Supporting Human Examiners in Facial Image Manipulation Detection

Published Version:
Dadi, M., Franco, A., Maltoni, D. (2024). Supporting Human Examiners in Facial Image Manipulation Detection. New York : IEEE [10.1109/MetroXRAINE62247.2024.10797036].

Availability:
This version is available at: https://hdl.handle.net/11585/1001381 since: 2025-01-13

Published:
DOI: http://doi.org/10.1109/MetroXRAINE62247.2024.10797036

Supporting human examiners in facial image manipulation detection

Maichol Dadi, Annalisa Franco, Davide Maltoni
Dipartimento di Informatica — Scienza e Ingegneria (DISI)
University of Bologna
Cesena, Italy
{maichol.dadi2, annalisa.franco, davide.maltoni}@unibo.it

***Abstract*—Many countries lack live enrollment procedures for facial images in identity documents, allowing citizens to provide their own photos. This operational procedure opens avenues for fraudulent attacks, primarily through digital manipulations of facial images, which have an impact on automated identity verification performance. Detecting such manipulations before image inclusion into documents is critical, especially given the difficulty in identifying them visually. To address this, a software tool is proposed to assist human examiners in analyzing facial images for electronic ID documents. The tool aligns with guidelines from the Facial Identification Scientific Working Group, implementing measures covering various facial components to ease the comparison between the ID photo and a live-captured image during enrollment.**

***Index Terms*—Electronic Machine-Readable Travel Documents (eMRTD), face image manipulation, facial image analysis, face beautification, face morphing.**

I. Introduction

One of the main challenges of biometric identification is to ensure that the identity data stored in a document, such as a passport or an ID card, correspond to the real identity of the person who owns it. This is especially important for electronic documents that embed biometric features, such as face images, fingerprints, or iris patterns, which can be used for automated verification at border crossings or other security checkpoints. However, the reliability of these biometric features depends on their integrity and authenticity, which can be compromised by various types of attacks, such as presentation attacks or digital manipulations of the images. Indeed, in many countries the enrollment procedure for identity documents doesn't include live enrollment of the face image, thus allowing citizens to bring their own photo printed on photographic paper. Unfortunately, such images might have been somehow digitally altered or manipulated, for instance to improve facial appearance (e.g. improving skin texture), or could present some geometric distortions introduced by acquisition devices or due to an uncareful printing process, or might finally have been intentionally manipulated with criminal intent. Among others, the morphing attack recently emerged as a serious security threat [1], [2], and several real cases have been reported by the authorities of different countries to confirm this risk. Such manipulations have a negative impact on the performance of automated identity verification, and it is therefore crucial to detect them before the image is included into the document. Unfortunately, tests conducted with human examiners [3] have shown that this type of manipulation is very difficult to detect with the naked eye. While digital image beautification is easier to spot since it introduces quite visible texture alterations (e.g. very smooth skin), other manipulations producing geometric distortions are very hard to identify.

The tool design and implementation are inspired by the guidelines defined by the Facial Identification Scientific Working Group for the morphological analysis recommended for manual facial comparison [4]. Several measures have been implemented, covering the different facial components. The single functionalities will be described in detail in the following section. It is worth noting that the aim of the tool itself is not to detect possible manipulations or to influence the human examiner's decision but only to facilitate the facial analysis and comparison.

This project received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 883356. Disclaimer: this text reflects only the author's views, and the Commission is not liable for any use that may be made of the information contained therein.

II. Tool overview

The main components of the software tool are outlined in Figure 1. The tool allows to compare two images, the one which should be included into the ID document D and a live captured image L, used as a ground truth to check if the first image has been somehow manipulated. The two faces are first aligned based on the facial landmarks automatically extracted, and then each face is segmented into its main components (eyes, nose, mouth, face region, etc.) in order to facilitate the comparison of the single parts. Moreover, further functionalities are implemented to allow the user to measure and compare specific distances between reference points (automatically detected or manually labeled). More in detail, the image comparison workflow includes several steps, described in detail in the following subsection.

