Portraits of No One: An Interactive Installation

Tiago Martins[0000−0003−2638−237X], João Correia[0000−0001−5562−1996], Sérgio Rebelo[0000−0002−7276−8727], João Bicker[0000−0002−0670−1217], and Penousal Machado[0000−0002−6308−6484]

CISUC, Department of Informatics Engineering,
University of Coimbra
{tiagofm,jncor,srebelo,bicker,machado}@dei.uc.pt

Abstract. Recent developments on artificial intelligence expedited the computational fabrication of visual information, especially photography, with realism and easiness never seen before. In this paper, we present an interactive installation that explores the generation of facial portraits in the borderline between the real and artificial. The presented installation synthesises new human faces by recombining the facial features of its audience and displays them on the walls of the room. The array of faces displayed in the installation space is contaminated with real faces to make people question about the veracity of the portraits they are observing. The photo-realism of the generated faces makes it difficult to distinguish the real portraits from the artificial ones.

Keywords: Interactive Installation, Media Art, Artificial Intelligence, Computer Vision, Image Generation, Computer Graphics

1 Introduction

Portraiture has always been an important form of art. Although its main goal is to immortalise the image of people, it also depicts more than the people traits, thus representing their physical and intellectual possessions [1]. When

Fig. 1. Environment of the installation *Portraits of No One*. Photo © José Paulo Ruas / DGPC 2019.

photography was invented, it soon became the major medium for portraiture by excellence. The portraits, which were formerly an expressive luxury, became affordable for almost everyone [2]. Gradually, portraits became a suitable way to prove the identity of someone and, nowadays, they are part of the identification documents present in everyone's wallet. On the other hand, modern artists started to portrait whatever pleased them in the way they like, aligned with the zeitgeist of the time and the aesthetics of the different art movements [3].

Advances in computer vision enabled the development of more sophisticated detection and recognition tools. This enabled the emergence of artificial intelligence techniques that are able to quickly generate imagery, especially photos, with unprecedented realism. Consequently, these techniques allow the generation of fake content, sometimes used for propaganda, influence, and defamatory purposes. In the case of photos, this is a very important issue, since, by tradition, they are considered proofs of evidence [4].

In this paper, we present the interactive Media Art installation *Portraits of No One* (see figure 1). This installation generates and displays portraits of human faces by recombining the facial features of its audience. The generated portraits exhibit a high level of photo-realism, living in the borderline between the real and artificial. This leads users to question themselves about the veracity of the faces they are seeing. This installation was designed for *Sonae Media Art Award 2019*, being selected as a finalist artwork and, consequently, exhibited at the National Museum of Contemporary Art, in Lisbon, Portugal.

Portraits of No One employs the generative system presented in the work *X-Faces* [5,6] to create face portraits. Therefore, each portrait is created through the recombination of parts of different faces. These parts are retrieved from the faces of the users and assembled using computer vision and computer graphics techniques. Users are invited to give their faces to the installation using a capturing box that is placed in the middle of the installation room. This capturing box also records a sound sample during each face capturing. This enables the installation to maintain an ever-changing audiovisual environment consisting of ephemeral portraits that are projected on the walls combined with intertwined sounds produced by the real people behind the portraits.

Whenever users give their faces to the installation, they will immediately see parts of their faces contaminating the portraits being displayed on the walls. In a way, one could say that the installation feeds on the faces of people who interact with it, as new faces become feedstock that enables the installation to grow. This enables the development of a symbiotic interaction between users and the installation, wherein the identities of the users become part of the artwork. Also, this creates a more engaging experience for users by allowing them to direct the emergence of new artificial identities in the form of portraits. Relationships and interactions between users are also discovered as the faces of different users are blended together to create portraits of no one.

The remainder of this paper is organised as follows. Section 2 presents related work. Section 3 presents a comprehensive description of the installation, both in terms of hardware and software. Finally, section 4 draws final conclusions and points future work.

2 Related Work

In the Media Arts context, portraiture is widely explored using video capturing, face detection techniques, and custom-made software. Some early examples include, *e. g.*, the installations *Video Narcissus* (1987) by Jeffrey Shaw [7], *Interfaces* (1990) by Eduardo Kac [8], or *Solitary* (1992) by Simon Biggs [9].

