

Received January 26, 2022, accepted March 11, 2022, date of publication March 17, 2022, date of current version March 24, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3160596

Psychophysical Evaluation of Human Performance in Detecting Digital Face Image Manipulations

ROBERT NICHOLS^{ID}, CHRISTIAN RATHGEB^{ID}, PAWEŁ DROZDOWSKI^{ID},
AND CHRISTOPH BUSCH^{ID}, (Senior Member, IEEE)

da/sec—Biometrics and Internet Security Research Group, Hochschule Darmstadt, 64295 Darmstadt, Germany

Corresponding author: Robert Nichols (robert.nichols@h-da.de)

This work was supported in part by the German Federal Ministry of Education and Research; and in part by the Hessian State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

ABSTRACT In recent years, increasing deployment of face recognition technology in security-critical settings, such as border control or law enforcement, has led to considerable interest in the vulnerability of face recognition systems to attacks utilising legitimate documents, which are issued on the basis of digitally manipulated face images. As automated manipulation and attack detection remains a challenging task, conventional processes with human inspectors performing identity verification remain indispensable. These circumstances merit a closer investigation of human capabilities in detecting manipulated face images, as previous work in this field is sparse and often concentrated only on specific scenarios and biometric characteristics. This work introduces a web-based, remote visual discrimination experiment on the basis of principles adopted from the field of psychophysics and subsequently discusses interdisciplinary opportunities with the aim of examining human proficiency in detecting different types of digitally manipulated face images, specifically face swapping, morphing, and retouching. In addition to analysing appropriate performance measures, a possible metric of detectability is explored. Experimental data of 306 probands indicate that detection performance is widely distributed across the population and detection of certain types of face image manipulations is much more challenging than others.

INDEX TERMS Digital forensics, face manipulation, human performance, manipulation detection, presentation attacks, psychophysical evaluation, spoofing.

I. INTRODUCTION

Face recognition is one of the most developed biometric technologies with endeavors dating back as far as 1966 [1]. In addition to algorithms, we often rely on the human ability to compare pairs of unfamiliar faces to verify identities, often in public security settings, such as border control. However, many studies have shown this to be a difficult task [2] with individual performance to be widely distributed [3], regardless of professional training [4].

This circumstance underlines vulnerabilities regarding identity verification processes, which are further exacerbated by a rather recent type of fraud which incorporates digitally manipulated face images to fraudulently obtain

genuine (FOG) documents. This has led to growing interest in detection of manipulated images that are submitted in the application procedure for a new identity document - a task human inspectors are traditionally expected to perform with ease, even in difficult visual conditions. Nonetheless, initial research suggests that human examiners frequently exhibit poor performance in detecting certain types of face image manipulations [5]–[7].

In light of the potential security implications, face image manipulations present a serious threat in these operational contexts, thus warranting further research to better understand the process of detecting manipulated face images and human ability to do so. To contribute to a better understanding of this topic, we propose an interdisciplinary approach to testing and evaluating human performance to address limitations of previous works and incorporating findings within the

The associate editor coordinating the review of this manuscript and approving it for publication was Vincenzo Conti^{ID}.

vast body of research in the field of psychophysics [8]–[10]. Specifically, the main contributions of this article are as follows:

- A comprehensive overview and literature survey of works that have evaluated human performance in detecting face image manipulations.
- Proposal of a methodological approach incorporating two well-established test procedures in the area of sensory science for this visual detection task, which better addresses limitations such as response bias and allow for a more robust evaluation of human performance. Furthermore, a computational measure to estimate task difficulty is used to facilitate interdisciplinary compatibility.
- An evaluation of human performance based on an expanded perceptual experiment incorporating multiple types and methods of face image manipulations to examine the capability of both novice and experienced individuals to detect these manipulations.

Our work is organized as follows. Relevant background information is outlined in Section II including face image manipulation types and processes, as well as a summary of basics and relevant methods in the area of visual psychophysics. In Section III, related work regarding human detection performance of manipulated face images is comprehensively summarized. Section IV introduces our experimental approach and presents results thereof, before we conclude our paper with an outlook in Section V.

II. BACKGROUND

Numerous everyday situations arise in both private and public sectors where identity and possibly corresponding characteristics (e.g. age, place of residence) must be verified to comply with local laws and ordinances, e.g. for gaining entry to a site or purchase of restricted goods. To that end, requesting a form of photo ID, confirming requirements and verifying legitimacy by comparing the submitted photo to the candidates' appearance is commonplace. Another critical example where the human ability to recognize faces is relied upon are checks on cross-border traffic as a matter of national security. The task of comparing a passport reference image with a trusted probe image in a border control scenario is typically performed by border guards and, more and more regularly, by additional automated border control (ABC) systems. In a routine border crossing setting, a border guard inspects the presented travel document and compares the hard copy photograph or the digital image extracted from the RFID chip to the appearance of the traveler in question, essentially comparing an unfamiliar face to a presented photograph. Studies [2], [11] have consistently shown individual performance in this task to be widely distributed across the population. Given practical and theoretical implications, determining the source of these differences has become a relevant issue, in turn prompting development of standardized tests such as the




FIGURE 1. Examples of bona fide source images and resulting manipulated image on the right. (a) Face swapping (b) Morphing (c) Retouching (Original image sources: Hochschule Darmstadt, and FERET [image publication permitted under fair use policy], and FRGC v2.0 [image publication permitted under fair use policy]).

Glasgow Face Matching Test (GMFT [2]) and the Cambridge Face Perception Test (CFPT [12]), amongst others, which have enabled the investigation of relationships within individual differences in performance across distinct recognition tasks regarding unfamiliar faces [3]. An attack leveraging digitally manipulated face images assumes the presentation of fraudulently obtained genuine (FOG) documents to a border guard as part of the biometric verification process, highlighting the role of humans in detecting manipulations, and thus preventing illicit border crossings. As such attack attempts are assumed to be rare in relation to the number of legitimate transactions [13], the ability of professional inspectors and border guards to reliably detect face image manipulations remains a core interest, as does which role, if any, experience has in this context.

A. DIGITAL FACE IMAGE MANIPULATIONS

Traditionally, digital face editing as an image processing task has received high levels of professional and private

interest, most commonly in the context of advertisement or entertainment. As such, it has recently become the focus of societal concerns, such as the spread of misinformation via social media [14]. Broadly speaking, face image manipulation involves the modification of specific facial attributes (e.g. age, sex, or hair), adding inconspicuous, yet disruptive artifacts, and swapping or morphing two distinct face images [15]–[17]. Other manipulations include modification of facial expressions as well as the addition of decorative components such as glasses or tattoos and even generating fully fictitious faces [18].

As automated photo and video editing tools such as deep learning-based face image manipulation software are publicly available and increasingly popular [19], [20], the looming potential for attacks based on digitally manipulated face images has recently prompted first regulatory mitigation attempts by procedural changes¹ in addition to increased efforts in the field of traditional media forensics [22], [23]. In such a scenario, the attacker utilizes digital image manipulation techniques to either distort his own facial image to avoid identification (identity concealer) or combines two (or more) face images, aiming to successfully purport an identity other than his own (impostor). Subsequently, the attacker attempts to acquire legitimate documents with the manipulated image, which would go undetected due to the perceived similar appearance of the attacker and resulting image, ultimately enabled by the circumstance that many countries do not require the photo to be taken on-premise [24]. An abundance of digital image manipulation techniques are available with varying levels of both required interaction and resulting quality, of which we briefly introduce three prominent examples as shown below.

Face swapping describes an image processing task with the goal of transferring the identity of a source face image onto a target face image while preserving target attributes. There are many approaches to accomplishing this task, ranging from manual image editing (e.g. Adobe PhotoshopTM) to automated deep learning-based techniques, for example by changing identity information within encoded features of the target image towards identity information from source image features [25]. An example of such an automatic approach for this technique is shown in Figure 1a.

Morphing is a multi-step process of combining two or more face images to one, which represents each contributing face to a degree relative to the respective contributions to the result [24]. An example with two face images is shown in Figure 1b. In principle, this is accomplished by applying generalized warping techniques and subsequent blending. In terms of required interaction, approaches range from manual [26], [27] or semi-automated [28] to

fully automated morphing processes [29]. Often an additional post-processing step (e.g. correcting contrast and brightness, manual adjustments) is carried out owing to visual artifacts in the images from previous steps.

Retouching refers to a diverse set of methods of face image processing with the general goal of improving visual appearance, a task commonly performed in commercial settings such as the fashion or advertisement industry. Aside from semi-automatic methods, automatic methods which often rely on machine learning are increasingly common, with approaches ranging from skin smoothing to reduce visible imperfections [30], digital make-up application [31], to methods for reversing age effects [32]. See Figure 1c for an example of automatic attractiveness improvement by, e.g. a sharpened jawline.

Recognition of face identity is critical to many everyday contexts, such as eyewitness identifications and inspection of identity documents or passports. In the past, there have been numerous reported instances of identity falsification [33]–[35], demonstrating a willingness of bad actors to engage in the general practice of claiming an alternative identity, also known as *identity spoofing*. In the context of biometric systems, spoofing refers to a type of a presentation attack, a term describing an external attack on the biometric data capture subsystem by subverting intended operation, where the subject aims to be identified as another individual by means of presenting manipulated input to the biometric sensor [36]. With more and more extensive deployment of face recognition technology in security-critical areas, specifically for applications in law enforcement and border control [37]–[39], face image manipulations present a serious threat in these operational contexts and to face recognition processes in general. This in turn has prompted sustained interest and work relating to detection methods for certain face image manipulations [40], in addition to aforementioned regulatory measures [21].

B. VISUAL PSYCHOPHYSICS

Psychophysics is concerned with relating quantifiable physical characteristics, referred to as a stimulus, to the perceptual sensation evoked in humans when they are subjected to said stimulus. Today, psychophysics mainly focuses on four perceptual primitives, namely detection, discrimination, identification, and scaling [41]. For our work, we focus on the detection and discrimination primitives, i.e. detecting if a stimulus is present or not and the magnitude of perceptual differences. Thresholds in the psychophysical context describe a specific stimulus intensity, which is required for detection of stimulus presence [9]. Scaling, on the other hand, describes the differential relation of intensities of two or more given stimuli [9].

¹Germany bans digital doppelgänger passport photos, Reuters (2020) [21].