A. Facial landmark detection

A preliminary alignment stage is necessary to deal with different face size in the two images to compare; facial landmarks automatically extracted using the Dlib library [5] are exploited to rescale and align the two images based on the position. In particular, the two images are rotated, in order to

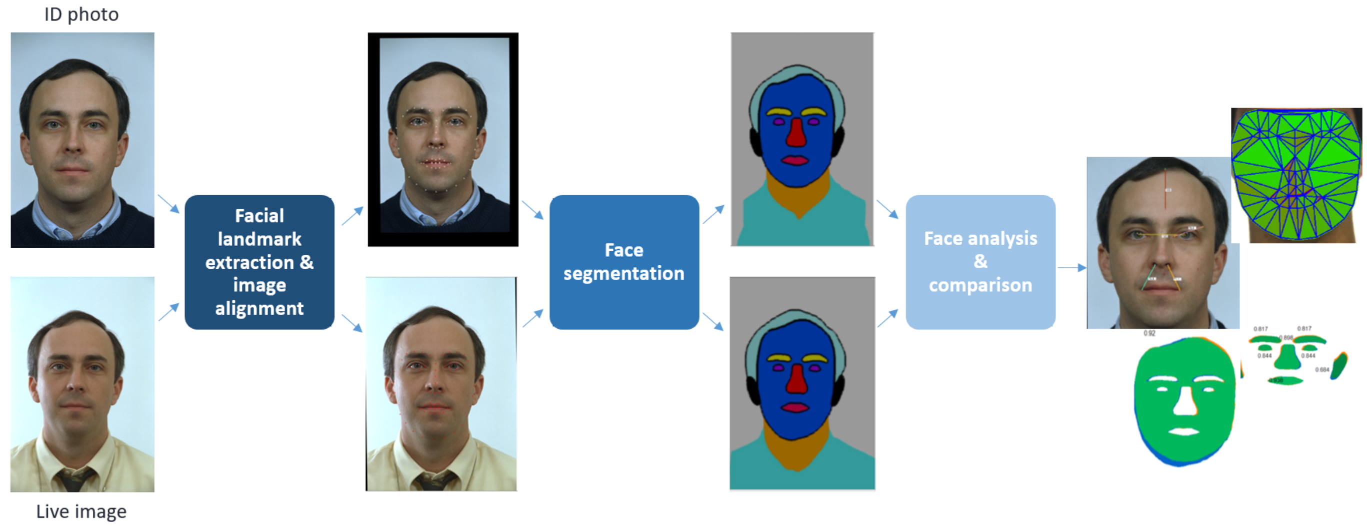

Fig. 1. The image processing workflow adopted by the tool. A pre-processing is applied to the two input images to be compared to align them; then the two images are segmented to identify the single facial components and finally a comparison utility is provided based on the facial landmarks, the face triangulation and the segmetation computed in the previous stages.

have a horizontal eyes line, and then rescaled based on the eyes distance to the smaller one. As further information for face comparison, the tool performs the Delaunay triangulation [6] based on the facial landmarks previously identified.

B. Face segmentation

After alignment, the two face images are segmented into their main components: face region, hairs, eyes, eye browses, nose, mouth, ears, neck, and dresses. This segmentation stage is useful to estimate the overlap of the single face regions in the two images. An open-source tool [7], based on a Bilateral Segmentation Network (BiSeNet v1), has been used to this purpose. The segmentation network works on images of size 512×512, so the aligned face images are resized to that reference dimension to be segmented and the resulting segmentation mask is then resized to the original image size to preserve the facial aspect ratio.

C. Facial analysis and comparison

Several measures have been implemented to compare the two faces from different points of view. We tried to take into account the different facial components, including regions, such as the ears, which are typically ignored by automatic face processing techniques but actually contain useful information for image analysis and comparison.