Recently, this subject has become more popular due to the technological advances in computer vision and, consequently, in face detection. In the scope of this work, we are interested in art installations that use these technologies with three main purposes: *i)* to create more engaging artistic experiences that enable users to feed the artwork; *ii)* to create interfaces that enable users to interact between themselves; and *iii)* to generate artificial human faces.

In some interactive installations, face detection methods are employed to gather content to be used in the creation process of the artwork and, at the same time, to create a more engaging artistic experience for the users. A case in point is *Reface*, an interactive installation developed in 2007 by Golan Levin and Zachary Lieberman, which generates group portraits of its users [10]. The portraits are generated through the combination of videos' excerpts of mouths, eyes, and brows, from different users. Also, the users interact with the artwork using eye blinks. In 2015, Sarah Howorka developed a digital mirror that displays an average face of all people who have been in front of installation [11]. More recently, in 2019, Sameh Al Tawil presented the installation *IDEMixer* that generates live portraits of the participants by blending pairs of faces captured by two different cameras [12]. In the same year, Cecilia Such presented the interactive multimedia performance *I, You, We* [13]. During this performance, users are invited to take a photo. The photographs are, subsequently, rendered and combined between them, and with other visuals, in a live video that runs in the background. Simultaneously, the performer changes the colours of the video through improvisation with a violin and/or a cello.

Some art installations capture the faces of the users to create interfaces that promote novel ways of interaction. An example is the installation *Sharing Faces* developed by Kyle McDonald, in 2013 [14]. This installation behaves as a mirror between two locations: Anyang, in South Korea, and Yamaguchi, in Japan. When users approach the installation, in one of the locations, it tracks their facial expressions and matches them, in real-time, with other expressions of other people who have already stood in front of the installation in the other location. The artists Karen Lancel and Hermen Maat developed the installation *Saving Face*, in various geographical contexts, between 2012 and 2016 [15]. In this artwork, portraits of users are blended and displayed in public screens. The intensity of each blend is determined by the number of times users made gestures of "take care" on their faces. In 2015, Rafael Lozano-Hemmer designed the interactive installation *Level of Confidence* [16], which uses a face recognition system trained to find similarities between a given face and the faces of forty-three disappeared students in a cartel-related kidnapping in Iguala, Mexico, in 2014. Whenever users stand in front of the installation, it captures their faces and unveils the most similar student and how accurate this match is. In a similar way, Antonio Daniele developed the installation *This is not Private*, which enables users to see some of their facial features in the faces of someone else [17]. When users stand in front of a screen watching an interview, their faces slowly merge with the face of the person who is being interviewed. The intensity of the merge is determined by the similarity of the expressions between the user and the interviewee. This installation can recognise six basic facial expressions in the faces of users.

Furthermore, some artworks are created only with the goal of generating artificial portraits. An example is the *Portraits of Imaginary People* series, developed by Mike Tyka, in 2017 [18]. These artificial portraits are generated with a generative adversarial network trained with thousands of photos of human faces. In 2017, João Martinho Moura and Paulo Ferreira-Lopes developed an installation that generates portraits based on the false-positives results of facial detection processes [19]. This way, this installation evolves images in an unsupervised way from a series of erroneous matches. The outputs resemble ghost faces. Physically, the installation uses two computers placed against each one, wherein one computer generates visual noise and the other tries to find faces. In 2018, Mario Klingemann developed the *Neural Glitch* technique [20]. He developed this technique randomly altering, deleting or exchanging the resulting networks of fully trained generative adversarial networks. He developed a series of artworks using *Neural Glitch*, some of them focused on artificial portraiture generation.

3 *Portraits of No One*

Portraits of No One is an interactive installation that generates and presents artificial photo-realistic portraits. This is achieved by automatically capturing and extracting the facial features of the visitors, followed by their recombination in a visually seamless way.