1) PSYCHOPHYSICAL METHODS

Psychophysical methods define stimulus positioning, i.e. in which order stimuli are presented and calculation rules pertaining to the threshold or a different parameter, which can be derived from subjects' responses. Generally, these methods can be assigned to two main categories, the classic psychophysical methods, which were first developed by Fechner [42], and the more recent, adaptive methods [9], [43]. Below, we briefly outline three main classic methods, and broadly describe adaptive methods:

Method of constant stimuli In this method, stimuli of varied intensities are presented in random order, so that expectation errors are eliminated. The experimenter first selects a set of stimuli (e.g. five to nine) that the subject must judge in random order. The range of intensities is chosen so that the lowest intensities are never correctly recognized while the highest intensities are always recognized. After each presentation, the subject must state whether he perceived the respective stimulus or not. Thus, a psychometric function (see Section II-B3) can be collected empirically [44]. Ehrenstein and Ehrenstein emphasize that *"although the method of constant stimuli is assumed to provide the most reliable threshold estimates, its major drawback is that it is rather time-consuming and requires a patient, attentive observer because of the many trials required"* [44]. A further drawback of this method is the fact that as *"only stimuli near the threshold provide relevant information, many of the stimuli presented are too far away from the threshold to be of use."* [44]

Method of limits In this method, the stimulus is adjusted up or down so that it is just before being perceptible or scarcely past being perceptible. The subjects must indicate the point at which they first perceive, or no longer perceive, the stimulus. In an ascending sequence, the stimulus is initially set weakly, which means the subject cannot yet perceive it, and this is increased until the subject perceives the stimulus. In a descending sequence, a high stimulus strength is set at the beginning. This is reduced until the stimulus can no longer be perceived by the subject [44]. Ehrenstein and Ehrenstein found that in the *"ascending and descending series of trials"* there are often *"slight but systematic differences in thresholds"* [44]. For this reason, the ascending and descending sequences are to be performed several times and alternately. Finally, the recorded stimulus strengths of all runs are averaged to obtain the threshold estimate [44]. Zwisler proposes to include so-called catch trials, interspersed in the experiment, where no stimulus is presented. The attained frequency of perceived stimuli in such trials

could be used as a corrective measure for certain response inclinations [45].

Method of adjustment In this method, the stimulus is controlled by the subject and can vary continuously. Ehrenstein and Ehrenstein believe that *"the simplest and quickest way to determine absolute and difference thresholds is to let a subject adjust just the stimulus intensity until it is just noticed or until it becomes just unnoticeable (in the case of measurements of the absolute threshold) or appears to be just noticeably different from, or to just match, some other standard stimulus (to measure a difference threshold)."* [44]. The experimenter controls stimulus presentation and monitors the level of magnitude set by the subject. For example, a controller can be used to adjust the intensity of a sound. The subject has the task of adjusting the stimulus until it is at the threshold of perception [44]. Meanwhile, *"the stimulus intensity is recorded to provide an estimate of the observer's threshold."* [44]. The mean of the settings serves as an estimate for the absolute threshold [44]. Some advantages of the method of adjustment are that it is more rapid than the method of limits and thus easier to use, but it may have more limited applicability because it estimates only certain points on the psychometric curve, and habituation and expectation errors are still present [46].

Adaptive methods In contrast to the previous three classic methods, for adaptive testing methods, the goal is to keep stimulus intensity close to the threshold level. They are designed so that the selection of comparison stimuli is dependent on the response behavior of the assessor. This allows for shorter procedures and thus, makes adaptive procedures particularly efficient as the authors emphasize in [44]. Ehrenstein and Ehrenstein cite the staircase method as an example of such an adaptive method first used by Békésy [47]. This method is considered a modification of the threshold method. After each stimulus and the subject's response, the stimulus intensity is changed in the direction in which the stimulus threshold is assumed to be based on the response. If the subject's response is positive, the stimulus is increased; if the response is negative, the stimulus is decreased. This process is repeated until the step size reaches the preset minimum value [44]. Weber believes that stimulus strengths thus *"[...] approach the threshold stimulus (threshold convergence)"* and settle *"around the threshold"* after a few stimulus presentations [48].

2) RESPONSE BIAS

All psychophysical methods indirectly make the assumption that stimulus intensity is the sole determinant of the observer's judgment. It appears quite likely, however, that

there are other factors at play which lead to a distortion of responses yet are included in that judgment. This results in a so-called response bias. Such biases can arise, for example, because the observer does not behave cooperatively in the kind of elaborate threshold determination methods encountered in settings in use, or because the certainty of the judgments can be variable. There may be observers who jump in to answer “signal present” at a very low certainty level and others who will only give the same verdict at a high level [9]. This ultimately results in different response patterns and ultimately different psychometric functions.

3) PSYCHOMETRIC FUNCTION

An important illustration of the human sensory ability with regard to detection and discrimination is the psychometric function, which can be estimated from psychophysical data obtained by sampling observers’ performance on a given psychophysical task at different stimulus intensities. Such function, as described by [49] in Equation 1, correlates a stimulus x to the probabilities P_Ψ of detecting this stimulus, where α and β describe parameters for location and sensitivity regarding the base function f and relate to the threshold and subject sensitivity for such a detection experiment [49]. A common choice for the base function f is either a cumulative normal distribution or a logistic function. The parameters γ and λ represent positive and negative guessing error rates, for instance, when an observer indicates a positive detection when no stimulus is present and fails to detect a stimulus at a level above the detection threshold.

$$\Psi(x; \alpha, \beta, \gamma, \lambda) = \gamma + (1 - \gamma - \lambda) * f(x; \alpha, \beta) \quad (1)$$

4) TEST PROCEDURES

Another significant aspect to consider for development of psychophysical experiments is the choice of test procedure or trial design, as used synonymously hereafter, to refer to the task itself, i.e. the number of stimuli in each trial, how they are presented and methods for data collection. The small subset of procedures described in this section are performance-based, in that they aim to measure perceptual aptitude; in other words, how well one performs at a specific visual task. In this paper, we refer to forced-choice procedures in a broad sense as all procedures that require the observer to make a choice from two or more predefined response options. This must not necessarily be dependent on the number of stimuli present in a given trial, as single stimulus trials with two response choices are equally possible and thus considered forced-choice.

Yes/No This class of procedures covers many forms of tests and is commonly employed to determine detection thresholds. Standard practice is for half of all trials to include stimuli with the target signal present and half with the target absent (i.e. noise trials). The observers task is to respond “Yes” or “No” on each trial, dependent on the target being perceived or not. The order of signal and noise trials must be

random, in any case avoiding repetitious sequences of either trial type [9]. As a criterion-dependent procedure, Yes/No tasks are exceptionally prone to effects of response bias [8], as previously described. Therefore, the proportion of correct responses is often an inferior measure of observer sensitivity and the signal detection measure d' is generally the metric of choice for analysis [9].

2AFC The paired comparison, in this case specifically a two-alternative forced-choice (2AFC) task, is considered one of the most popular procedure designs in psychophysics [9]. In this task, observers are presented with two stimuli A & B, where the target signal can be either of both, and subsequently required to decide which represents the target signal. In the form we commonly refer to in this paper, both stimuli are presented side-by-side simultaneously (i.e. spatial arrangement); likewise, a form in which stimuli are displayed in temporal order is possible but is outside the scope of this work. This task is appropriate when measuring specific differences between two samples in one specific known attribute, with results providing directional information of the specified difference [43]. As a comparatively simple and intuitive task, sensory fatigue and carryover effects are reduced and the effects of response bias are considered to be minimal, given randomized balanced serving sequences. Limitations, however, can arise both when differences extend beyond the single specified sensory dimension, which could affect other sensory attributes, and when observers become frustrated due to a large number of required trials to fully sample the psychometric function [50].

ABX This discrimination task is a sequential match-to-sample test, where observers are, on each trial, initially presented with two reference samples A & B, both target signal present and target absent. After the inspection phase, a third sample X is presented, which can be of type A or B and the observer is asked to indicate whether the sample X is a positive match to type A or B. As a discrimination task, the main objective is to determine if a perceptual difference exists between two sample categories [43]. The main advantages of this procedure are that no prior knowledge of the stimuli is required and that measurement of general sensory differences is enabled by allowing observers to inspect both samples, and thus independently discover attributes of difference and weight them accordingly in subsequent decisions [51]. Conversely, this could be a drawback as well, as the difference is not specified and observers might be influenced by overt, yet random and irrelevant attributes [43]. Other possible disadvantages could stem from the inherent reliance on memory,

as opposed to strictly perceptual representations of difference and a potentially strong bias effect towards “ $X = B$ ” in categorical perception, when reference samples are presented sequentially during the inspection phase.

III. RELATED WORK

As the present topic is a fairly recent subject of investigation, literature is relatively sparse. Nevertheless, initial studies have indicated that humans display inadequate proficiency in detecting certain manipulated face images. However, the existing research still has problems in terms of chosen approaches, trial design and appropriate performance measures. See Table 1 for an overview of relevant studies reviewed in this section.

The face morphing attack is assumed to have the highest security impact, mainly due to the possibility of FOG documents which would sufficiently resemble two individuals and thus subvert biometric authentication and achieve positive verification for both individuals. Therefore, human detection performance has been previously assessed almost exclusively for this type of face image manipulation. However, given the rapid development of technology at the disposal of potential adversaries, other possible scenarios and approaches should be considered, for instance so-called deepfakes and face swapping (i.e. replacing identifying information in an image with another person’s depiction) or beautification and digital make-up application. One such example is when an individual aims to avoid automatic verification altogether by means of removing identifying information from the image. This causes the process to fail and forces potentially unreliable human verification. With increasing availability, capability and consequently popularity of specialized software for face image manipulation, attacks utilizing alternate manipulation techniques become more probable.

One of few studies examining alternative manipulation methods is [54], where the human ability to distinguish deepfake (i.e. face swapping) images from originals was assessed. The face images were sourced from YouTube videos and subsequently manipulated with both graphics-based and learning-based approaches, yielding detection performance levels below chance for more sophisticated manipulations. Another assessment of deepfakes was conducted in [56], where short videos rather than images were presented to participants. Additionally, the authors manually assigned stimuli to five categories, representing different levels of difficulty to detect. In line with the findings of [54], their results also showed that participants are easily deceived by a good quality deepfake video. Human performance in [55] was reported at chance level for detecting face images, which were automatically manipulated primarily to improve appearance or attractiveness. Stimuli were both manually and automatically generated, with the manual editing performed by professional artists being less subtle and thus yielding higher detection rates.

Two main approaches in the discussed works are predominant for determining morph detection performance. In [15] and for individual experiments conducted in [5], [7], [53], [57], [58], participants are tasked with comparing unfamiliar faces, given combinations of either mated, non-mated, and morphed face images, essentially determining if the two face images displayed represent the same person. Considering previous work on the unfamiliar face comparison task has consistently shown humans, including experienced individuals, to be unreliable [2], [4], [11], [60]–[62]; this approach possibly introduces an unintentional bias by including face recognition performance as a component in the results. Moreover, as the general relationship of performance in unfamiliar face comparison and performance in detecting manipulated face images remains unclear, the conclusiveness of such findings may be limited in this context. The other approach as seen in [6], [52], [54], [56], [59] and experiments in [7], [53], [57], [58] consists of a direct morph detection task comprised of bona fide and morphed face images. As this approach exclusively targets perceptual detection and discrimination, it is less affected by the aforementioned bias. Nonetheless, this approach requires very conscious decisions in terms of experiment design, as many factors can substantially influence outcome. In some cases, the chosen designs do not satisfy all requirements for the subsequent evaluation, and thus reliability of results may be limited.

