

Special Section on 3DOR2024

From coin to 3D face sculpture portraits in the round of Roman emperors

Umberto Castellani^{a,*}, Riccardo Bartolomioli^a, Giacomo Marchioro^b, Dario Calomino^b^a University of Verona, Department of Computer Science, Verona, 37134, Veneto, Italy^b University of Verona, Department of Cultures and Civilisations, Verona, 37129, Veneto, Italy

ARTICLE INFO

Keywords:

3D face
Morphable model
Model fitting
3D scanning

ABSTRACT

Representing historical figures on visual media has always been a crucial aspect of political communication in the ancient world, as it is in modern society. A great example comes from ancient Rome, when the emperor's portraits were serially replicated on visual media to disseminate his image across the countries ruled by the Romans and to assert the power and authority that he embodied by making him universally recognizable. In particular, one of the most common media through which ancient Romans spread the imperial image was coinage, which showed a bi-dimensional projection of his portrait on the very low relief produced by the impression of the coin-die. In this work, we propose a new method that uses a multi-modal 2D and 3D approach to reconstruct the full portrait in the round of Roman emperors from their images adopted on ancient coins. A well-defined pipeline is introduced from the digitization of coins using 3D scanning techniques to the estimation of the 3D model of the portrait represented by a polygonal mesh. A morphable model trained on real 3D faces is exploited to infer the morphological (i.e., geometric) characteristics of the Roman emperor from the contours extracted from a coin portrait using a model fitting procedure. We present examples of face reconstruction of different emperors from coins produced in Rome as well as in the imperial provinces, which sometimes showed local variations of the official portraits centrally designed.

1. Introduction

The visual representation of historical figures is a research topic of great prominence in the study of the human past for its many implications in the political, economic and cultural life of ancient societies. The Roman Empire was the first state of the ancient world, spanning three continents and lasting for almost five centuries, in which a system was implemented for designing and disseminating the official image of the ruler throughout its entire dominion on a regular basis. The portraits of the emperor and of the members of his family were the most ubiquitous symbols of Roman power: they were displayed in the main public venues and buildings of the cities of the empire so that everyone would be able to recognize them and acknowledge their authority. The serial reproduction of the imperial image is a unique phenomenon of visual duplication which is extensively documented, especially on coinage and in sculpture. The process by which they were designed, produced and transferred began with the creation of common models (or types), which were replicated for use and distribution throughout the empire on different media. The process would have been centrally controlled, with types introduced at different stages of each emperor's reign, possibly marking special occasions and celebrations. However, the relationship between these media (as well as other ones less documented in archaeology, such as paintings, gems and other objects),

and sometimes even their dependency from shared models, is only vaguely understood. This holds especially true in the provinces, where the Roman prototypes that inspired the provincial copies may have sometimes been altered by local artists and adapted to local needs [1].

In this work, we investigate the relation between 3D face sculptures and coins to recover the former from the latter. The use of 3D imaging is particularly aimed at recreating the process by which a sculptor would have modelled a portrait on a bi-dimensional reference. In order to reduce the potential incidence of the artist's subjectivity, we needed to design a different methodology to achieve our goal: an entirely data-driven process. This task is particularly challenging due to the lack of 3D information in the coin portraits, where only an extremely low-relief is available, in comparison to the complete portraits in the round of sculpted statues and busts. Therefore, the problem is highly ill-posed and reliable prior information is required to regularize the solution. An additional hindrance is posed by the process of reconstructing the image of non-rigid subjects such as the human face, which can appear with very different morphological characteristics and facial expressions. We propose a well-defined 3D reconstruction pipeline that combines 3D scanning and geometry processing techniques to recover a full 3D face sculpture of Roman emperors from the 2D portraits depicted on ancient coins. We exploit a data-driven approach using a morphable model for 3D human faces to obtain a simple parametric representation of the face-space. In this fashion, the generation of a plausible

* Corresponding author.

E-mail address: umberto.castellani@univr.it (U. Castellani).

3D face of a Roman emperor can be easily obtained by estimating the correct parameter of the morphable model employing a model fitting procedure from the morphable model to the coin. The proposed pipeline consists in the following main steps: (i) coin digitization, (ii) landmark and contour extraction, (iii) morphable model definition, (iv) model fitting, and (v) hair and beard refinement. Note that step (v) is a proper procedure to fix beard and hair that are not encoded by the morphable model but, most interestingly, can be automatically inferred from the coin relief. Finally, a user interface is designed to tune the facial morphological characteristics that remain uncertain. In this fashion, we allow the potential intervention of archaeologists and specialized graphic designers to make the reconstruction even more plausible and reliable.