This installation was exhibited in the *Sonae Media Art Award 2019* held at the National Museum of Contemporary Art, in Lisbon, Portugal. In this exhibition, the installation was enclosed in a room specifically made for the purpose, with six meters long by four meters wide. As depicted in figure 2), the space of the installation comprises two main areas: *i)* input area and *ii)* output area. The input area, which is located next to the entrance to the room, contains the capturing box attached to one side of a pillar. The capturing box allows users to capture their faces and this way feed the installation. The output area is located at the back of the room, after the input area, and is surrounded by an immersive video projection of twelve meters wide on three walls. This projection contains an array of portraits that are continuously generated by the installation. Each portrait is displayed with a size similar to that of a real face. This area also includes a speaker attached to the ceiling of the room that reproduces an

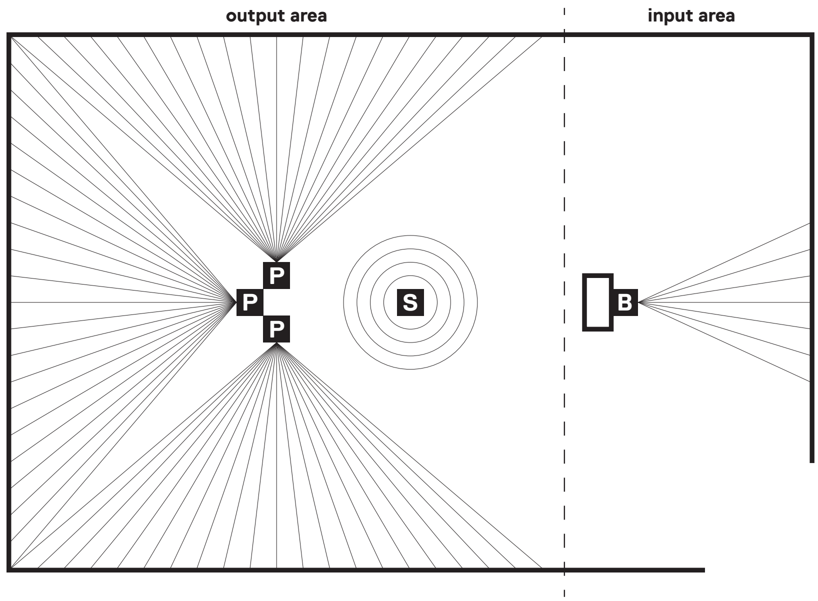

Fig. 2. Floor plan of the installation *Portraits of No One* with the input area (right) and output area (left) along with the main components: capturing box (`B`), video projectors (`P`) and speaker (`S`). Both the video projectors and the speaker are attached to the ceiling of the room. The capturing box is attached to one side of a pillar, with its centre point at a height of 160 centimetres. There is another key component of the installation, the computer, which is hidden in a space next to the room.

intertwined combination of sounds recorded during the interactions of the users with the capturing box.

The user interaction is simple and is described as follows: *i)* the user enters the space of the installation; *ii)* the user approaches the capturing box and observes her/his face on the screen, in a mirror-like fashion; *iii)* if the face is recognised by the system, the button becomes white and blinks fast; otherwise, the button remains red and blinks slowly; *iv)* when the button is white and the reflected image pleases the user, he/she presses the button to capture the face; *v)* a countdown of three seconds begins and then the face is captured; *vi)* the capturing box pauses for a couple of seconds to avoid too frequent captures and the button turns off until a new face can be captured; and, finally, *vii)* the user steps into the output area to see new portraits being generated using her/his facial parts.

In the following subsections, we describe the body (hardware) and behaviour (software) of the presented installation while considering the corresponding inputs and outputs.

3.1 Hardware

The hardware that builds the installation can be divided into input and output hardware. Input hardware is responsible for collecting data from the user, while output hardware is responsible for displaying audiovisual outputs generated by the installation or providing feedback to the users. The core component in this data flow is the computer, which is responsible for connecting all hardware parts of the installation and managing the inputs and outputs. We use a micro-controller to allow the computer to communicate with the push-button through which users trigger the capture of the face. Figure 3 presents the main hardware components and their data flow.

