Jenkins, and A Mike Burton, “Gfmt2: A psychometric measure of face matching ability,” *Behavior research methods*, vol. 54, no. 1, pp. 252–260, 2022.
- [13] Mirta Stantic, Rebecca Brewer, Bradley Duchaine, Mi-chael J Banissy, Sarah Bate, Tirta Susilo, Caroline Catmur, and Geoffrey Bird, “The oxford face matching test: A non-biased test of the full range of individual differences in face perception,” *Behavior Research Methods*, vol. 54, no. 1, pp. 158–173, 2022.
- [14] Naphtali Abudarham, Idan Grosbard, and Galit Yovel, “Face recognition depends on specialized mechanisms tuned to view-invariant facial features: Insights from deep neural networks optimized for face or object recognition,” *Cognitive science*, vol. 45, no. 9, pp. e13031, 2021.

- [15] Martin Knoche, Stefan Hormann, and Gerhard Rigoll, “Cross-quality lfw: A database for analyzing cross-resolution image face recognition in unconstrained environments,” in *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*. IEEE, 2021, pp. 1–5.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [17] Jiankang Deng, Jia Guo, Jing Yang, Niannan Xue, Irene Kotsia, and Stefanos Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 5962–5979, 2022.
- [18] Yaoyao Zhong and Weihong Deng, “Face transformer for recognition,” *CoRR*, vol. abs/2103.14803, 2021.
- [19] Martin Knoche, Mohamed Elkadeem, Stefan Hörmann, and Gerhard Rigoll, “Octuplet loss: Make face recognition robust to image resolution,” in *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*. IEEE, 2023, pp. 1–8.
- [20] Grigorios G Chrysos, Stylianos Moschoglou, Giorgos Bouritsas, Yannis Panagakis, Jiankang Deng, and Stefanos Zafeiriou, “P-nets: Deep polynomial neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7325–7335.
- [21] Martin Knoche, Torben Teepe, Stefan Hörmann, and Gerhard Rigoll, “Explainable model-agnostic similarity and confidence in face verification,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 711–718.
- [22] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang, “Racial faces in the wild: Reducing racial bias by information maximization adaptation network,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 692–702.
- [23] P Jonathon Phillips and Alice J O’toole, “Comparison of human and computer performance across face recognition experiments,” *Image and Vision Computing*, vol. 32, no. 1, pp. 74–85, 2014.
- [24] Robert A Johnston and Andrew J Edmonds, “Familiar and unfamiliar face recognition: A review,” *Memory*, vol. 17, no. 5, pp. 577–596, 2009.