1) Face components: An overall measure of the similarity between the two images can be obtained by comparing the two related segmentation masks. In particular, for each single face component in the two images the Intersection over Union (IoU) is computed as the ratio between the area of overlap between the two corresponding regions and the area of their union. The IoU value is a value in the range [0;1] where 1 indicates a perfect overlap. A comparison of the IoU values for the single components will focus the human examiner's attention on specific regions to detect possible deviations. An example of IoU computation is given in Figure 2 for a bona fide image pair as well as for some image pairs taken from the questionnaire described in [3] where one of the two images is manipulated. In particular, an example for each kind of alteration is provided: pincushion geometric distortion, barrel geometric distortion, digital beautification and morphing. For the bona fide pair, high values of IoU are observable for the two images, despite of small changes in pose or lighting. The good overlap is also visually confirmed; all facial components are quite well aligned. On the contrary, in all manipulation cases, lower values of IoU are observed and also visually the overlap between the single components is not optimal. The typical effects produced by the single manipulations are observable in the IoU values; for instance, face beautification acts mainly on eyebrowse, nose and mouth.

The same measure of Intersection over Union is computed for each pair of corresponding triangles in the two images and the resulting value (between 0 and 1) is visually represented by a color map (ranging from red to green, according to the degree of overlap). This kind of visualization is particularly useful for some kinds of manipulation which act mainly on certain facial regions. An example of application is given in Figure 3 for a bona fide image pair as well as image pairs with manipulation. In the IoU visualization, the IoU value is mapped into a color map where, as expected, green means high overlap and red low overlap.

Even though some limited deviations are observable even for the bona fide pair, due for instance to small pose changes, all the IoU values measured for this specific case are quite high. As a term of comparison, the IoU results computed for face triangulation on the manipulated image pairs present a higher variability and a stronger presence of red areas with respect to the genuine example is evident. In particular, noticeable deviations are visible in the eyes, nose, mouth, and