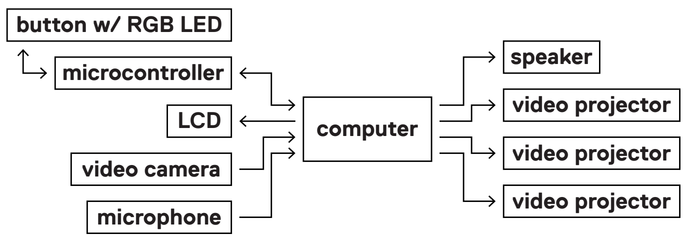

Fig. 3. Diagram of the main hardware components and their data flow. The components on the left are located inside the capturing box, while the ones on the right are attached to the ceiling of the room. The computer is hidden in a space next to the room.

Input Hardware The input hardware obtains data needed for the installation to operate, including images of users' faces and sounds produced by them. This data is retrieved through the capturing box that contains the following input hardware: *i)* a video camera to capture the faces; *ii)* a square LED ring that fits around the camera to provide even illumination with few shadows visible in the captured faces; *iii)* an LCD display to show users a live preview of their faces before capturing; *iv)* a microphone to record some seconds of audio during each interaction; and *v)* a push-button, with an RGB LED, that allows users to trigger the face capturing while providing colour feedback to users.

The video camera and the LED ring are positioned on the top of the display, horizontally centred on the capturing box. The microphone is positioned on one side of the video camera, oriented forward to pick up sounds produced in front of the capturing box. The push-button is placed on the bottom of the screen, also horizontally centred on the capturing box. Figure 4 explains the anatomy of the capturing box.

The capturing box is attached to one side of a pillar, with its centre point at a height of 160 centimetres. The wall immediately in front of the capturing box is covered with black flannel. This way, we are able to create a completely

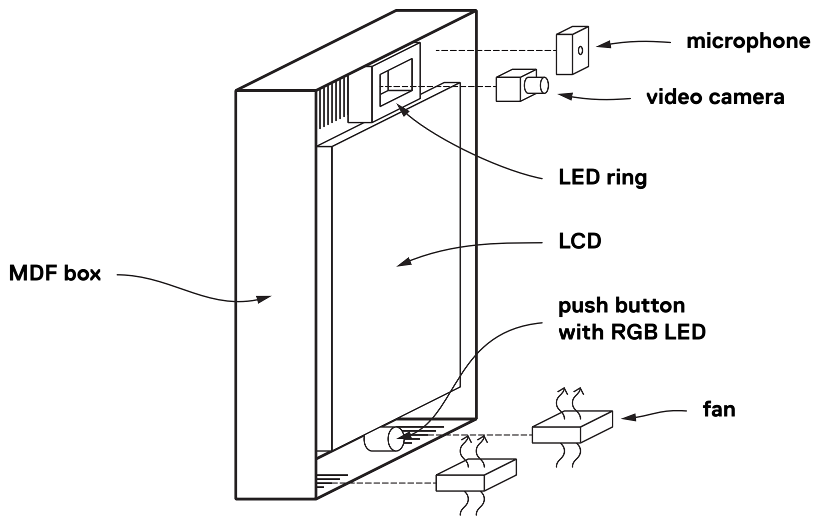

Fig. 4. Schematic of the capturing box showing its anatomy. In addition to the components illustrated in the figure, the box also contains two voltage regulators that allow the adjustment of the brightness of the LED ring and the rotation speed of the fans.

black background in the portraits. This, in combination with the uniform light produced by the LED ring, provides a coherent photographic style across all portraits. The capturing box was entirely designed and developed in-house, allowing us to control all its features. Figure 5 shows a series of snapshots taken during different stages of the making process, from the laser cutting to the assemblage of custom-made parts that build the capturing box.

Output Hardware The output hardware consists of components that allow the installation to express itself to users based on the data it collected from them. This expression is audiovisual and uses images of generated portraits and sounds of different users who gave their faces to the installation. The installation employs three video projectors and one loudspeaker as output hardware to provide an ever-changing audiovisual environment to the users.