The main contribution of this work is twofold:

- to propose a theoretically-based methodology that generates a 3D digital reconstruction of historical figures from ancient coins. This provides archaeologists, art historians and ancient historians with an opportunity to get a better understanding of the image-making process and of an artist's workflow in the production of imperial portraits in different media and formats, especially in relation to the use of one or more shared models.
- to design a novel methodological pipeline that optimizes the combination of automatic procedures based on data-driven approaches with the contribution of users, especially experts in this field of research (but potentially also members of the generic public), to obtain more accurate and credible results.

Our experimental section provides several reconstruction results of different Roman emperors derived from the 3D scans of Roman coins struck in copper alloy in the mint of Rome and of some provincial cities, such as Alexandria in Egypt. The rest of the paper is organized as follows: Section 2 introduces the state of the art related to this work. Section 4 describes the proposed methods by highlighting each step in the pipeline. Section 5 reports several reconstruction results and describes the user interface to refine the output of the automatic procedure. Finally, Section 6 presents the conclusions and discusses the potential for future work.

2. State of the art

Digital techniques in cultural heritage. Digital representation of archaeological artefacts are very important for cultural heritage [2,3]. The modelling and digitization are often combined with learning [4] and geometry processing procedures [5] to solve advanced tasks such as content based 3D object retrieval [6].

More recently, digital acquisition has been exploited for non-rigid shapes, such as the full human body or the human face [7–9]. In [7] a proper modelling tool has been proposed to generate and animate historical characters in a reconstructed environment of archaeological sites such as the case of Pompeii [10]. The accurate reconstruction of detailed morphological characteristics of the human face requires high expertise and manual work. In [11] the 3D face reconstruction has been recovered from the CT scan of a real skull by combining forensics techniques, graphics knowledge and geometry processing. A similar approach has been proposed in [12] for the face sculpture reconstruction of Ramesses II from Egypt. Other approaches estimate the 3D face structure directly from 2D images. In [13] authors proposed the reconstruction of the 3D face of ancient Kabuki drawings. A non-rigid shape registration was proposed to deform a generic 3D face to the input point using Radial Basis Functions.

Morphable models. Face reconstruction in cultural heritage relies on advanced models, namely *morphable models* [8,9,14]. For instance, in [15] a morphable model has been proposed to restore ancient paintings where the reconstruction is exploited to complete and correct the damaged parts. In general, morphable models are largely employed

in computer vision and machine learning for face reconstruction and recognition [16–23]. One of the main aspects is the definition of the model fitting procedure. In [23] only sparse geometric information was considered using hard correspondences constraints for both landmarks and edges. In [22] authors proposed an inequality-constrained optimization strategy as an alternative to standard l_2 regularization approaches. In [21] an efficient face fitting method was defined exploiting linearization of the cost terms that provides a fast closed form solution of the optimization procedure. Other interesting methods proposed the integration between different morphable models for different shape parts (e.g., face and head) [19], the use of morphable models at different resolutions (i.e. multi-resolution morphable model [17]), and the design of a different learning strategy for the estimation of basis face using a dictionary learning framework [18]. More recent methods are based on the advances in deep learning techniques [16,20,24,25]. In [20] authors proposed a GAN network architecture to improve the reconstruction of face details starting from the output of a standard morphable model. A volumetric convolution approach was proposed in [24] to deal with 3D face regression. In [16] a geometric deep learning method was proposed to improve the extraction of 3D features using a spectral-based graph convolution procedure. Open access datasets of 3D faces such as [25] can be used to rely on a sufficiently large sample for the training phase.

Our approach. In this work we focus on ancient coins that is a very interesting research area to exploit information from several historical areas [26]. We focus on accurate morphological reconstruction without the use of appearance properties such as colour and texture [14]. We emphasize the use of information enclosed in the coins along the face and head contours, and coin relief. Our method strongly believes in the morphable model, even though it is not suitable to model beard and hair [8]. Therefore, unlike previous work, we complete the bald and shaved face provided by the morphable model with beard and hairs estimated from a new automatic procedure properly design to work on coins.

3. Historical perspectives

The present study rests within the scope of the broader collaborative project RESP (The Roman Emperor Seen from the Provinces), which investigates the genesis, the replication, and the diversification of Roman imperial portraits in the provinces during the first three centuries of the Empire (c. 31 BC – AD 297) to get a better understanding of this phenomenon.

The project is combining traditional research methodology with digital advances in three-dimensional reproduction especially to reconstruct how the portraits of the same emperor might have originated from a shared model and were reproduced and adapted on different media, with a major focus on the comparison between coins and sculptures. Coins in particular represent an incredible resource for spreading the imperial image and for observing how imperial portraits changed over time: they were serially produced by state mints (mainly Rome) as well as by local mints in the provinces, carrying miniaturized versions of the official portraits designed by the central authority which circulated in every corner of the Roman world [27].