In [6], an extended version of the two-out-of-five test procedure is used, which is an appropriate method for determining if a significant, yet unspecified difference exists between two categories of stimuli [63]. In this extended version, participants were asked to choose manipulated images out of 10 presented face images, of which 5 were manipulated, basically corresponding to a five-out-of-ten test. The main advantage of this sorting task is the low chance probability of 1/10 and 1/252 for the standard and extended versions respectively. In [64], however, the authors underline the importance of task instructions regarding this test design and reported that different terminology (e.g. *same pair* or *grouping*) can lead to greater proportion of incorrect responses. Taking this into account, when the previously described experiment was replicated with more sophisticated face image manipulations (i.e. harder to detect) in [7], the task instruction was changed from originally “Which of these images have been digitally manipulated?” to “Which of these images are face morphs (a blend of two faces)?”, which could be an additional factor in the reported performance reduction, despite participants’ receiving specialized training. Furthermore, if the instructions or guidance include detailed descriptions, suggestions and illustrations of visual cues as e.g. [7] described, it remains unclear if the perceptual component of detecting face manipulations is discounted and a shift towards the visual search component with an additional memory task takes place. Another factor potentially facilitating such a task shift is the absence of a time limit for stimulus presentation or inspection phases in most studies, with [55] representing one of few exceptions with a correct implementation of a spatial

TABLE 1. Overview of relevant studies of human performance in detecting manipulated face images.

Reference	Manipulations	Stimuli	Test class ⁱⁱ	Chance probability	Observers	No. trials	Performance
Ferrara <i>et al.</i> (2016) [15]	Morphing	10 Bona fide 20 Manipulated	Yes/No	N/A ⁱ	44 (Expert) 543 (Non-expert)	30	MAR 9.57-100% FRR 0-22.27%
Makrushin <i>et al.</i> (2017) [52]	Morphing	7 Bona fide 23 Manipulated	Yes/No	N/A ⁱ	42 (unspecified)	30	MAR 44.6% FRR 43.64% ACC 55.62%
Robertson <i>et al.</i> (2017) [5]	Morphing	49 Bona fide 245 Manipulated	Yes/No	N/A ⁱ	28 (Non-expert) 42 (Non-expert)	49	MAR _b 68% MAR _g 21%
Robertson <i>et al.</i> (2018) [6]	Morphing	49 Bona fide 245 Manipulated	Sorting	1/252	80 (Non-expert)	32	HR _b 48% FA _b 23-27% HR _t 79% FA _t 7% HR _g 71% FA _g 7%
			Forced-choice	1/2			ACC _{tt} 89%
Kramer <i>et al.</i> (2019) [7]	Morphing	127 Bona fide 245 Manipulated	Sorting	1/252	80 (Non-expert)	32	HR _b 38-42% FA _b 30-33% HR _t 41% FA _t 37% HR _g 40% FA _g 29%
			Forced-choice	1/2			ACC _{tt} 51.3%
			Yes/No	N/A ⁱ	40 (Non-expert) 1410 (Non-expert)	60 1	ACC 51.3-57.1% MAR 49%
Makrushin <i>et al.</i> (2019) [53]	Morphing	11 Bona fide 19 Manipulated	Yes/No	N/A ⁱ	49 (Expert) 230 (Non-expert)	30	FNR 15.41% TNR 82.35-83.27% MAR 34.65% TAR 77.3%
Rössler <i>et al.</i> (2019) [54]	Deepfake Face swapping	140 Bona fide 140 Manipulated	Yes/No	N/A ⁱ	204 (Non-expert)	60	HR 68.69% TNR 78.19%
Wang <i>et al.</i> (2019) [55]	Retouching	100 Bona fide 100 Manipulated	Forced-choice	1/2	40 (Non-expert)	35	ACC 53.5% ACC 71.1%
Korshunov <i>et al.</i> (2020) [56]	Deepfake	60 Bona fide 60 Manipulated	Yes/No	N/A ⁱ	60 (Non-expert)	40	TPR 24.5-71.1% TNR 82.2%
Makrushin <i>et al.</i> (2020) [57]	Morphing	15 Bona fide 20 Manipulated	Yes/No	N/A ⁱ	21 (Expert) 168 (Non-expert)	40	TPR 65.4% TNR 88.47% MAR 57.88% FRR 23.54%
Nightingale <i>et al.</i> (2021) [58]	Morphing	216 Bona fide 108 Manipulated	Yes/No	N/A ⁱ	500 (Non-expert)	108	ACC 54.1% (d' .21) ACC _t 60.4% (d'_t .53)
Zhang <i>et al.</i> (2021) [59]	Morphing	15 Bona fide 75 Manipulated	Yes/No	N/A ⁱ	19 (Expert) 52 (Non-expert)	90	HR 64.31-88.25% TNR 79.21-97.14% ACC 81-86%
Ours	Morphing Face swapping Retouching	59 Bona fide 64 Manipulated	Forced-choice Matching	1/2	83 (Expert) 144 (Non-expert)	50	ACC 75.43% (d' 1.07) ACC 62.92% (d' 1.02)

ACC: Classification accuracy (*percentage of correct answers*)HR Hit Rate, TPR True Positive Rate (*proportion manipulated images correctly classified as manipulated*)TNR: True Negative Rate ($1 - FA$)FA False Alarm Rate, FPR False Positive Rate (*proportion bona fide images falsely classified as manipulated*)MAR: Morph Acceptance Rate, FNR: False Negative Rate ($1 - TPR$; 50/50 morphs, if applicable)TAR: True Acceptance Rate (*proportion mated facial image pairs correctly classified as from the same source*)TRR: True Rejection Rate (*proportion non-mated facial image pairs correctly classified as different*)FAR: False Acceptance Rate (*proportion non-mated facial image pairs falsely classified as from the same source*)FRR: False Rejection Rate (*proportion mated facial image pairs falsely classified as different or bona fide images as manipulated*) b = baseline/pre-training performance t = post training performance tt = performance on training task g = post guidance/tips performanceⁱIn Yes/No testing, the individually set criterion may vary, resulting in a probability not equal to 1/2 [43]ⁱⁱsee e.g. [43]

2AFC test procedure, limiting stimulus presentation time to 6 seconds on each trial. Notably, [54] utilized time limits for stimulus inspection, although as intervals were chosen at random (2, 4, or 6 seconds) for each trial to simulate real-world conditions, this choice is not common in discrimination testing, arguably due to the possibility of inter- and even intra-individual biases. Most studies have generally followed test procedures of the “Yes/No” class of psychophysical tests [43]. The specific reported procedures vary across studies, both in type and degree of conformity with the recommended protocol. Despite this design being suitable to measure detection thresholds, the limitations inherent to this task type must be considered and thus evaluated accordingly.

For example, [15] implements the “Yes/No” one-alternative forced-choice (1AFC) “same-different” test procedure, in which observers are presented with combinations of the only two stimuli involved. Due to the asymmetry of the discriminands (‘same’ and ‘different’) this trial design is especially prone to response bias. In this case, this is a tendency for observers to respond ‘Same Subject’ or ‘Different Subject’, and therefore analyzing data from such trials with signal detection measures to account for any bias would be favorable [9]. Furthermore, an imbalanced number of presented stimulus combinations (10 ‘Same Subject’ and 20 ‘Different Subject’) constitutes another source of bias.

The other common “Yes/No” procedure seen in present studies is the 1AFC symmetric task with one stimulus per trial. Examples for this procedure can be seen in [53], [54], [56]–[59], where participants were presented with a face image on each trial, belonging to either of two discriminands (“real/fake” or “bona fide/manipulated”, respectively). A requirement for this procedure to be considered unbiased and subsequent evaluation by $\llproportion\ correct\gg$ as a performance metric as stated in [9], is for both discriminands to be “symmetric”. In terms of stimuli, this implies equivalence along all dimensions except for one, in which the difference shall be the respective opposite. This is trivial for simple visual stimuli (e.g. “arrow left” and “arrow right”); however, despite defining manipulations as the opposite of genuine, face images and manipulations thereof are complex, multidimensional visual stimuli which are not easily classified as opposite in absence of objective metrics to determine the degree of divergence. An argument could be made for classifying test procedures implemented in e.g. [52], [53] as the “A, not-A” (“Yes/No Response Choice”) test; however, to comply with any of the “A, not-A” test variations referenced in the literature, either a training phase or sequentially presented reference sample is required, and subsequent evaluations require different statistical models [43]. Further variations of this procedure were implemented in e.g. [5] as a composite, where observers had an additional alternative as a response option, effectively representing an aggregate trial design, and in [7] (see Experiment 3) as a follow-up trial after observers performed face comparison and a decision reflection phase. Problematically for the latter layout, observers completed multiple dependent trials

on the same stimuli, leading to results that are likely biased. In the instance of the former design, possible qualified models or corrective metrics of performance are undetermined.

Additionally, as for all “Yes/No” procedures, for minimal bias the discriminands should be presented with equal probability on each trial and as highlighted in [9], it is important for the observers to be aware of this. As proper implementation of this procedure is admittedly difficult and thus, rare, treatment of this procedure as free from bias and evaluation by $\llproportion\ correct\gg$ as a performance measure is inadequate and analysis of signal detection measures are preferable. Indeed, in the second experiment conducted in [7], the aforementioned procedure design was employed with equal occurrences of discriminands (i.e. 30 manipulated, 30 bona fide stimuli in 60 trials) and subsequent analysis of signal detection measures. Nevertheless, in light of questions regarding symmetry and inherent unknown intensities (i.e. detection difficulty) of manipulated stimuli, presenting each participant with a random sample of each 30 from 60 potential stimuli, rather than each participant assessing the same set of images, could exacerbate this issue. Furthermore, no indication was made whether participants were in fact aware of discriminands appearing with equal probability.