The video projectors are attached to the ceiling of the room, each one facing a different wall (see figure 2). The video projectors are strategically positioned and oriented to create a single continuous projection. The result is a surrounding video projection of twelve meters wide, on three walls, with a total image size of 5760 by 1080 pixels. This image size allows the presentation of a large array of portraits with proper quality.

The loudspeaker is also attached to the ceiling of the room, in a central position, facing down. This way, we can fill the room with an ambient sound that complements the installation environment.

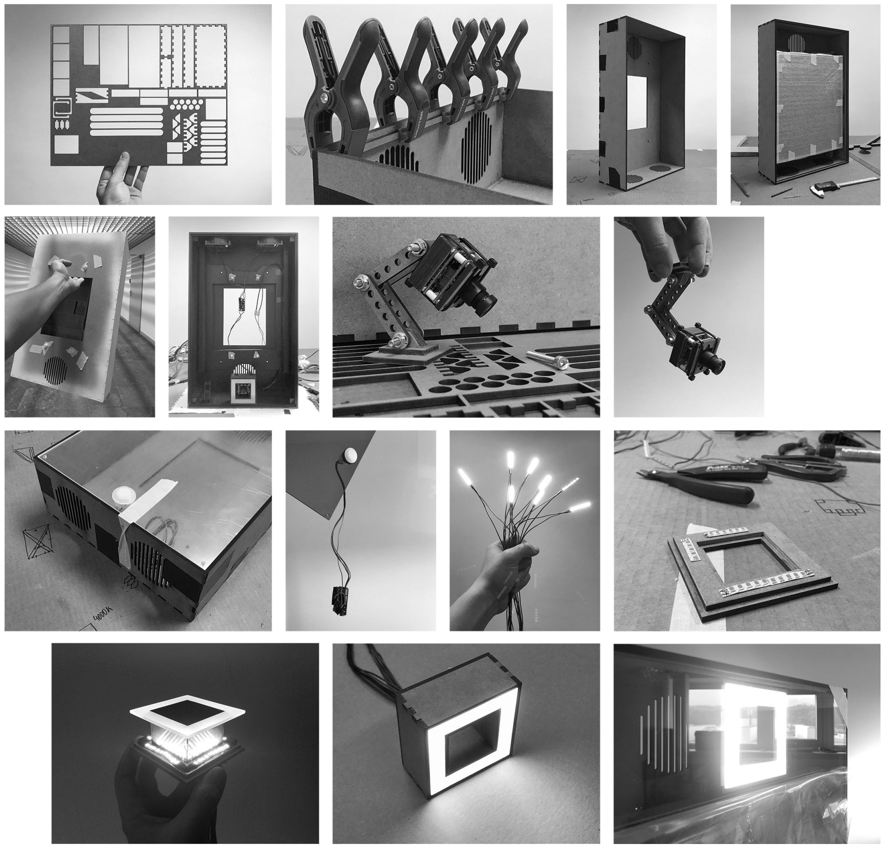

Fig. 5. Snapshots of the making process of the capturing box depicting different stages from the laser cutting to the assemblage of custom-made parts, including the body of the box, the adjustable arm for the video camera, and the LED ring that fits around the camera.

One could say the LCD screen and the RGB LED on the push-button can also be considered as output hardware, since they provide visual feedback to users while they interact with the capturing box. However, we think that it is easier to make this division between input (capturing box) and output hardware (video projectors and loudspeaker).

3.2 Software

The software that runs the installation is mostly implemented in Java, using the open-source libraries Processing, OpenCV and Dlib. We also use Max for sound recording, processing, and reproduction. The micro-controller that establishes the communication between the computer and the push-button is programmed in Arduino.

The installation relies on a computer system that collects users' data using the input hardware, processes this data, and expresses new portraits through the output hardware.

Input Software The input software is essentially the part of the system that manages the capturing box (see figure 6). One of the key features implemented by this part of the system is real-time face detection. We resort off-the-shelf solutions from OpenCV and Dlib libraries to analyse each frame of the camera feed and detect the faces of users, and their facial landmarks, when they stand in front of the capturing box. In the context of this work, we require neutral frontal poses for the portraits. This is not only advantageous to the accurate extraction of the facial features but also to ensure visual coherence between different portraits. To achieve this, we implemented a set of detection filters that determine whether or not a detected face should be considered as valid.