4. Method

In this work, we consider only coins featuring the portrait of the emperor seen in profile. This is the most common representation of the obverse of Roman coins. The profile can be easily outlined, especially along the nose, the forehead, and the neck. The morphological characteristics of this 2D profile are used to guide the generation of the full 3D face. The prior knowledge to complete the unavailable 3D information is referred to the morphable model properly trained on 3D faces of real subjects. Then, the main idea consists in employing a variational approach to fit this morphable model to the face profile observed on



Fig. 1. Hadrian's coin (a) with its corresponding mesh wireframe (b). The mesh consists of 153,646 vertices and 306,589 faces.

the actual coin. Finally, an automatic procedure is introduced to add hair and beard to the face, exploiting the information enclosed in the coin relief.

The proposed method to reconstruct a portrait in the round from coins relies on a processing pipeline that involves the following steps: (i) coin digitization, (ii) landmark and contour extraction, (iii) morphable model definition, (iv) model fitting, and (v) hair and beard refinement.

4.1. Coin digitization

Ancient coins were digitized using 3D scanning techniques. Although the coins' metal surface (mainly dark patina of copper and bronze metal alloys) can prompt problems with light reflection or absorption, the acquisition procedure has produced high-quality digitization of coins with accurate details on the portrait (face, neck, back of the head), while hair and beard are slightly harder to retrieve. Since the colour information is not available, only geometry data have been collected for this task. The coins were scanned using a 3D scanner based on conoscopic holography principle [28]. They were sampled every 50 μm using a ConoPoint-3 range finder with a nominal distance measuring accuracy of about 6 μm . The array of distances produced by the scanning system was later converted into a mesh and exported in STL format. The output results in a high-resolution 3D triangular mesh that is shown in Fig. 1.

4.2. Landmark and contour extraction

From the portrait of the coin, we can easily identify few facial landmarks and the face profile that characterize the main morphological properties of the portrait. In this part of the process, the user needs to manually provide these landmarks and contours using the point-pick tool available in Meshlab.¹ In more detail, we extract:

- **Landmark extraction:** 5 landmarks are extracted on face profile, namely (i) *nasion*, (ii) *pronasale*, (iii) *subnasale*, (iv) *stomion*, and (v) *gnathion* [29]. Moreover, a further important landmark is the *eye*. These points are also identified to the morphable model in a hard one-to-one correspondence. Fig. 2.a shows these landmarks.
- **Head profile:** several points are sampled along the face profile. The *visage profile* is extracted as the segment from the forehead to the neck (Fig. 2.b). Then a further set of 6 segments are extracted namely (i) *front-neck*, (ii) *jaw*, (iii) *upper lip*, (iv) *nose and forehead*, (v) *back-neck*, and (vi) *hair* (Fig. 2.c). Note that the visage profile is the union of jaw, upper lip, nose, and forehead. The same head profile sampling are also manually identified to the morphable model.

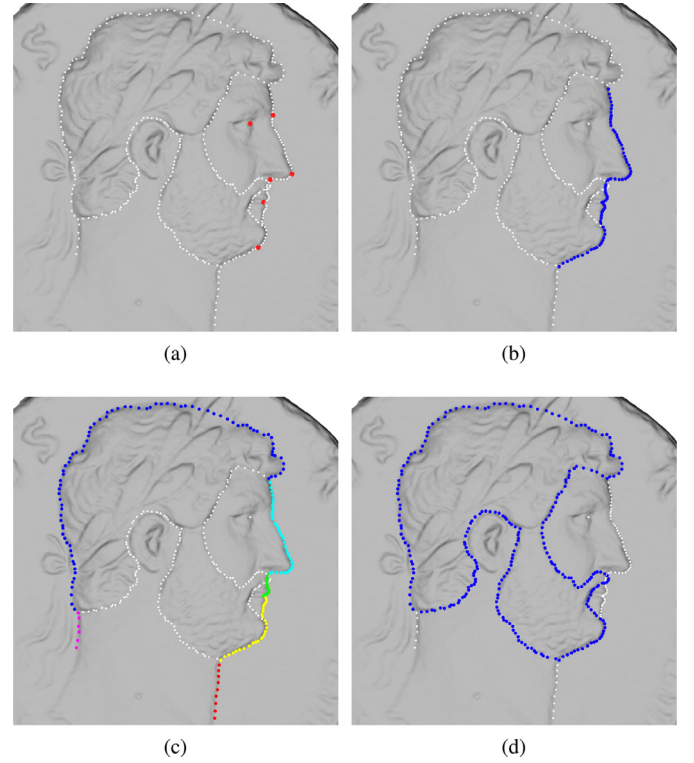


Fig. 2. Points manually selected from the coin and their subdivisions. (a) Landmarks, in red: nasion, pronasale, subnasale, stomion, and gnathion. (b) Visage region, in blue. (c) Head profile regions: front-neck (red), jaw (yellow), upper lip (green), nose and forehead (cyan), back-neck (pink), and hair (blue). Beard and hair contour, in blue.