A relatively rare choice for studies discussed in this section is the “Forced choice” test class, which can be described as a directed comparison of one or more pairs of stimuli, thus often implemented as (spatial-) 2AFC or m -alternative forced-choice (m -AFC) procedures, respectively. This test method is appropriate for studies in which the two (or more) sample categories are distinguished by only one specific sensory attribute, which is known a priori. Given these prerequisites, this procedure is described as efficient and powerful [43], which speaks for this test design as an excellent choice for evaluating perceptual performance when specifying manipulation as the distinctive sensory attribute. As previously mentioned, [55] employed a spatial 2AFC test procedure with “image modification” as the discriminating sensory attribute. Furthermore, this design was applied to the morph detection training task in [6] and subsequently in [7]. As this test procedure is assumed to be less prone to response bias, typically $\llproportion\ correct\gg$ is an adequate measure of performance; yet, in certain cases (i.e. positional bias), signal detection measures might be more appropriate. Noteworthy, reported training task performance of 89% accuracy for less sophisticated morphs in [6] and 51.3% for high quality morphs in [7], is in line with findings in [55], where 71.1% accuracy is reported for more noticeable, manual manipulations (e.g. attractiveness) by professional artists, as opposed to 53.5% for intentionally subtle, automatic manipulation methods.

Many of the reviewed studies used performance measures based on the *proportion of correct* responses. As previously mentioned, with regard to the particular test procedure, this is a potentially unreliable measure for observer sensitivity to the target stimulus due to effects of bias or an individually set decision criterion. For the two-out-of-five test, [63] proposes

chi-squared or binomial statistics for analysis. Since [6] and [7] employed a modified form of this procedure with five stimuli with the target signal present and five without, it is possible such recommendations do not fully extend to these procedures. In contrast, a signal detection approach as pursued, thus calculating d' sensitivity and criterion location c for analysis of variance (ANOVA). This represents a valid course of action in the context of Yes/No tasks [8], [65], which was appropriately followed in the remaining experiments of [7]. However, results arising from alternate test procedures and corresponding theoretical models need to be interpreted with caution, as substitute models and thus corrections [66] and performance measures might prove more robust in terms of accuracy. Few studies use d' sensitivity and c criterion as measures of performance for analysis, despite implementing test procedures where these would be the measures of choice. For instance, in [58] experiments are of the Yes/No class, specifically both same-different and 1AFC symmetric test types. The authors report d' sensitivity and β for response bias in addition to $\langle\langle\textit{proportion correct}\rangle\rangle$. The reported low sensitivity for detecting single morphed images is noteworthy, and likely can be ascribed to the high quality and computational selection of similar individuals for morph generation. The measure β is a historically popular choice to represent response bias, which is based on the assumption that observers make decisions based on a likelihood ratio of obtaining a given value on either signal or noise trials. Nonetheless, arguments have been made that criterion location c represents a more plausible choice to measure response bias, as it presumes a response is directly based on the decision variable [65].

The main issues and challenges in the previously discussed studies pertaining to human detection performance can briefly be summarized as follows:

- *Approach*: evaluating human performance in either a single or differential detection scenario, therefore weighting inclusion of face comparison task performance in results against potentially less precise results.
- *Stimuli*: focus on one type of manipulation (morphing) in addition to unquantified differences in quality of manipulations (i.e. stimulus intensity), which greatly complicates comparison and reduces practical relevance.
- *Test procedure*: incomplete or erroneous implementation of test procedures, in some cases modifying well studied procedures, resulting in unknown effects on robustness of findings.
- *Performance measures*: an inadequate choice of performance measures for subsequent analysis w.r.t. the implemented test procedure may not take possible biases into account, and thus could impact reliability of findings.

IV. EXPERIMENTAL EVALUATION

In the following section, we elaborate on the details of the experiment by giving an overview of the experiment properties, trial structure and selected stimuli, define performance

measures, and briefly outline limitations in this regard. We then go on to report the results, putting our findings in perspective with regard to the context of the related works.

A. EXPERIMENTAL SETUP

The experiment is designed as a web-based survey that is available at <https://dasec.h-da.de/unclassifyd>. Any device with a modern web browser can be used to participate. After completing a brief registration with the possibility to self-report participant characteristics and following email confirmation to discourage improper use, access to the experiment is granted. The self-reported participant characteristics include age, gender, and relevant professional experience. The data used in this work was collected on a voluntary basis and in accordance with General Data Protection Regulation (GDPR) [67] provisions, thus providing required details for participants' informed consent during registration.

B. PERFORMANCE MEASURES

Previous studies, depending on either employing a face comparison task or detection task, mostly relied on general classification accuracies and standard performance measures of biometric systems, which often mirror common measures in sensory science. The variety of overlapping definitions make comparison across studies difficult; therefore our aim is to follow the recommended measures and models in the area of sensory science and note equivalencies with standard metrics of biometric performance where applicable. The comprehensive discussion in [68] points out various procedure designs may require different approaches for evaluation, and thus different measures or different d' calculations, depending on perceptual decision models.

As our work focuses on the detection of manipulations as a binary classification problem and purposefully abstains from face comparison tasks, standard metrics are employed to describe detection performance. Manipulated images are understood as positive samples (i.e. signal stimuli s) and bona fide images as negative samples (i.e. noise stimuli n). Consequently, we use the following performance measures:

Classification accuracy (ACC) is defined as the proportion of correct responses over all trials.

True positive rate (TPR) or hit rate (H) is defined as the proportion of signal stimuli (manipulated images) correctly classified as manipulated.

False positive rate (FPR) or false alarm rate (F) is the proportion of noise stimuli (bona fide images) falsely classified as manipulated, denoted as bona fide presentation classification error rate (BPCER) in the context of biometric presentation attack detection (PAD) performance [69].

False negative rate (FNR) and true negative rate (TNR) the reverse metrics of TPR and FPR, which can enable better comparison with previous studies. FNR is denoted as attack presentation classification

TABLE 2. Stimulus-response table relating possible stimulus sequences and possible observer responses in the 2AFC procedure.

		Response alternative	
		" $s = A$ "	" $s = B$ "
Stimulus sequence	$< s, n >$	Hits	(False alarms)
	$< n, s >$	False alarms	(Hits)

error rate (APCER) in the biometric PAD performance context [69].

1) 2AFC PROCEDURE

Data from the 2AFC procedure is represented in a 2×2 stimulus-response table as shown in Table 2. For each trial, the observer reports in which location the signal stimulus occurred, as both stimuli – signal and noise – are present in each trial. The two stimulus sequences, in the sense of spatial order, are denoted by angle brackets and correspond to the rows of the table. In this case, as stimuli are displayed side-by-side, $< s, n >$ represents a manipulated image in the "A" location to the left, and a bona fide image on the right ("B" location). For each row, the number of hits and false alarms can be understood as proportions of the total number of corresponding trials. Observers can either respond " $signal = A$ " ("A is manipulated") or " $signal = B$ " ("B is manipulated"); thus only one response column is necessary to determine remaining proportions of hits and false alarms, as the sum in each row is equal to the number of trials of the respective sequence. Hence, we arbitrarily define *hit* (H) and *false alarm* (F) rates as

$$\begin{aligned} H &= P("s = A" | < s, n >) \\ F &= P("s = A" | < n, s >), \end{aligned} \quad (2)$$

and denote this by parenthesis for the numbers in the " $signal = B$ " response column. Generally, if observers do not favor one or the other of the alternatives a priori in the spatial 2AFC procedure, the proportion of correct responses (i.e. classification accuracy *ACC*) can be considered a performance measure relatively unaffected by bias. Although rare, this design does not preclude individual observers adopting extreme decision strategies or responding to one display position more often than another; therefore we additionally calculate d' by taking into account the added difficulty compared to the Yes/No task discussed in [68], and obtain sensitivity with

$$d' = \frac{1}{\sqrt{2}}[z(H) - z(F)], \quad (3)$$

where $z(x)$ expresses the z-score, as in the relative location where the standard normal distribution is equal to x .

TABLE 3. Stimulus-response table relating possible target stimulus occurrence and observer responses in the ABX procedure.

		Response alternative	
		"S"	"N"
Stimulus alternative	$X = s$	$P(S s)$ Hit	$P(N s)$ Miss
	$X = n$	$P(S n)$ False alarm	$P(N n)$ Correct rejection

2) ABX PROCEDURE

Similarly, data arising from the ABX procedure design is represented in Table 3. Observers for each trial report if they believe the stimulus ("X") is a signal or noise stimulus after viewing references of both types in temporal order. In this case, the table rows correspond to both signal and noise trials with manipulated and bona fide target stimuli, respectively. Possible responses include " $X = signal$ " (or " X is manipulated") or " $X = noise$ " (or " X is bona fide"), thus allowing us to define *hit* (H) and *false alarm* (F) rates in this context as

$$\begin{aligned} H &= P(S|s) \\ F &= P(S|n). \end{aligned} \quad (4)$$

Our experiment includes samples of different types of face image manipulations and each is of varying subjective quality. In particular, each manipulation category features samples which are both more and less obvious to spot. For this reason, we must assume that any given observer may continuously adopt different thresholds in view of the level of sensory evidence required for a positive decision. Therefore, we conduct a signal detection analysis as recommended in [70], applying the *differencing strategy*, which [71] indicated is likely a suitable model for ABX experiments. For our experiment, rather than looking up corresponding values for d' in tables as included in MacMillan and Creelman's book [68], we calculate d' directly by the outlined method for differencing. First, we find the largest value of $\ll proportion correct \gg$ for constant d' , where responding is unbiased under the assumption of equal variance, from

$$P(C)_{max} = \Phi \left[\frac{z(H) - z(F)}{2} \right], \quad (5)$$

where Φ represents the standard normal cumulative distribution function. We can then find d' by inserting the previously obtained $P(C)_{max}$ into

$$P(C)_{ABX} = \Phi \left(\frac{d'}{\sqrt{2}} \right) \Phi \left(\frac{d'}{\sqrt{6}} \right) + \Phi \left(\frac{-d'}{\sqrt{2}} \right) \Phi \left(\frac{-d'}{\sqrt{6}} \right) \quad (6)$$

and solving for d' . In line with [68], bias was calculated as the observer criterion location c , which is defined as

$$c = \frac{z(H) + z(F)}{-2}. \quad (7)$$




FIGURE 2. Example of displayed decision feedback on a 2AFC trial (Original image sources: FERET [image publication permitted under fair use policy]).

C. PROCEDURE

Each trial consisted of four components:

- 1) a textual task description
- 2) a stimulus inspection phase
- 3) a decision or response step
- 4) feedback display.

Firstly, participants were made aware of the current trial procedure and objective, before inspecting the respective stimuli. Afterwards, participants were asked to indicate their assessment and to rank how confident they were in their decision on a 5-point scale from 0 (unsure) to 4 (certain). Feedback was given immediately after each trial decision to improve task comprehension and maintain motivation, while detailed study of the stimuli during feedback was limited by a reduced display duration before progressing to the next trial. An example is given in Figure 2.