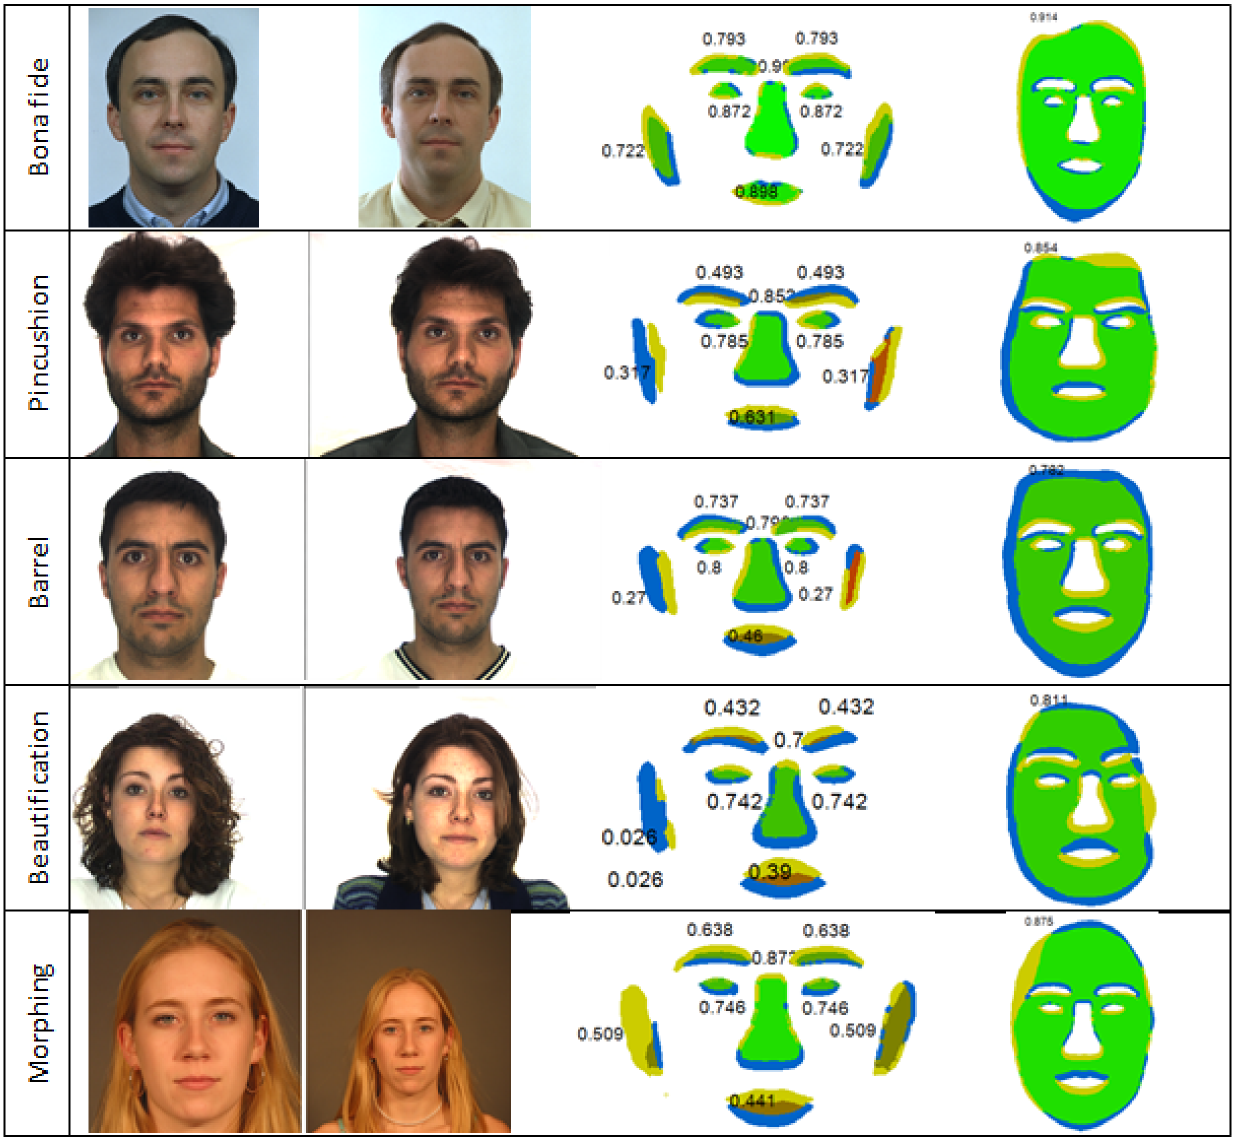

Fig. 2. Examples of IoU computed on face segmentation for some image pairs taken from the questionnaire [3]; in particular, a bona fide pair as well as one example for each manipulation is reported: pincushion geometric distortion, barrel geometric distortion, digital beautification and morphing.

chin regions.

2) Facial proportions: As suggested by the FISWG document [4] some general proportions hold for the different face parts (see Figure 4):

- the face can be divided into three equal parts: hairline to nasal root, nasal root to nasal base, and nasal base to chin;
- the width of the chin is the same as the width of the mouth;
- the distance between the inner corners of the eyes is equal to the width of one eye;
- the mentolabial sulcus is half the distance from the nasal base to the chin.

Even if not all faces conform to these proportions, such measures can be a valid contribution to the detection of image alterations since any kind of manipulation changing the face geometry, will be reflected in a modification of these values.

3) Eyes analysis: The FISWG document suggests some other interesting reference measures in the eye's region (see Figure 5). In addition, other reference measurements such as the intercanthal distance (distance between inner corners of the right and left eyes) and eyes width are considered in the tool. The tool offers therefore the possibility to observe these measures; reference points for measurement computation are automatically calculated starting from the detected landmarks, but the user can manually modify the points to label them with higher accuracy.