- **Hair and beard contours:** the *hair* and *beard* contours are extracted to the coin to separate the visible skin part from the rest of the face (Fig. 2.d). Note that morphable model does not deal with hairs and beard. Therefore, we need to design a proper procedure for that.

4.3. Morphable model

The modelling and synthesis of 3D faces is an important task that is usually performed by expert artists. It is a challenging task due to the complexity of the human face that is characterized by a large variation of morphological characteristics. Human perception is very sensitive to these shape properties that enable us to recognize an individual and distinguish him or her from the others. Usually, in the context of face animation, specialized artists adopt a blend-based approach where a fixed set of basis-face are modelled by the artist and then a generic face can be obtained by their linear combination. A similar approach was proposed in the context of face recognition, where the basis-faces are automatically estimated from real examples using a statistical learning framework. This brings the idea of morphable models where each 3D face is encoded by a high dimensional vector and the basis-faces are computed as eigenvectors (or eigenfaces) of this face-space. A data-driven approach is exploited starting from a set of 3D faces of real subjects acquired with 3D scanning techniques. These scans are aligned to a face-template to guarantee the coherence in representing each 3D face by a vector with a fixed dimension. The basis-faces are computed using Principal Component Analysis (PCA) that provides an orthogonal basis.

More formally, given the following vectors:

- the template mesh or mean shape $\mathbf{T} \in \mathbb{R}^{3N}$, where N is the number of vertices,

¹ <https://www.meshlab.net/>

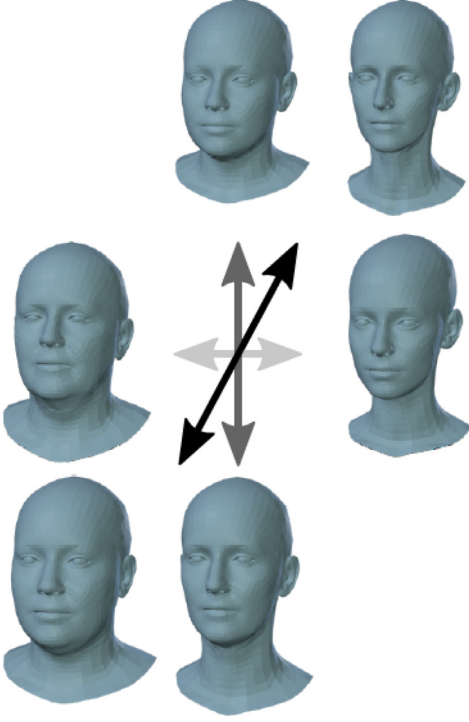


Fig. 3. The face-space. Changing the set of basis parameters $\hat{\beta}$ the morphable model generates different shapes with different face identities.
Source: The image is taken from [9].

- the set of K basis-face $S = [S_1, S_2, \dots, S_K]$, with $S_i \in \mathbb{R}^{3N}$,
- the shape coefficients $\hat{\beta} = [\beta_1, \beta_2, \dots, \beta_K]$,

a new arbitrary face can be generated by the function $M(\hat{\beta}; T, S) \rightarrow \mathbb{R}^{3N}$:

$$M(\hat{\beta}; T, S) = T + \sum_{k=1}^K \beta_k S_k \quad (1)$$

Fig. 3 shows the face-space whereby changing the set of basis parameters $\hat{\beta}$ of the morphable model generates different face identities. In practice, the output vector $M \in \mathbb{R}^{3N}$ is re-arranged in the set of mesh vertices $V \in \mathbb{R}^{3 \times N}$ which is ready to be visualized by standard graphics procedures.

In our work we use FLAME [9] which is a morphable model properly designed for 3D faces. FLAME was trained from 3800 scans of human heads spanning a wide range of ages, ethnicities, and both genders. The generated face is composed by $N = 5023$ vertices, and the available basis-shapes are $K = 300$. Although FLAME involves further parameters to deal with changes of pose and face expression, we mainly consider only shape variation since we are interested in the face identity that appears in neutral expression and almost standard pose. The vertices of FLAME are organized with a fixed triangular connectivity, for which the UV-map is also provided. Fig. 4 shows the template shape T and the UV-map of FLAME.