On each 2AFC trial, both stimulus alternatives were presented in a random spatial order, designated with labels *A* and *B*. Participants were instructed to specify the position label of the manipulated image. Similarly, on ABX trials, two reference stimuli designated with corresponding labels *Bona fide* and *Manipulated* were displayed in sequential (temporal) order, followed by the target stimulus, specified as *X*. Participants were instructed to specify if the last image (*X*) was either manipulated or bona fide. In visual psychophysics stimulus presentation is typically focused on very short intervals, e.g. 50 to 1000 milliseconds [72], [73]. Faces in particular represent a group of stimuli which have been found to be processed very quickly; in terms of identity representations as fast as 30 milliseconds [74], [75]. Nevertheless, to date, detection of manipulated face images represents a novel task in a psychophysical context. We therefore chose comparatively long intervals to account for presumed greater stimulus complexity, and thus greater task difficulty.



FIGURE 3. Interactive example trial explaining the stimulus presentation phase with placeholder stimuli to ensure participants understood and were prepared for the task at hand.

A 90-second time limit was chosen for the task description, although participants had the option to manually proceed to stimulus inspection before the timeout had elapsed. Stimuli were displayed for 8 seconds on spatial 2AFC trials and for 6 seconds each for sequential ABX trials. Participants were given 60 seconds to respond; failure to submit their assessment within the given time resulted in a nondecision and automatic progression to the next trial. Feedback display was limited to 10 seconds for previously stated reasons. Participants each completed 23 ABX and 27 2AFC trials for a total of 50 experiment trials and two preceding instruction screens with optional interactive example trials, as illustrated in Figure 3. Note, the proportion of 2AFC trials was greater owing to our goal of collecting a broad sample of all stimulus intensities. As the example trials purpose was strictly to clarify tasks and methodology to minimize participant confusion, only stimulus placeholders were presented in these example trials. The experiment was based on the method of constant stimuli, a choice made in light of the advantages, despite possible drawbacks, as described in Section II-B, particularly referring to accuracy and the relative procedural simplicity. As such, all trials were completed in pseudorandom order. Textual instructions and timeout parameters were uniform across all trials of the same type. Upon completion of the procedure, participants were thanked, and overall performance feedback was displayed.

D. PARTICIPANTS

A total of 306 participants of different ages, genders, and levels of professional experience were recruited in an online campaign spanning multiple predominantly security and biometrics-focused email newsletters and social media. Based on information collected during the experimental phase, exclusion criteria was established to avoid distortion of results, as intermittent unresponsiveness on trials and delayed reactions indicated inattentive participants. Individual results were excluded when one or more of the following conditions were met:




FIGURE 4. Participant demographics: Age distribution.




FIGURE 5. Participant demographics: Gender distribution.

- The procedure was not completed within 6 hours
- Less than 23 and 27 responses were recorded for ABX and 2AFC trials, respectively
- The participant previously completed the experiment

Of 306 registered participants, 249 completed the full experiment. Furthermore, 22 results failed to meet remaining criteria, allowing only 227 participant results to be included in our evaluation. Two independent sample t-tests were performed to compare age and gender differences between novice (experience levels “None” through “Intermediate”) and expert (levels “Expert” through “Specialized professional”) groups. There was both a significant difference in age between novice ($N = 144$, $M = 2.13$, $SD = .998$) and expert ($N = 83$, $M = 2.71$, $SD = .957$) participants $t(225) = -4.272$, $p = .00$; as well as in gender $t(225) = 2.121$, $p = .035$. An overview of participant demographics is given in Figures 4 and 5, and characteristics related to experience in Figures 6 and 7.

E. STIMULI

Stimulus facial images were derived from the FERET [76] and FRGC v.2.0 [77] databases. The utilized in-house database of manipulated images consisted of different automatically generated manipulations of selected images and image pairs. For each face manipulation (i.e. face swapping, morphing, retouching) two methods were used: FewShotFace² and SimpleFS,³ FaceFusion⁴ and UBO

²<https://github.com/shaoanlu/fewshot-face-translation-GAN>

³https://github.com/Jacen789/simple_faceswap

⁴<https://www.wearemoment.com/FaceFusion/>

TABLE 4. Average ArcFace distances for each manipulation type and difficulty, used as an approximation of stimulus intensity or difficulty.

Stimulus class			Avg. ArcFace score
Face swapping	hard	Fewshotface	0.37
	easy	SimpleFS	0.14
Morphing	hard	FaceFusion	0.55
		UBO	0.56
	easy	FaceFusion	0.27
		UBO	0.28
Retouching	hard	Instabeauty	0.60
		Fotorus	0.72
	easy	Instabeauty	0.45
		Fotorus	0.41




FIGURE 6. Participant demographics: Age distribution in relation to experience.

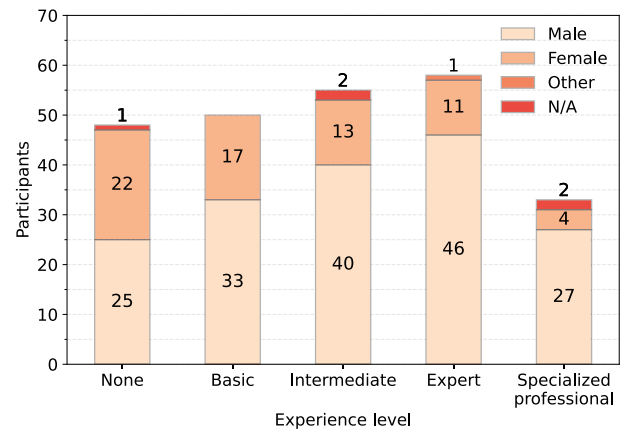


FIGURE 7. Participant demographics: Gender distribution in relation to experience.

Morpher,⁵ FotoRus [78] and InstaBeauty [79], respectively. Thus, in total six distinct manipulation classes were available. Similarity scores were computed for each manipulated face image from a comparison with the bona fide source image, by means of ArcFace face recognition system (FRS), based on a deep convolutional neural network (DCNN) approach utilizing Additive Angular Margin Loss [80], for

⁵<http://biolab.csr.unibo.it/morphedfacegenerationtools.html>



FIGURE 8. Example stimuli of the hard category (a) Morphing: UBO (b) Retouching: Fotorus (Original image sources: Hochschule Darmstadt, and FERET [image publication permitted under fair use policy], and FRCG v.2.0 [image publication permitted under fair use policy]).

which the code and pre-trained model “LResNet100E-IR, ArcFace@ms1m-refine-v2” were used.⁶ Based on these scores, we further distinguished each of these classes into easy and hard categories; see Table 4 for clarification. Representative examples for both hard and easy categories are shown in Figure 8 and Figure 9, respectively. Assuming that similarity of a bona fide source image and its manipulated counterpart can be interpreted as an indicator for detection difficulty, we surmise that low similarity scores in theory would be more noticeable, as the greater distance to the bona fide source image would suggest the image was manipulated to a higher degree. This is undoubtedly a loose approximation, as it would fail to properly reflect difficulty in terms of detection for specific manipulations and resulting signal images of high quality. As an example, considering a face swap manipulation with few visible artifacts, the similarity score may be low due to significant differences in the image. However, as the depicted face is unfamiliar to the observer, the manipulation may be hard to detect. Ten images of each class and difficulty were systematically selected for a total of 120 manipulated face images as target signal stimuli. For adequate demographic diversity within our data set, we built upon the subjective assessment of the authors. Subsequently, 60 bona fide and 60 manipulated reference face images were automatically selected to preclude any subject overlap. From this data set, image combinations balanced across manipulation types and difficulty were randomly selected as stimuli according to trial type. To summarize, our experiment included 59 bona fide and 50 manipulated target face images with an additional 14 manipulated images as reference.

F. RESULTS

Results were gathered starting 2021-08-09 and the last procedure result considered was obtained on 2021-12-13. On average, participants completed the experiment in 26 minutes. Summary statistics for data from the experiment and

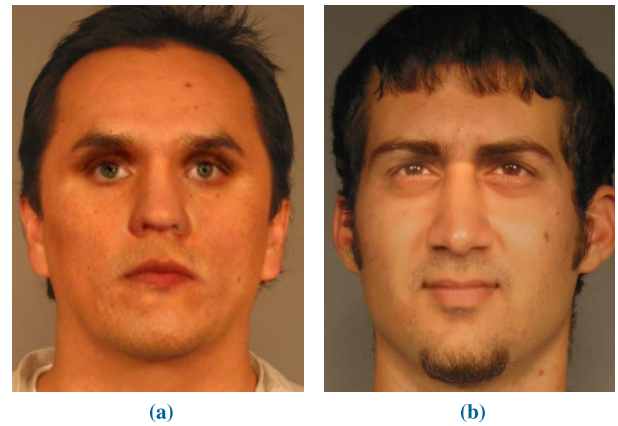


FIGURE 9. Example stimuli of the easy category (a) Face swapping: Fewshotface (b) Morphing: FaceFusion (Original image sources: Hochschule Darmstadt, and FERET [image publication permitted under fair use policy], and FRCG v.2.0 [image publication permitted under fair use policy]).

TABLE 5. Summary statistics for data from all measures used in our experiment. Perceptual tasks included spatial 2AFC and sequential ABX trial procedures.

Measure	Mean	Std. Dev.	Min.	Max.
2AFC procedure				
Accuracy (%)	75.43	±10.99	40.74	96.30
H (%)	76.01	±12.70	42.86	96.43
F (%)	25.50	±13.69	3.85	76.92
d'	1.07	±0.54	-0.39	2.29
ABX procedure				
Accuracy (%)	62.92	±10.99	34.78	91.30
H (%)	68.41	±14.63	28.57	96.43
F (%)	45.77	±17.92	11.11	88.89
d'	1.02	±0.95	-1.50	3.20
c	-0.16	±2.25	-13.13	8.13

corresponding measures are presented in Table 5. Average classification accuracy on 2AFC trials was 75.43% ($SD = 10.99\%$), and 62.92% ($SD = 10.99\%$) on ABX trials. Although participants generally performed better on the differential 2AFC task than on the sequential ABX task, mean performance was above chance level (1/2) in both cases. Notably, with regard to previous findings indicating individual performance in the associated unfamiliar face comparison task is widely distributed in the normal population (see e.g. [3], [81], [82]), our data suggests a similar distribution referring to the wide range of participant performance measures, though we deliberately did not replicate any previously reported face recognition tasks. This result ties well with the similar pattern of results that was obtained in the 2AFC training trials in [6], where $SD = 15\%$ and a range of 35 – 100% was reported. Figure 15 shows the mean classification accuracy for each trial type across self-reported experience levels. In line with previous findings (e.g. [15], [53], [57]), professional experience appears not to be a main factor predictive of performance in perceptually detecting digitally manipulated face images. Interestingly, for the ABX trials, the self-reported ‘Expert’ group did perform

⁶<https://github.com/deepinsight/insightface>




FIGURE 10. Average classification accuracy (with 95% confidence interval) across self-reported confidence levels.