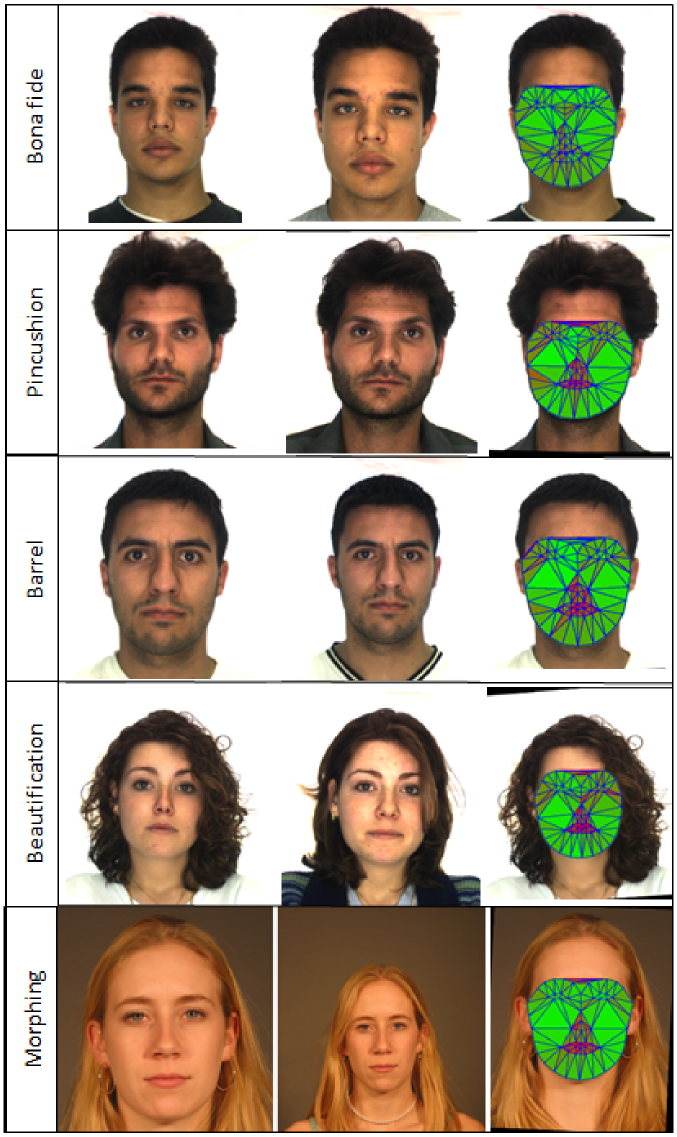

Fig. 3. Examples of IoU computed on face triangulation for some image pairs taken from the questionnaire in [3]; in particular, a bona fide pair as well as one example for each manipulation is reported: pincushion geometric distortion, barrel geometric distortion, digital beautification and morphing.

III. Validation of the Tool Measurements

In order to validate the main measures provided by the tool and evaluate their effectiveness to facilitate the detection of image manipulation, we computed some statistics aimed at estimating the impact of image manipulations on the different facial measurements. A set of 63 images has been considered (21 for barrel distortion, 21 for pincushion distortion, 21 for digital beautification). We first analysed in Figure 6 the IoU values observed for the different facial components obtained by the face segmentation approach (see Section II-B). The ideal value is 1 which represents a perfect overlap between the two compared images. The graphs report the results for geometric distortions (barrel and pincushion), and for digital beautification (applied with three increasing levels of intensity). These results mainly reflect the changes in the facial components area for the different manipulations. The deviation from ideal value of 1 is quite limited for some regions, but for others such as the ears or the eyes region, noticeable differences can be observed. Moreover, a higher intensity of beautification corresponds to a higher deviation of the IoU value. These results are in line with the expectations since the manipulations have a higher impact on some face with respect to other.

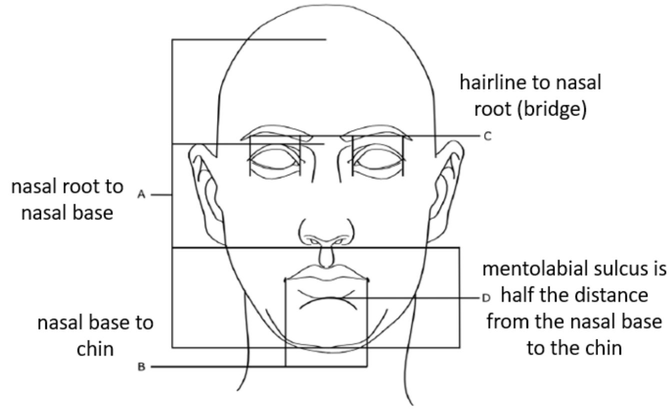

Fig. 4. Some of the traditional canons for ideal facial proportions described in [4].