4.4. Model fitting

The model fitting procedure is designed to estimate the morphable model parameters that generate a new synthetic face with the morphological properties observed on the coin. A variational approach is introduced, defining an *energy* or *error* function whose terms encode a specific behaviour that we impose to the generative framework.

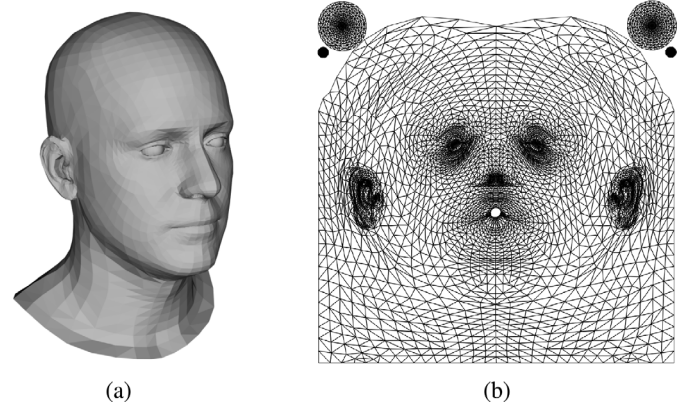


Fig. 4. FLAME template T (a) and its UV-map (b).

Therefore, the fitting process consists in solving an optimization procedure for the estimate of *free* parameters that minimize the given error function. In more details, we define two main fitting steps: (i) coarse fitting, and (ii) fine fitting.

Coarse fitting. Coarse fitting provides a global alignment with scale estimation of the landmarks and face profile extracted from the coin to the morphable model in neutral shape, i.e., $M(\hat{\beta}; T, S)$ with $\hat{\beta} = \mathbf{0}$. Only the rotation matrix $R \in \mathbb{R}^{3 \times 3}$, translation vector $t \in \mathbb{R}^3$, and scale $s \in \mathbb{R}$ are estimated. According to Section 4.2 we define the 5 profile landmarks as $L^c = [L_1^c, \dots, L_5^c]$, with $L_i^c \in \mathbb{R}^3$ and $L^m = [L_1^m, \dots, L_5^m]$, with $L_i^m \in \mathbb{R}^3$ for landmarks extracted respectively on the coin and morphable model. Therefore, we define the following *error* term for landmarks:

$$E_L = \sum_{i=1}^5 \frac{\|s(RL_i^c + t) - L_i^m\|^2}{5} \quad (2)$$

This term imposes a full overlap between landmarks extracted from the coin and those identified on the morphable model in neutral shape. Note that the latter are fixed once for all the coins. Similarly, we define the N points on *visage* profile as $P^c = [P_1^c, \dots, P_N^c]$, and $P^m = [P_1^m, \dots, P_N^m]$ on coin and morphable model respectively. In this case, it is difficult to find hard one-to-one correspondences between points on coin and points on morphable model. Therefore, it is more convenient to encode the *visage* profile on the morphable model as a poly-line composed by the segments of subsequent points, i.e., $l(P^m) = \bigcup_{i=1}^{N-1} (P_i^m - P_{i+1}^m)$. In this fashion, we can compute the projection operator $\mathcal{P}_l^m(P_i^c)$ that project the point P_i^c observed on the coin to the polyline $l(P^m)$ defined from the points on the morphable model.

Therefore, we define the following *error* term for *visage* profile:

$$E_V = \sum_{i=1}^N \frac{\|s(RP_i^c + t) - \mathcal{P}_l^m(s(RP_i^c + t))\|^2}{N} \quad (3)$$

This term implements a point to line distance that avoids a hard point-to-point overlapping constraint.

Finally, the final energy for coarse fitting is obtained by:

$$E_{\text{coarse}} = \omega_L E_L + \omega_V E_V \quad (4)$$

where ω_L and ω_V are parameters manually fixed by the user to tune the importance of involved energy terms. The involved free parameters are estimated as:

$$(\theta^*, v^*, t^*, s^*) = \arg \min_{\theta, v, t, s} E_{\text{coarse}}(\theta, v, t, s) \quad (5)$$

where the rotation matrix R involved in the energy terms is encoded by the rotation angle θ and the unit rotation vector v to guarantee the orthogonality constraint on R during the optimization procedure.

Fine fitting. The fine fitting step optimizes the generation of the 3D portrait in the round, allowing the shape of the morphable model to deform in overlapping along the landmarks and face profiles observed on the coin after the coarse fitting step. To this aim, in the fine fitting, we estimate the K shape parameters of $\hat{\beta}$. Moreover, to improve the global alignment we also optimize the rotation angles of the neck-joint $\hat{\theta} = [\theta_1^{nk}, \theta_2^{nk}, \theta_3^{nk}]$, to deal with examples where the head is slightly moved up or down. Therefore, we use a slightly extended version of the morphable model, i.e., $M(\hat{\beta}, \hat{\theta}; \mathbf{T}, S)$. We also introduce further constraints to allow the alignment of the eyes and to impose a *smoothness* behaviour of the generated shape.