FIGURE 11. Average classification accuracy (with 95% confidence interval) correlated to time in seconds participants required to submit a decision.

better than examiners with no prior experience (average ACC 64.6% vs. 61.5%), however the previously mentioned potential bias resulting from prior knowledge of the unmanipulated source data should be considered here. Furthermore, examiners who indicated they were specialized professionals performed worse in both ABX and 2AFC trials than those who reported only basic experience (average ACC_{ABX} 60.6% vs. 64% and ACC_{2AFC} 74.6% vs. 76%). The correlation of decision confidence and classification accuracy is shown in Figure 10. Participants' confidence in having made a correct decision is reflected in accuracy. Interestingly, for the ABX trial, there appears to be a higher level of uncertainty in some cases, reflected by near chance level accuracy on the lower half of the confidence scale. This may point to participants' tendency to report confidence near the mean when they are unsure. Presumably, this is the case with ABX trials, which require a high level of attention and quick responses based on perception, as opposed to visual study of given stimuli.

Conversely, classification accuracy declines the more time participants require for responding. Figure 11 illustrates this correlation. Fast responses between 0 and 3 seconds yield the highest accuracy, indicating the participants are certain of




FIGURE 12. Estimated psychometric function based on 2AFC data. ArcFace distance scores (i.e. higher values indicate substantial differences to bona fide source images) are used as a measure of stimulus intensity.




FIGURE 13. FPR and FNR measures for each type of manipulation in ABX trials.

their assessment. However, when participants require more than 8 or 9 seconds to respond, accuracy approaches levels of random guessing, in line with previous interpretation.

To further explore our assignment of difficulty levels to manipulated face images as complex visual stimuli, we pursued an evaluation based on classification accuracy scores of the 2AFC trials, as these offered the widest range of hypothetical difficulty levels. The scores were computed for each trial corresponding to one manipulation type and respective difficulty by averaging the correct responses of all observers. Given these data, we conducted a Bayesian psychometric function estimation, building on the comprehensive work of [83] and utilizing the provided implementation.⁷ The resulting estimated psychometric function Ψ is visualized in Figure 12. Indeed, this result reflects our previous assumptions and findings quite well, estimating a detection threshold of 0.493 in terms of ArcFace distance at the corresponding accuracy within the margin of error. Outlier values demonstrate the previously discussed issues regarding limited applicability of these measures of difficulty, at least for certain manipulations. Note, variance is relatively high, while the number of trials, and therefore sampled difficulties,

⁷<https://github.com/wichmann-lab/python-psignifit>

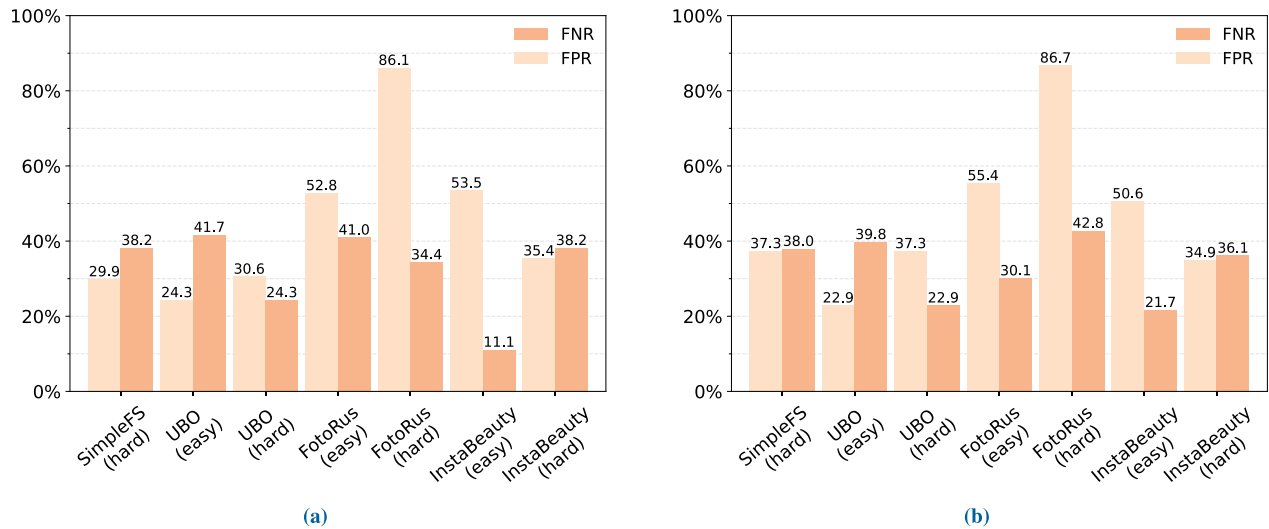


FIGURE 14. Error rates expressed as FPR and FNR for select manipulation methods from ABX trials. Results were grouped for following experience levels: (a) Novice ("None" - "Intermediate") and (b) Expert ("Expert" - "Specialized professional").

is limited. Nonetheless, we believe that this approach is well justified to consider for further exploration and refinement.

In terms of error rates, Figure 13 shows FPR and FNR in relation to different categories of previously defined manipulation types. Generally, in line with Table 4, an increase in error rates can be observed from face swap to retouch. This confirms preliminary assumptions based on visual appearance of stimuli, but more importantly sheds light on the relationship regarding computationally determined difficulty scores. Notably, compared to the general face swap category, the presumed more difficult class of morphing experiences a reversal of FPR and FNR between the easy and hard categories. This indicates participants more willingly accept morphed images as bona fide at normalized difficulty scores, in terms of Euclidean distance, of 0.27–0.28. However, when presented with morphed images with scores of 0.55–0.56, the opposite appears to be the case: at this level, participants predominantly classify bona fide face images as morphs. This points to the previously mentioned limitation regarding similarity scores as detection difficulty measures of manipulated unfamiliar face images in Section IV-E. Thus, although a relatively low similarity score might indicate substantial differences between the source and resulting manipulated image, it does not preclude a sophisticated manipulation that would not be detected by an examiner unfamiliar with the identity. Indeed, more intuitively with retouching as presumed to be the most difficult to detect manipulation, on average 53.1% and 60.8% of presented bona fide stimuli are incorrectly thought to be manipulated when compared to easy and hard difficulty reference stimuli, corresponding to 0.45–0.41 and 0.60–0.72 difficulty scores, respectively. A more detailed breakdown of FPR and FNR error measures is presented in Figure 14, where select manipulation methods are shown for novice experience levels in Figure 14a and for the expert experience group in Figure 14b. Here, previous observations are largely reflected for both groups. Remarkably, the expert group displays a higher FPR in some cases, which supports




FIGURE 15. Average classification accuracy for each level of self-reported experience level and trial type.

the assumption of a lower threshold to detect manipulations, as speculated earlier. More interesting however, are the higher FNR rates, specifically for the InstaBeauty easy and FotoRus hard manipulations, indicative of a possible advantage of novice examiners. Note the high FNR rates for some manipulations, which translates to nearly a third of manipulations going undetected. One limitation of our implementation is that our experiment scope is relatively ambitious, in that we have attempted to include many stimuli classes in addition to two trial procedures. Therefore, trials are not optimally balanced, as it was necessary to consider total experiment duration and participant fatigue, thus yielding uneven numbers of trials, which is expected to affect bias to some extent. Another limitation in terms of bias involves the issue of demographics of our sample. Due to the nature of our goal, recruiting ample numbers of good-faith participants with authentic professional experience in either biometrics or other security-related areas, we inherently reflect these demographics in terms of gender and age with regard to level of experience. Moreover, as the participant population includes a substantial number of experts of the biometric field, and the stimuli




FIGURE 16. Confidence score as reported for each decision, grouped by experience levels for each type of manipulation, i.e. face swapping, morphing and retouching.

are taken from well-known, public datasets, potential bias can arise from prior knowledge of the unmanipulated source data. Participants were consistently more confident in their decisions regarding face swap stimuli than morph or retouch stimuli. Figure 16 illustrates our finding across all experience levels, that participants in fact perceive detecting face swap manipulation as an easier task than e.g. morphs. Note, these data only reflect participant subjective assessment in terms of certainty of a previous decision, regardless of actual validity. Nevertheless, perceived difficulty positively correlates with stimulus difficulty measures, i.e. ArcFace distance scores, previously presented in Table 4.

V. CONCLUSION

This paper proposes a novel approach to utilizing well established methods in the area of psychophysics, implemented as a web-based experiment, to examine the human ability to perceptually detect digitally manipulated face images. We obtained an average ACC of 75.43% in 2AFC trials and 62.92% in the ABX trials from 227 participants. We did not find a significant difference in general performance related to the characteristic of professional experience, even given a relatively large range of options for participant self-assessment. Unexpectedly, specialized professionals were in some cases less successful than participants with little or no experience. This can be explained by two probable factors. The first is the relatively small sample size and inherent effects of gender and age bias; the second is the possibility that trained professionals have a tendency to assume manipulated images based on other internal criteria beyond basic face perception.

We have, however, shown that there are significant differences in terms of FPR and FNR among different manipulation types, suggesting different levels of difficulty can be determined. Importantly, our results indicate the use of DCNN based FRS to estimate detection difficulty of manipulated images might be a component in future attempts to improve the mapping of stimuli in face perception tasks. On average, performance was poor in terms of error rates, especially relatively high FPR and FNR across all experience levels on ABX

trials. Nevertheless, we found individual participants who performed exceptionally well, with classification accuracy reaching 96.30% and hit rates (TPR) of 96.43%.

Collectively, our results are consistent with previous findings on professional experience in detecting face image manipulations and the related topic concerning face comparison and perception. This experiment adds to a growing corpus of research showing that human detection ability in this regard is limited; however, some individuals display extraordinary performance, and further study of these high-performing individuals is desirable for future work. Additionally, it will be important that future research further investigate the adequacy of different test methodologies and evaluation protocols in relation to this category of visual perception tasks.