Fig. 6. Photographs of the installation *Portraits of No One* showing someone interacting with the capturing box. On the background of the photographs, one can see walls crowded with faces generated and displayed by the installation system. Photos © José Paulo Ruas / DGPC 2019.

When a valid detection is found, the system draws a rectangle around the face, on the live preview shown on the screen. In situations where more than one valid detection is found, the system focuses on one face based on two preference criteria: *i)* faces horizontally centred in the captured image; and *ii)* faces with large bounds.

The rectangle that is drawn around the detected face indicates the area of the captured image that will be cropped and used to generated new portraits. The aspect ratio of the crop rectangle is always 3:4 (vertical). The size and position of the rectangle is determined in relation to the detected facial landmarks. This way, we ensure the faces present in the portraits are always roughly aligned between them.

In addition to drawing a rectangle around the valid face detection, the system invites the user to capture the face by pressing the push-button. To do so, the system changes the colour of the push-button from red, which indicates that no face is being detected, to white and presents a message on the screen encouraging the user to so.

When the user presses the push-button, the system begins a countdown of three seconds, shown on the top of the screen close to the video camera. When the countdown ends, the system captures the face and saves the related data needed to generate new portraits. At this moment, the system turns off the LED of the push-button, during a couple of seconds, until a new face can be captured. This is part of a delay mechanism that we implemented to avoid bursts of photos.

The data that is recorded when a face is captured includes the image of the face, the sound sample recorded from a few seconds before the face being captured to a few seconds later, the points of the facial landmarks that were calculated during detection, and a timestamp indicating the time at which that face was captured. After absorbing and processing the data, the installation can now feed on the new captured face and synthesise new portraits using the generative process described in the next subsection.

Output Software The output software consists in the part of the system that maintains the audiovisual environment formed by portraits and sounds. This part of the system integrates three modules: the *generator* of artificial portraits based on a set of input faces, the *selector* of faces to generate portraits, and the *producer* of ambient sound.

The *generator* of artificial portraits uses the generative system presented in the work *X-Faces* by Correia et al. in [6,5]. The input of the generator is a set of five faces and their facial landmarks. One face will be our *target* face, *i. e.* the face whose elementary parts (eyebrows, eyes, nose and mouth) will be replaced with parts from other faces. The other four faces will be our *source* faces, *i. e.* the faces that will provide those parts to be blended onto the *target* face. For each face, we calculate the Delaunay triangulation of each face using points based on its landmarks. The triangulation of these points allows us to divide each face image into triangles, which tend to cover corresponding facial features between faces. The next step is to warp each *source* face triangle, *i. e.* the pixels contained in it, to the corresponding *target* face triangle using affine transformations. This results in the *source* faces that are fully aligned with the *target* face. Then, we check if the warped source image part is compatible with the *target*, by calculating the Intersection over Union value of the bounding boxes with a certain threshold. If the value of the Intersection over Union is above the threshold, we consider the part compatible. If they are compatible, we use a seamless cloning algorithm to blend one region of each warped *source* face into the *target* image. The regions of the different facial parts of the warped *source* faces are delimited with masks that are calculated using specific vertexes of the Delaunay triangles. If at least one compatible swap has occurred, the new resulting face is returned. Otherwise, no face is returned. This way, we ensure the seamless swap of facial parts between faces, minimising artefacts and enhancing the realism of the resulting face.

Fig. 7. Photograph of the installation *Portraits of No One* with someone observing the walls crowded with faces. Photo © José Paulo Ruas / DGPC 2019.