In more details, after the coarse fitting step we move all the points of the coin face profile employing the global alignment procedure $\mathbf{P}_i^{c'} = s^*(\mathbf{R}^*\mathbf{P}_i^c + \mathbf{t}^*)$, where s^* , \mathbf{R}^* , and \mathbf{t}^* have been estimated in the coarse fitting step. According to the semantic profile segmentation defined in Section 4.2 we get the following contours $\mathbf{P}_r^{c'} = [\mathbf{P}_{r,1}^{c'}, \dots, \mathbf{P}_{r,N}^{c'}]$, where $r \in S = \{\text{front-neck, jaw, upper-lip, nose\&forehead, back-neck, hair}\}$. For all these points, we have the corresponding points on the morphable model that we have already jointed in the polyline $l(\mathbf{P}^m)$. Note that now the polyline is extended to the points of the entire face (i.e., visage, neck, and hair), namely, $l(\mathbf{P}^m) = \bigcup_{r=1}^{\#S} \bigcup_{i=1}^{N_r-1} (\mathbf{P}_{r,i}^m - \mathbf{P}_{r,i+1}^m)$.

We define an *energy* term for face profile overlapping:

$$E_{\text{pro}} = \sum_{r \in \{S \setminus \text{hair}\}} \omega_r \sum_{i=1}^{N_r} \frac{\|\mathbf{P}_{r,i}^m - \mathbf{P}_i^{c'}\|^2}{N_r} \quad (6)$$

Note that we excluded the points on hair profile since the hair is not modelled by the morphable model and therefore a strict constraint on the hair overlap may cause the generation of incoherent shape.

On the other hand, to use the information on hair profile we introduce the following *energy* term:

$$E_{\text{hair}} = \sum_{i=1}^{N_r} \frac{\text{clamp}(\|\mathbf{P}_{r,i}^m - \mathbf{P}_i^{c'}\|^2 - \epsilon_h, 0)}{N_r} \quad (7)$$

where only points on hair profile are involved (i.e., $r = \{\text{hair}\}$) and a more relaxed error functions are proposed which is tuned by a tolerance value ϵ_h manually defined by the user.

A further important shape overlapping constraint is based on the eye alignment. Given the eye-point on the coin $\mathbf{L}_e^{c'}$ and the respective left eye \mathbf{L}_e^m and right eye points on the morphable model \mathbf{R}_e^m , we define the following *energy* term:

$$E_{\text{eye}} = \frac{\|\mathbf{L}_e^{(x_e', y_e')} - \mathbf{L}_e^{(x_m, y_m)}\|^2 + \|\mathbf{L}_e^{(x_e', y_e')} - \mathbf{R}_e^{(x_m, y_m)}\|^2}{2} \quad (8)$$

Note that the global alignment brings the face profile along the YZ axis. The eye coordinates is useful to fix only the (x, y) coordinates of the morphable model, since the depth (i.e., the z coordinates) is not encoded in the coin. This is the reason why only the (x, y) components are involved in this energy term. Note also that the eye position on the coin is available only on one side, but it can be easily propagated to both eyes in the morphable model by imposing symmetry.

A last prior is proposed to introduce a smoothness behaviour along the estimated surface. Given the set of mesh vertices $\mathbf{V} \in \mathbb{R}^{3 \times N}$, and given the mesh Laplacian operator \mathbf{L} [30], the smoothness term is defined as:

$$E_{\text{sm}} = \|\mathbf{V}\mathbf{L}\|^2 \quad (9)$$

Note that this term encourages a better propagation of the constraints based on shape overlap also in the areas where the overlap is not defined.

Finally, the final energy for fine fitting is given by:

$$E_{\text{fine}} = E_{\text{pro}} + \omega_{\text{hair}} E_{\text{hair}} + \omega_{\text{eye}} E_{\text{eye}} + \omega_{\text{sm}} E_{\text{sm}} \quad (10)$$

The resulting 3D portrait is generated by the morphable model $M(\hat{\beta}^*, \hat{\theta}^*; \mathbf{T}, S)$, whose free parameters are estimated by solving the following optimization procedure:

$$(\hat{\beta}^*, \hat{\theta}^*) = \arg \min_{\hat{\beta}, \hat{\theta}} E_{\text{fine}}(M(\hat{\beta}, \hat{\theta}; \mathbf{T}, S)) \quad (11)$$