REFERENCES

- [1] W. W. Bledsoe, *The Model Method in Facial Recognition*, vol. 15. Palo Alto, CA, USA: Panoramic Research, 1966, p. 2.
- [2] A. Burton, D. White, and A. McNeill, "The Glasgow face matching test," *Behav. Res. Methods*, vol. 42, pp. 91–286, Feb. 2010.
- [3] J. M. McCaffery, D. J. Robertson, A. W. Young, and A. M. Burton, "Individual differences in face identity processing," *Cognit. Res., Princ. Implications*, vol. 3, no. 1, p. 21, Jun. 2018, doi: [10.1186/s41235-018-0112-9](https://doi.org/10.1186/s41235-018-0112-9).
- [4] D. White, R. Kemp, R. Jenkins, M. Matheson, and A. Burton, "Passport officers' errors in face matching," *PLoS ONE*, vol. 9, no. 8, Aug. 2014, Art. no. e103510, doi: [10.1371/journal.pone.0103510](https://doi.org/10.1371/journal.pone.0103510).
- [5] D. J. Robertson, R. S. S. Kramer, and A. M. Burton, "Fraudulent ID using face morphs: Experiments on human and automatic recognition," *PLoS ONE*, vol. 12, no. 3, Mar. 2017, Art. no. e0173319, doi: [10.1371/journal.pone.0173319](https://doi.org/10.1371/journal.pone.0173319).
- [6] D. J. Robertson, A. Mungall, D. G. Watson, K. A. Wade, S. J. Nightingale, and S. Butler, "Detecting morphed passport photos: A training and individual differences approach," *Cognit. Res., Principles Implications*, vol. 3, no. 1, p. 27, Jun. 2018, doi: [10.1186/s41235-018-0113-8](https://doi.org/10.1186/s41235-018-0113-8).
- [7] R. Kramer, M. Mireku, T. Flack, and K. Ritchie, "Face morphing attacks: Investigating detection with humans and computers," *Cognit. Res., Princ. Implications*, vol. 4, pp. 1–4, Dec. 2019.
- [8] D. M. Green and J. A. Swets, *Signal Detection Theory Psychophysics*. New York, NY, USA: Wiley, 1966.
- [9] F. A. Kingdom and N. Prins, "Psychophysics—A practical introduction," in *Psychophysics*, 2nd ed., F. A. Kingdom and N. Prins, Eds. San Diego, CA, USA: Academic, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780124071568000013>
- [10] S. Stevens and L. E. Marks, *Psychophysics: Introduction to Its Perceptual, Neural, Social Prospects*. Evanston, IL, USA: Routledge, 2017.
- [11] A. J. Dowsett and A. M. Burton, "Unfamiliar face matching: Pairs outperform individuals and provide a route to training," *Brit. J. Psychol.*, vol. 106, no. 3, pp. 433–445, Aug. 2015.
- [12] B. Duchaine and K. Nakayama, "The Cambridge face memory test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants," *Neuropsychologia*, vol. 44, pp. 85–576, Dec. 2006.
- [13] M. Ngan, P. Grother, K. Hanaoka, and J. Kuo, "Face recognition vendor test (FRVT) part 4: Morph-performance of automated face morph detection," *Nat. Inst. Standards Technol., Gaithersburg, MD, USA, Tech. Rep.* 8292, Mar. 2020.
- [14] V. Dan, B. Paris, J. Donovan, M. Hameleers, J. Roozenbeek, S. van der Linden, and C. von Sikorski, "Visual mis and disinformation, social media, and democracy," *J. Mass Commun. Quart.*, vol. 98, no. 3, pp. 641–664, 2021, doi: [10.1177/10776990211035395](https://doi.org/10.1177/10776990211035395).
- [15] M. Ferrara, A. Franco, and D. Maltoni, *On Effects Image Alterations Face Recognition Accuracy*. Cham, Switzerland: Springer, 2016, pp. 195–222, doi: [10.1007/978-3-319-28501-6_9](https://doi.org/10.1007/978-3-319-28501-6_9).
- [16] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, A. Morales, and J. Ortega-Garcia, "Deepfakes and beyond: A survey of face manipulation and fake detection," *Inf. Fusion*, vol. 64, pp. 131–148, Dec. 2020.

- [17] C. Rathgeb, R. Tolosana, R. Vera, and C. Busch, *Handbook of Digital Face Manipulation and Detection: From DeepFakes to Morphing Attacks* (Advances in Computer Vision and Pattern Recognition), 1st ed. Cham, Switzerland: Springer, 2022.
- [18] Z. Akhtar, D. Dasgupta, and B. Banerjee, "Face authenticity: An overview of face manipulation generation, detection and recognition," in *Proc. Int. Conf. Commun. Inf. Process. (ICCIP)*, 2019, pp. 1–8, doi: [10.2139/ssrn.3419272](https://doi.org/10.2139/ssrn.3419272).
- [19] F. Richter. (Jul. 2019). *Faceapp Craze Reaches New Heights*. [Online]. Available: <https://www.statista.com/chart/18769/estimated-worldwide-faceapp-downloads-by-platform/>
- [20] Airnow. (Jul. 2021) *Leading Entertainment APP Titles in the Apple App Store in Spain*. [Online]. Available: <https://www.statista.com/statistics/699448/leading-iphone-entertainment-apps-in-spain-by-downloads/>
- [21] Reuters. (Jun. 2020). *Germany Bans Digital Doppelganger Passport Photos*. [Online]. Available: <https://www.reuters.com/article/us-germany-tech-morphing/german-bans-digital-doppelganger-passport-photos-idUSKBN23A1YM>
- [22] M. C. Stamm and K. R. Liu, "Forensic detection of image manipulation using statistical intrinsic fingerprints," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 3, pp. 492–506, Dec. 2010.
- [23] H. Farid, "Image forgery detection," *IEEE Signal Process. Mag.*, vol. 26, no. 2, pp. 16–25, Mar. 2009.
- [24] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport," in *Proc. IEEE/IAPR Int. Joint Conf. Biometrics*, Dec. 2014, pp. 1–7.
- [25] R. Chen, X. Chen, B. Ni, and Y. Ge, "SimSwap: An efficient framework for high fidelity face swapping," in *Proc. 28th ACM Int. Conf. Multimedia*, 2020, pp. 2003–2011.
- [26] G. Wolberg, *Digital Image Warping*, vol. 10662. Los Alamitos, CA, USA: IEEE Computer Society Press, 1990.
- [27] T. Beier and S. Neely, "Feature-based image metamorphosis," in *Proc. 19th Annu. Conf. Comput. Graph. Interact. Techn.*, New York, NY, USA, 1992, pp. 35–42, doi: [10.1145/133994.134003](https://doi.org/10.1145/133994.134003).
- [28] T. Guanglai, C.-C. Chen, S. Gengxin, and B. Sheng, "Research on face image morphing based on automatic feature detection," in *Proc. 2nd Eurasia Conf. Biomed. Eng.*, 2020, pp. 172–174.
- [29] N. Damer, A. M. Saladié, A. Braun, and A. Kuijper, "Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network," in *Proc. IEEE 9th Int. Conf. Biometrics Theory, Appl. Syst. (BTAS)*, 2018, pp. 1–10.
- [30] C. Lee, M. T. Schramm, M. Boutin, and J. P. Allebach, "An algorithm for automatic skin smoothing in digital portraits," in *Proc. 16th Int. Conf. Image Process. (ICIP)*, 2009, pp. 3149–3152.
- [31] A. Dhall, G. Sharma, R. Bhatt, and G. M. Khan, "Adaptive digital makeup," in *Proc. ISVC (Lecture Notes in Computer Science)*, vol. 5876, G. Bebis, R. D. Boyle, and B. Parvin, Eds. Springer, 2009, pp. 728–736, doi: [10.1007/978-3-642-10520-3](https://doi.org/10.1007/978-3-642-10520-3).
- [32] Y. Novskaya, R. Sun, H. Zhu, and L. Ma, "Floor-ladder framework for human face beautification," in *Proc. 6th Int. Conf. Anal. Images, Social Netw. Texts*, Moscow, Russia, in *Lecture Notes in Computer Science*, vol. 10716, W. M. P. van der Aalst, D. I. Ignatov, M. Y. Khachay, S. O. Kuznetsov, V. S. Lempitsky, I. A. Lomazova, N. V. Loukachevitch, A. Napoli, A. Panchenko, P. M. Pardalos, A. V. Savchenko, and S. Wasserman, Eds. Cham, Switzerland: Springer, Jul. 2017, pp. 255–266, doi: [10.1007/978-3-319-73013-4](https://doi.org/10.1007/978-3-319-73013-4).
- [33] D. Kravets. (Aug. 2017). *Driver's License Facial Recognition Tech Leads to 4,000 New York Arrests*. [Online]. Available: <https://arstechnica.com/tech-policy/2017/08/biometrics-leads-to-thousands-of-a-ny-arrests-for-fraud-identity-theft/>
- [34] A. Boyd. (Apr. 2021). *CBP's Facial Biometrics Program has Caught 26 Alleged Imposters*. [Online]. Available: <https://www.defenseone.com/technology/2018/11/cbps-facial-biometrics-program-has-caught-26-alleged-imposters/152993/>
- [35] T. Costello and E. Sacks. (Aug. 2018). *New Facial Recognition Tech Catches First Impostor at DC Airport*. [Online]. Available: <https://www.nbcnews.com/news/us-news/new-facial-recognition-tech-catches-first-impostor-d-c-airport-n903236>
- [36] *Information Technology—Biometric Presentation Attack Detection—Part 1: Framework*, Standard ISO/IEC 30107-1, International Standards Organization, 2016.
- [37] (Jan. 2020). *Facial Recognition Statistics in Airports*. [Online]. Available: <https://www.reservations.com/blog/resources/facial-recognition-airports-survey/>
- [38] M. Xie. (Sep. 2018). *Council Post: The Future of Biometric Facial Recognition*. [Online]. Available: <https://www.forbes.com/sites/forbestechcouncil/2018/09/19/the-future-of-biometric-facial-recognition/>
- [39] FastPass. (Aug. 2021). *A Harmonized, Modular Reference System for all European Automated Border Crossing Points*. [Online]. Available: <https://www.fastpass-project.eu/>
- [40] E.-V. Pikoulis, Z.-M. Ioannou, M. Paschou, and E. Sakkopoulos, "Face morphing, a modern threat to border security: Recent advances and open challenges," *Appl. Sci.*, vol. 11, p. 3207, Apr. 2021.
- [41] B. B. Strickland, *The Gale Encyclopedia of Psychology*. Farmington Hills, MI, USA: Gale, 2001.
- [42] G. T. Fechner, *Elemente der Psychophysik*, vol. 2. Leipzig, Germany: Breitkopf & Härtel, 1860.
- [43] H. T. Lawless and H. Heymann, *Discrimination Testing*. New York, NY, USA: Springer, 2010, pp. 79–100, doi: [10.1007/978-1-4419-6488-5_4](https://doi.org/10.1007/978-1-4419-6488-5_4).
- [44] W. H. Ehrenstein and A. Ehrenstein, *Psychophysical Methods*. Berlin, Germany: Springer, 1999, pp. 1211–1241, doi: [10.1007/978-3-642-58552-4_43](https://doi.org/10.1007/978-3-642-58552-4_43).
- [45] R. Zwisler. (1998). *Psychophysische Methoden*. [Online]. Available: <http://www.zwisler.de/scripts/methoden/node4.html>
- [46] S. Glasauer and Z. Shi, "The origin of Vierordt's law: The experimental protocol matters," *Psych J.*, vol. 10, no. 5, pp. 732–741, 2021. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pchj.464>
- [47] G. V. Békésy, "The sound pressure difference between the round and the oval windows and the artificial window of labyrinthine fenestration," *Acta OtoLaryngol.*, vol. 35, no. 3, pp. 301–315, 1947.
- [48] J. Weber, *Atlas der Computerperimetrie*. Berlin, Germany: Springer, 1993.
- [49] F. A. Wichmann and N. J. Hill, "The psychometric function: I. fitting, sampling, and goodness of fit," *Perception Psychophys.*, vol. 63, no. 8, pp. 1293–1313, Nov. 2001, doi: [10.3758/BF03194544](https://doi.org/10.3758/BF03194544).
- [50] Q. Yang and M. L. Ng, "Paired comparison/directional difference test/2-alternative forced choice (2-AFC) test, simple difference test/same-different test," in *Discrimination Testing Sensory Sci.*, L. Rogers, Ed. Amsterdam, The Netherlands: Elsevier, 2017, pp. 109–134.
- [51] R. E. Greenaway, "ABX discrimination task," in *Discrimination Testing in Sensory Science* (Woodhead Publishing Series in Food Science, Technology and Nutrition), L. Rogers, Ed. Amsterdam, The Netherlands: Elsevier, 2017, pp. 267–288.
- [52] A. Makrushin, T. Neubert, and J. Dittmann, "Automatic generation and detection of visually faultless facial morphs," in *Proc. 12th Int. Joint Conf. Comput. Vis., Imag. Comput. Graph. Theory Appl.*, 2017, pp. 39–50.
- [53] A. Makrushin, T. Neubert, and J. Dittmann, "Humans vs. algorithms: Assessment of security risks posed by facial morphing to identity verification at border control," in *Proc. 14th Int. Joint Conf. Comput. Vis., Imag. Comput. Graph. Theory Appl.*, 2019, pp. 1–4.
- [54] A. Rössler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner, "FaceForensics++: Learning to detect manipulated facial images," *Tech. Rep.*, Jan. 2019.
- [55] S.-Y. Wang, O. Wang, A. Owens, R. Zhang, and A. A. Efros, "Detecting photoshopped faces by scripting photoshop," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 10072–10081.
- [56] P. Korshunov and S. Marcel, "Subjective and objective evaluation of deepfake videos," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2021, pp. 2510–2514.
- [57] A. Makrushin, D. Siegel, and J. Dittmann, "Simulation of border control in an ongoing web-based experiment for estimating morphing detection performance of humans," in *Proc. Workshop Inf. Hiding Multimedia Secur.*, New York, NY, USA, Jun. 2020, pp. 91–96.
- [58] S. J. Nightingale, S. Agarwal, and H. Farid, "Perceptual and computational detection of face morphing," *J. Vis.*, vol. 21, no. 3, p. 4, Mar. 2021, doi: [10.1167/jov.21.3.4](https://doi.org/10.1167/jov.21.3.4).
- [59] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer, and C. Busch, "MIPGAN generating strong and high quality morphing attacks using identity prior driven GAN," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 3, no. 3, pp. 365–383, Jul. 2021.
- [60] R. Kemp, N. Towell, and G. Pike, "When seeing should not be believing: Photographs, credit cards and fraud," *Appl. Cognit. Psychol.*, vol. 11, no. 3, pp. 211–222, 1997.
- [61] B. Wirth and C.-C. Carbon, "An easy game for frauds? Effects of professional experience and time pressure on passport-matching performance," *J. Exp. psychol. Appl.*, vol. 23, no. 2, pp. 138–157, Apr. 2017.
- [62] D. Weatherford, D. Roberson, and W. Erickson, "When experience does not moderate expertise: Security professionals fail to detect low prevalence fake IDs," *Cognit. Res., Princ. Implication*, vol. 6, p. 45, Dec. 2021.