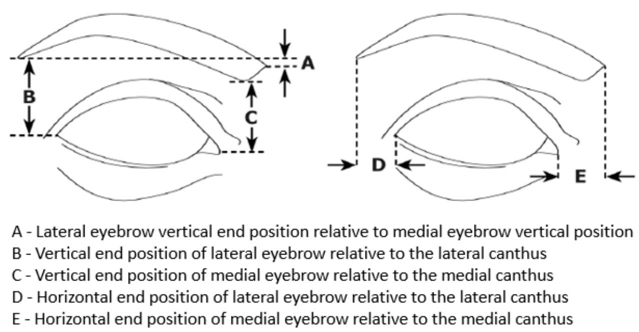

Fig. 5. Eyes measurements based on the analysis of the position of the eyebrow relative to the position of the eye opening.

With regard to facial proportions described in Figure 4 and II-C, we computed for different image pairs the deviation of each measure observed in the altered image with respect to the original one as follows:

$$dev = \frac{|d_{altered} - d_{original}|}{d_{original}} \quad (1)$$

where $d_{altered}$ is a specific distance or value computed in the manipulated image and $d_{original}$ is the same indicator computed in the original image. The higher the value, the greater the impact of manipulation. The deviations are analyzed in Figure 7 for geometric distortions and beautification. In this case the absolute deviation with respect to the unmanipulated image is computed for the different measures. The average results obtained are reported in the graphs. The results show

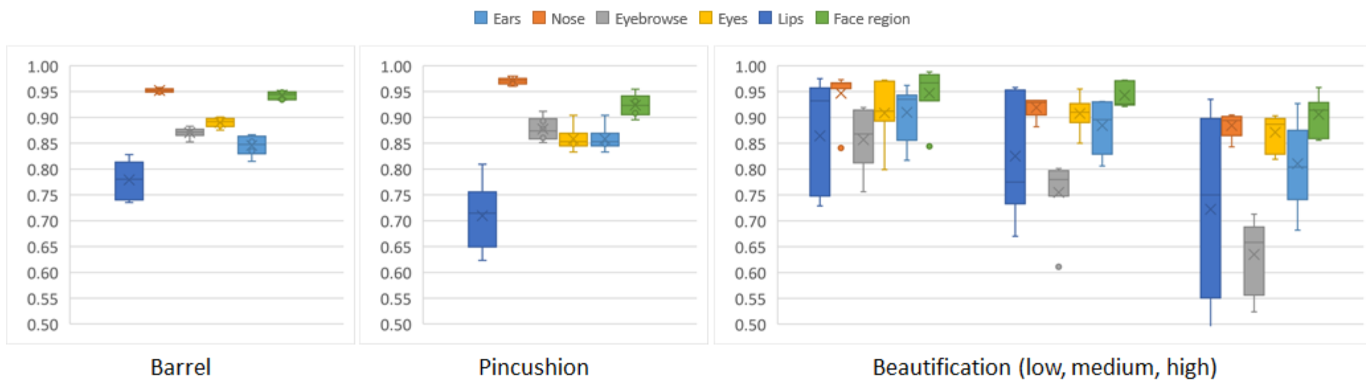

Fig. 6. Box plot of the face components IoU values for images manipulated with geometric distortion (barrel and pincushion) and digital beautification with different levels of intensity.