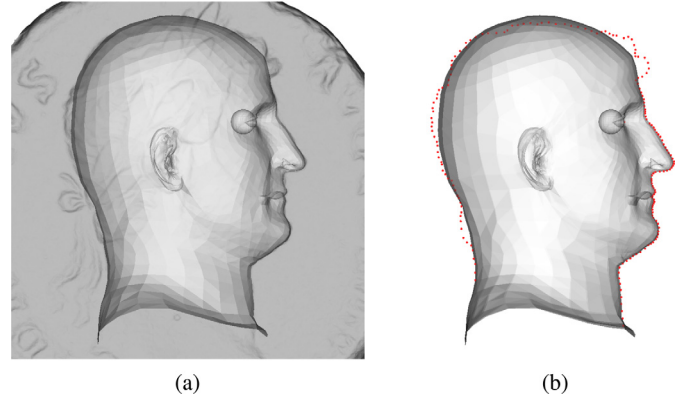


Fig. 5. The optimal model, generated by the automatic process, overlaid on the starting coin (a) and face profile points (b).

Fig. 5 shows an example of automatic estimation of the 3D face sculpture where the morphable model is correctly fitted to the coin.

4.5. Beard and hair

To complete the bald and shaved face provided by the morphable model with beard and hairs, and therefore complete the full 3D face reconstruction, we exploit the possibility of inferring this information from the coin relief. The main idea consists in projecting the morphable model to the coin organized as depth map and transferring these value on the mesh vertices as displacement vectors. The main step of our procedure are:

Step 1 : Grown model generation. In order to transfer the depth values from the coin to the morphable model, we need an accurate mapping between the two representations. The 3D portrait generated from the fine fitting step is not fully suitable, as the alignment was relaxed on hairs for purpose. In this procedure instead, we need a fully aligned version of the morphable model with accurate overlap also on the hair region. To this aim we estimate a new morphable model $M(\hat{\beta}_{\text{gr}}^*, \hat{\theta}_{\text{gr}}^*; \mathbf{T}, S)$

$$(\hat{\beta}_{\text{gr}}^*, \hat{\theta}_{\text{gr}}^*) = \arg \min_{\hat{\beta}, \hat{\theta}} E_{\text{fine}}(M(\hat{\beta}, \hat{\theta}; \mathbf{T}, S)) \quad (12)$$

where tolerance value on hair is neglected (i.e., $\epsilon_h = 0$), and the smoothness constrain is removed (i.e., $\omega_{\text{sm}} = 0$). We call this 3D version the *Grown* model.

As shown in Fig. 6, the output 3D face is still bald and shaved, but its profile is more adhered with the coin, especially on the hair.

Step 2 : Depth image generation for coin. The 3D mesh of the coin (as shown in Figs. 5 and 6) is scaled and moved according to global alignment parameters (s^* , \mathbf{R}^* , \mathbf{t}^*) estimated in the coarse fitting procedure. In this fashion, the coin-mesh and the morphable model are globally aligned. In particular, the XY axes are aligned with the main plane of the coin-mesh. Therefore, the z -coordinate of each coin-mesh vertex can be considered as the depth value computed from the orthographic projection to the XY plane. This provides us a depth image $\text{Depth}_c(i, j)$ of the coin that is aligned with the morphable model. Fig. 7 show the starting mesh and the respective depth image.

Step 3 : Image masking. In our task, we need the relief information only in the regions of the coin enclosing beard and hair. We already extracted these regions in Section 4.2 on the coin mesh, now we can do the same on the depth version.

Fig. 8 shows the mask obtained from the beard and hair contour and the respective regions in the depth image.

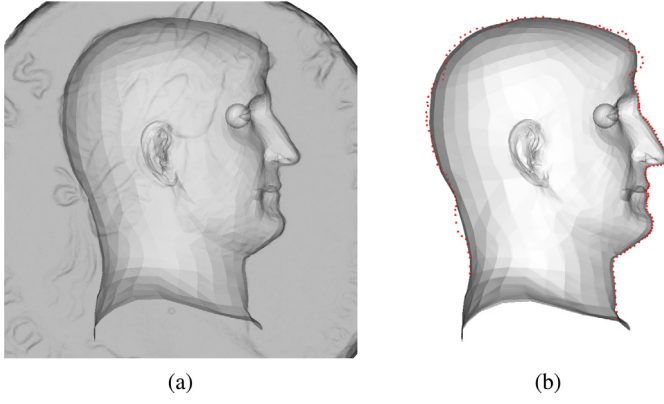


Fig. 6. The grown model, overlaid on the starting coin (a) and face profile points (b). This model overlaps the hair region better than the optimal one, at the cost of worse plausibility.