The *selector* of faces manages the faces that are used to compose the grid of 225 portraits that are displayed simultaneously on the walls of the installation. This module selects a random set of faces captured from the users and supplies them to the *generator* to create a new portrait using facial features of these faces. This selection process takes time into account the timestamps of the faces, tending to select faces that are more recent. Each portrait being projected on the walls has a lifespan of 10 to 20 seconds. After this lifespan is over, the *selector* selects a new set of faces, asks the *generator* to create another artificial portrait, and replaces the old portrait by the new one. This process is repeated whenever the lifespan of a portrait is over. Once in a while, the *selector* chooses a face at random and displays it. This mechanism contaminates the array of faces displayed in the installation with real faces to make users question about the veracity of the portraits they are observing. Also, the smooth transition between portraits makes users aware of the subtle changes on the walls of the installation that are populated with numerous portraits. Figures 7 and 8 show the environment of the installation with users observing the walls populated with the numerous portraits, real and artificial.

The *producer* of ambient sound combines sounds recorded by the capturing box and reproduces the result using the loudspeaker. The sound is generated by intertwining a total of twelve samples randomly selected from the set of all

Fig. 8. Photograph of the installation *Portraits of No One* with people observing the walls crowded with faces. On the right, one can see the capturing box. Photo © José Paulo Ruas / DGPC 2019.

recorded samples. Most of the interactions occur in silence, so we decided to play a base sound sample, in background, at a volume lower than the main samples. This base sound consists in a sample, in loop, of crowd noise that was previously recorded. The frequency and pitch of the recorded sound samples are attenuated and made homogeneous, digitally, through single-pole low-pass sound filters. Generally, the users inside the room mostly hear imperceptible sounds resulting from the mixing of captured noise. However, sometimes, they can recognise some words or even sentences. This provides a sonic environment wherein the user has the perception that the faces depicted in the numerous portraits are talking between them.

4 Conclusions

We have presented the Media Art interactive installation *Portraits of No One*. This installation generates and displays photo-realistic human portraits by recombining facial parts of its audience. This installation was designed for *Sonae Media Art Award 2019* and exhibited at the National Museum of Contemporary Art, in Lisbon, Portugal.

The installation employs the generative system *X-Faces* [6,5] to create the portraits. Each portrait is generated by the recombination of parts of different real faces. These parts are automatically annotated, retrieved, and assembled using techniques from computer vision and computer graphics. The real faces used in this process are portraits of the installation users.

The space of this installation comprises an input and an output area. In the input area, the users are invited to give their faces to the installation, using a capturing box attached to a pillar on the room. In addition to taking pictures, the capturing box also records a sound sample in each interaction. The capturing box was entirely designed, assembled and developed in-house. In the output area, the audiovisual outputs of the installation are presented. Portraits are projected on the walls while ambient sound is reproduced with a loudspeaker. The is generated by the intertwine of sounds of real people interacting with the capturing box. This creates an ever-changing audiovisual environment, contaminated with chaos and doubt.

Whenever users give their faces to the installation, they will immediately see their facial parts contaminating the portraits that are displayed on the walls. The installation feeds on the information of facial portraits from users who interact with it, allowing it to grow. As the installation takes advantage of users to feed itself and grow, the users' facial parts become part of the artwork. This aspect of the installation creates an engaging experience for users by allowing them to directly participate in the generation of new artificial identities in the form of portraits.

Moreover, the artificial portraits generated by the installation encourage critical thinking about how recent technological advances are changing our relationship with the others and the world. The photo-realism of the generated portraits places the audience in the borderline between the real and artificial. As a result, users tend to question the veracity of the faces that they are seeing. Besides, this opens a window of opportunity to discuss the veracity of the images that we see in other contexts and environments and if we should consider them as unquestionable proofs of the truth.

In the installation, one can also observe varied relationships and interactions between users, as well as between the users and the artificial identities generated by the installation. Most of the users envisage the portraits on the walls as other people, starting to judge and talk about their visual appearance. However, when they recognise themselves, or a familiar face part, in a portrait of no one, they change the behaviour. At this moment, users engender empathy by this artificial person and by the others who, like them, are blended to generate their portraits.

Future work will focus on *i)* set up and test the installation in different locations and *ii)* further improve the installation system according to feedback already obtained.