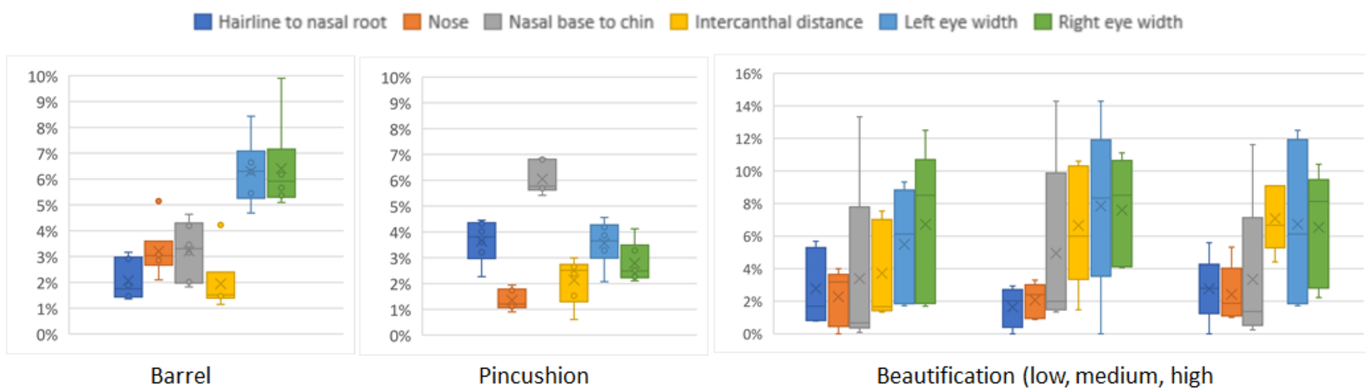

Fig. 7. Box plot of the deviations of facial proportions measurements in images manipulated by geometric distortions and digital beautification with respect to the original image.

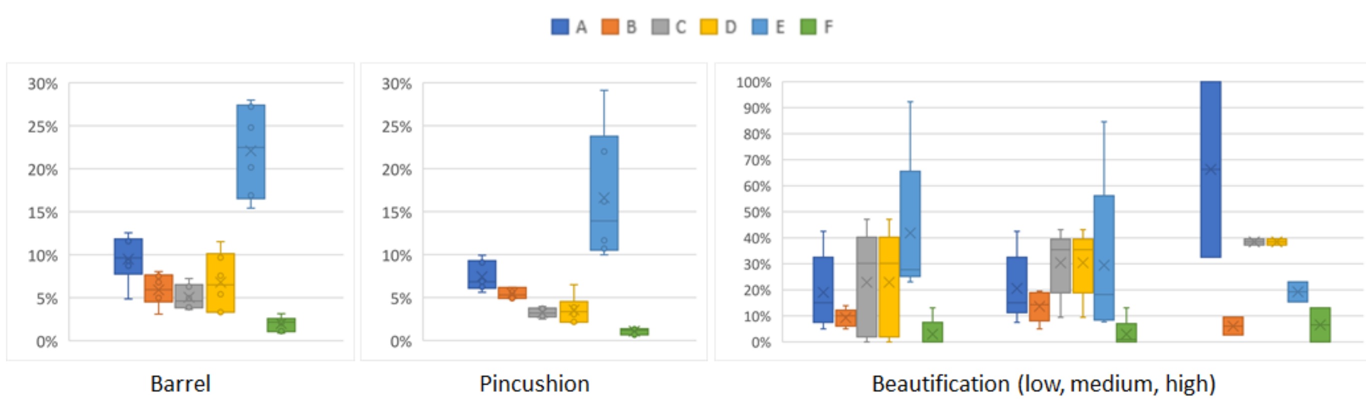

Fig. 8. Box plot of the deviations of eyes measurements 5 in images manipulated by geometric distortions and beautification with respect to the original image.

that, even in this case, some measures are not significantly affected, but the impact on others (e.g. eyes, intercanthal distance, nasal base to chin) is more significant.

The results of Figure 8 analyze the impact of image manipulation on the measures extracted in the eyes region. In particular, the distances analyzed in the graphs are described in Figure 5. The results show that geometric distortions have a lower impact on this region with respect to beautification, but still some measures show noticeable variations (distance E in particular). As clearly visible in the graph, digital beautification strongly affects the eyes region and very significant deviations from the non-manipulated image are observed.

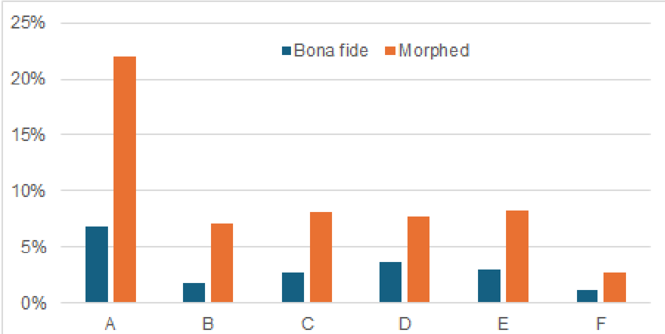

Fig. 9. Deviations of facial proportions measured on bona fide and morphed image pairs.