- [63] S. Billson, "Two-out-of-five test," in *Discrimination Testing in Sensory Science* (Woodhead Publishing Series in Food Science, Technology and Nutrition), L. Rogers, Ed. Chicago, IL, USA: Woodhead, 2017, pp. 171–181.
- [64] J. M. Ennis, B. Rousseau, and D. M. Ennis, "Proper task instructions and the two-out-of-five-test," *IFPress*, vol. 15, no. 4, pp. 3–4, 2012.
- [65] H. Stanislaw and N. Todorov, "Calculation of signal detection measures," *Behav. Res. Methods, Instrum., Comput. J. Psychonomic*, vol. 31, pp. 49–137, Mar. 1999.
- [66] M. J. Hautus, "Corrections for extreme proportions and their biasing effects on estimated values of d' ," *Behav. Res. Methods*, vol. 27, pp. 46–51, Mar. 1995, doi: [10.3758/BF03203619](https://doi.org/10.3758/BF03203619).
- [67] European Parliament and Council. (2016). (EU) General Data Protection Regulation. [Online]. Available: <http://data.europa.eu/eli/reg/2016/679/oj>
- [68] N. A. Macmillan and C. D. Creelman, *Detection Theory* East Sussex, U.L.: Psychology Press, 2004.
- [69] *Information Technology—Biometric Presentation Attack Detection—Part 3: Testing and Reporting*, Standard ISO/IEC 30107-3, International Standards Organization, 2017.
- [70] J. Boley and M. Lester, "Statistical analysis of abx results using signal detection theory," in *Proc. 127th Audio Eng. Soc. Conv.*, New York, NY, USA, Oct. 2009. [Online]. Available: <http://www.aes.org/e-lib/browse.cfm?elib=15022>
- [71] M. J. Hautus and X. Meng, "Decision strategies in the ABX (matching-to-sample) psychophysical task," *Perception Psychophys.*, vol. 64, no. 1, pp. 89–106, Jan. 2002, doi: [10.3758/BF03194559](https://doi.org/10.3758/BF03194559).
- [72] G. T. Buracas, "The relationship between task performance and functional magnetic resonance imaging response," *J. Neurosci.*, vol. 25, no. 12, pp. 3023–3031, Mar. 2005.
- [73] T. Elze, "Misspecifications of stimulus presentation durations in experimental psychology: A systematic review of the psychophysics literature," *PLoS ONE*, vol. 5, no. 9, Sep. 2010, Art. no. e12792.
- [74] J. Willis and A. Todorov, "First impressions," *Psychol. Sci.*, vol. 17, no. 7, pp. 592–598, Jul. 2006.
- [75] C.-C. Carbon, "The first 100 milliseconds of a face: On the microgenesis of early face processing," *Perceptual Motor Skills*, vol. 113, no. 3, pp. 859–874, Dec. 2011.
- [76] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The feret database and evaluation procedure for face-recognition algorithms," *Image Vis. Comput.*, vol. 16, no. 5, pp. 295–306, 1998.
- [77] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the face recognition grand challenge," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 1, Dec. 2005, pp. 947–954.
- [78] Fotoable. (2021). *Fotorus—Camera & Photo Editor*. [Online]. Available: <https://apps.apple.com/us/app/fotorus-camera-photo-editor/id457517348>
- [79] Fotoable. (2021). *Instabeauty—Makeup Camera*. [Online]. Available: <https://apps.apple.com/us/app/instabeauty-makeup-camera/id599534650>
- [80] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Dec. 2019, pp. 4685–4694.
- [81] V. Bruce, M. Bindemann, and K. Lander, "Individual differences in face perception and person recognition," *Cognit. Res., Princ. Implications*, vol. 3, p. 45, Dec. 2018.
- [82] J. P. Davis, K. Lander, R. Evans, and A. Jansari, "Investigating predictors of superior face recognition ability in police super-recognisers," *Appl. Cognit. Psychol.*, vol. 30, no. 6, pp. 827–840, Nov. 2016.
- [83] H. H. Schätt, S. Harmeling, J. H. Macke, and F. A. Wichmann, "Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data," *Vis. Res.*, vol. 122, pp. 105–123, May 2016.



ROBERT NICHOLS is currently pursuing the master's degree with the Faculty of Computer Science, Hochschule Darmstadt (HDA), Germany. He works as a Student Research Assistant with the da/sec—Biometrics and Internet Security Research Group, HDA. His research interests include biometrics, in particular facial recognition.



CHRISTIAN RATHGEB is currently a Senior Researcher with the Faculty of Computer Science, Hochschule Darmstadt (HDA), Germany. He is a Principal Investigator at the National Research Center for Applied Cybersecurity ATHENE. He has coauthored over 100 technical papers in the field of biometrics. His research interests include pattern recognition, iris and face recognition, security aspects of biometric systems, secure process design, and privacy enhancing technologies for biometric systems. He is a member of the European Association for Biometrics (EAB). He is a Winner of the EAB—European Biometrics Research Award 2012, the Austrian Award of Excellence 2012, the Best Poster Paper Awards (IJCB'11, IJCB'14, and ICB'15), the Best Paper Award Bronze (ICB'18), and the Best Paper Award (WIFS'21). He is the Program Chair of the International Conference of the Biometrics Special Interest Group (BIOSIG). He is an Editorial Board Member of *IET Biometrics* (IET BMT). He has served for various program committees and conferences, such as ICB, IJCB, BIOSIG, and IWBF; and journals as a Reviewer, such as *IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY*, *IEEE TRANSACTIONS ON BIOMETRICS*, *BEHAVIOR*, and *IDENTITY SCIENCE*, and *IET BMT*.



PAWEŁ DROZDOWSKI worked as a Senior Researcher with the Faculty of Computer Science, Hochschule Darmstadt (HDA), Germany. His research interests include biometrics, information security and privacy, pattern recognition, and algorithmic fairness. He has coauthored over 35 technical publications in the field of biometrics. He has won the European Biometric Industry Award of the European Association for Biometrics (EAB), in 2021, the Best Student Paper Runner-Up Award (WIFS'18), the Best Poster Award (BIOSIG'19), and the Best Paper Award (WIFS'21).



CHRISTOPH BUSCH (Senior Member, IEEE) is currently a member of the Norwegian University of Science and Technology (NTNU), Norway. He holds a joint appointment with the Hochschule Darmstadt (HDA), Germany. Further, he has been lecturing biometric systems at Denmark's DTU, since 2007. On behalf of the German BSI, he has been the Coordinator for the project series BioIS, BioFace, BioFinger, BioKeyS Pilot-DB, KBEinweg, and NFIQ2.0. He was/is a Partner of the EU projects 3D-Face, FIDELITY, TURBINE, SOTAMD, RESPECT, TreSPsS, iMARS, and others. He is also a Principal Investigator of the German National Research Center for Applied Cybersecurity (ATHENE) and a Co-Founder of the European Association for Biometrics (EAB). He is a Convenor of WG3 in ISO/IEC JTC1 SC37. He has coauthored more than 500 technical papers and has been a speaker at international conferences. He is a member of the Editorial Board of the *IET Biometrics* and formerly of the *IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY* journal. Furthermore, he chairs the TeleTrusT Biometrics Working Group as well as the German standardization body on biometrics.

...