5 Acknowledgments

The authors express their acknowledgement to Directorate General for Cultural Heritage of Portugal, National Museum of Contemporary Art - Museu do Chiado, and Sonae by commissioning the installation *Portraits of No One*. The authors would also like to express their gratitude to the exhibition production staff and to everyone involved in the development of the installation. This work is partially supported by national funds through the Foundation for Science and Technology (FCT), Portugal, within the scope of the project UID/CEC/00326/2019 and it is based upon work from COST Action CA15140: Improving Applicability of Nature-Inspired Optimisation by Joining Theory and Practice (ImAppNIO). The third author is funded by FCT under the grant SFRH/BD/132728/2017.

References

1. Berger, J., Blomberg, S., Fox, C., Dibb, M., Hollis, R.: Ways of Seeing. Penguin Books Ltd, London, UK (1972)
2. Benjamin, W.: A short history of photography. Screen 13(1), 5–26 (03 1972)
3. West, S.: Portraiture and modernism. In: Oxford History of Art: Portraiture, chap. 8, pp. 187–204. Oxford University Press, Oxford, UK (2004)
4. Sontag, S.: In plato's cave. In: On Photography, chap. 1, pp. 1–19. Farrar, Straus & Giroux, New York, NY, USA (1973)
5. Correia, J., Martins, T., Martins, P., Machado, P.: X-Faces: The eXploit Is Out There. In: Pachet, F., Cardoso, A., Corruble, V., Ghedini, F. (eds.) Proceedings of the Seventh International Conference on Computational Creativity (ICCC 2016). pp. 164–182. Sony CSL Paris, France (2016)
6. Correia, J., Martins, T., Machado, P.: Evolutionary Data Augmentation in Deep Face Detection. In: GECCO 2019 - Proceedings of the 2019 Genetic and Evolutionary Computation Conference. Prague, Czech Republic (2019)
7. Shaw, J.: Video narcissus (1987), https://www.jeffreyshawcompendium.com/portfolio/video-narcissus/, (accessed November 9, 2019)
8. Kac, E.: Interfaces (1990), http://www.ekac.org/sstv.html, (accessed November 10, 2019)
9. Biggs, S.: Solitary (1992), http://littlepig.org.uk/installations/solitary/solitary.htm, (accessed November 10, 2019)
10. Levin, G., Lieberman, Z.: Reface [portrait sequencer] (2007), http://www.flong.com/projects/reface/, (accessed November 8, 2019)
11. Howorka, S.: Average face mirror (2015), http://www.sarahhoworka.at/projects/average-face-mirror, (accessed November 10, 2019)
12. Tawil, S.A.: Idemixer (2019), http://samehaltawil.com/portfolio/idemixer/, (accessed November 10, 2019)
13. Suhr, H.C.: I, you, we: Exploring interactive multimedia performance. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 1147–1148. MM '19, ACM, New York, NY, USA (2019)
14. McDonald, K.: Sharing faces (2013), https://github.com/kylemcdonald/sharingfaces, (accessed November 10, 2019)
15. Lancel, K., Maat, H., Brazier, F.: Saving face: Playful design for social engagement, in public smart city spaces. In: Brooks, A.L., Brooks, E., Sylla, C. (eds.) Interactivity, Game Creation, Design, Learning, and Innovation. pp. 296–305. Springer International Publishing, Cham (2019)
16. Lozano-Hemmer, R.: Level of confidence (2015), http://www.lozano-hemmer.com/artworks/level_of_confidence.php, (accessed November 10, 2019)
17. Daniele, A.: This is not private (2015), http://www.letitbrain.it/letitbrain/?port=this-is-not-private, (accessed November 10, 2019)

18. Tyka, M.: Portraits of imaginary people (2017), http://www.miketyka.com/?s=faces, (accessed November 10, 2019)
19. Moura, J.M., Ferreira-Lopes, P.: Generative face from random data, on 'how computers imagine humans'. In: Proceedings of the 8th International Conference on Digital Arts. pp. 85–91. ARTECH 2017, ACM, New York, NY, USA (2017), http://doi.acm.org/10.1145/3106548.3106605
20. Klingemann, M.: Neural glitch (2018), http://underdestruction.com/2018/10/28/neural-glitch/, (accessed November 10, 2019)