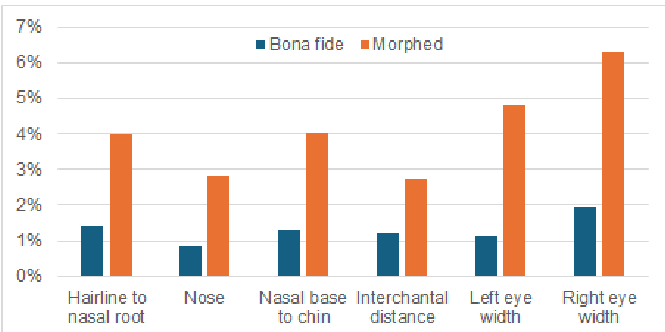

Fig. 10. Deviations of eyes measurements computed on bona fide and morphed image pairs.

A final analysis has been conducted about the effects of morphing on the facial proportions (Figure 9) and the eyes measurements (Figure 10). In particular, the two graphs compare the average deviation of each single measure on a set of bona fide pairs and the same average deviation computed on a set of image pairs where one of the two images is bona fide and the other one is morphed. In general, even if the deviations observed are limited, morphed pairs present a higher deviation than bona fide pairs, suggesting that the implemented measures can to some extent highlight possible discrepancies. In this case we also have to consider that the images used for this test are not always perfectly frontal, thus influencing the measurement outcome.

IV. Conclusions

The tool developed in implements several facial measurements typically used by human examiners for manual facial comparison. The evaluation carried out shows that some measurements are noticeably modified by digital manipulation and their analysis, supported by the tool, can help to detect manipulations. The visual interface is designed to focus the attention of the examiner on specific face regions thus facilitating and making more knowledgeable the human's final decision. We want to stress once more that the aim of the tool is not to influence the human examiner and his decision, but just to help him in the analysis. The tool presents some useful features but is also subject to some limitations, mainly related to the need of a good alignment between the two images to compare. Therefore, in the enrollment scenario, the live capture should be guided to select a properly aligned and frontal image. The tool could also be used in other identity verification scenarios; in that case, the effectiveness of the proposed measures could be affected by some factors such as aging or changes in face appearance; specific and more robust facial metrics and descriptors should be considered in this case to mitigate these challenges. A multi-stage analysis could be considered in this scenario. Initially, some general measures such as IOU and Delaunay triangulation could be considered; In cases where these scores fall below a predefined threshold, additional features available to the human examiner could be utilized to investigate any discrepancies further, enabling a faster real-time inspection process.

We are currently organizing an evaluation of the tool with human examiners to analyze the effectiveness of the proposed analysis and to have a feedback on the usability and on existing functionalities as well as suggestions for future possible improvements. As a further research direction, we will investigate the effectiveness of the facial measurements reported in this work as a contribution to automatic Morphing Attack Detection systems.

V. Acknowledgment

Authors are grateful to Frøy Løvåsdal from the National Police Directorate (Norway) for her insightful suggestions about facial image comparison by human examiners.

References

[1] M. Ferrara, A. Franco, and D. Maltoni, "The magic passport," in *IEEE International Joint Conference on Biometrics, Clearwater, IJCB 2014, FL, USA, September 29 - October 2, 2014*. IEEE, 2014, pp. 1–7. [Online]. Available: https://doi.org/10.1109/BTAS.2014.6996240

[2] M. Ferrara and A. Franco, *Morph Creation and Vulnerability of Face Recognition Systems to Morphing*. Cham: Springer International Publishing, 2022, pp. 117–137.

[3] A. Franco, F. Løvåsdal, and D. Maltoni, "On the human ability in detecting digitally manipulated face images," *2023 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE)*, pp. 531–536, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:267387160

[4] FISWG, "Facial Image Comparison Feature List for Morphological Analysis - version 2.0," Facial Identification Scientific Working Group, Tech. Rep., 09 2018.

[5] D. E. King, "Dlib-ml: A machine learning toolkit," *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.

[6] B. Delaunay *et al.*, "Sur la sphere vide," *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, vol. 7, no. 793-800, pp. 1–2, 1934.

[7] Face parsing. [Online]. Available: https://github.com/zllrunning/face-parsing.pytorch