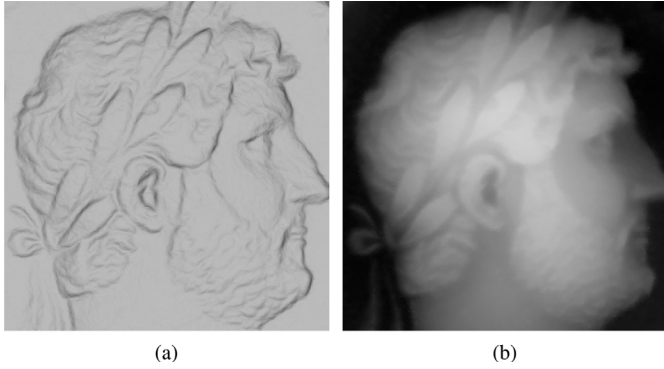


Fig. 7. The starting coin mesh (a) and the derived depth map $Depth_c(i, j)$ (b).

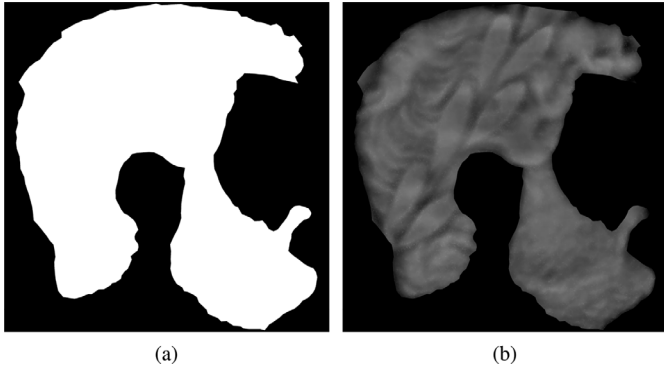


Fig. 8. The mask generated from contours described in Section 4.2, used to determine the hair and beard regions (a), and its application on the processed depth map $Depth_c(i, j)$ (b).

Step 4 : UV-map alignment. We project the mesh of *Grown* model to the image plane of the coin. In this fashion, each projected vertex \mathbf{V}_i of the morphable model can be associated to a pixel $\mathbf{p}_i^{Depth} = (i, j)$. Notably, as explained in Section 4.3 the morphable model is equipped with a fixed UV-map that is derived by the mesh connectivity (see Fig. 4.b). This map provides the mapping between each vertex of the 3D mesh and the pixel coordinates of the morphable model $texture\ Text(u, v)$. Therefore, we can associate to each pixel of the depth map $\mathbf{p}_i^{Depth} = (i, j)$ the corresponding pixel $\mathbf{p}_i^{Text} = (u, v)$ in the texture image, providing

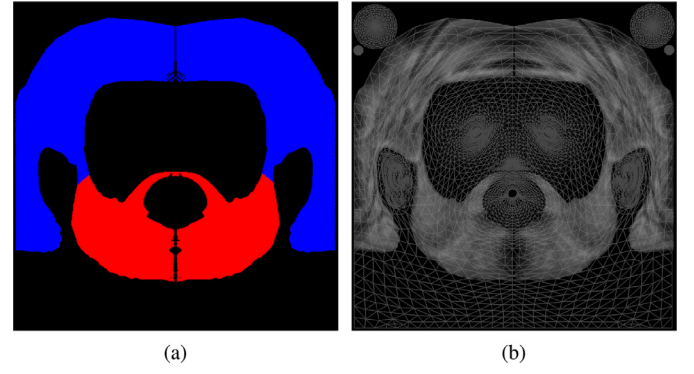


Fig. 9. (a) The projection of the mask using the model's UV-map, with a visual identification of the regions of the beard (red) and hair (blue). (b) The displacement map obtained by projecting the depth values using the model's UV-map (a representation of the model triangles is provided in overlay).

an alignment between the depth and the texture. In this fashion, the information encoded on the depth map can be easily transferred onto the 3D mesh using the UV-map. We exploit this strategy to generate a *displacement* map from the depth image using only the values enclosed in the beard and hairs mask. A simple processing procedure is employed to change the absolute values observed on the coin-depth in a relative value $d \in [0, 1]$ to be transferred to the morphable model. Note that information observed on one side of the coin (the obverse) are transferred on the other side, exploiting the symmetry of the UV-map. Note further that from this mapping we can easily control the resolution of the final mesh employing a mesh subdivision techniques whose displacement weights are taken from the (dense) displacement map.

Fig. 9 shows the result of projecting depth values onto the displacement map by using the model's UV-map. Both sides are drawn by exploiting the symmetry of the model.

Step 5 : Face with beard and hair. Finally, we complete the beard and hairs reconstruction, taking back the 3D portrait derived by $M(\hat{\beta}^*, \hat{\theta}^*; \mathbf{T}, \mathbf{S})$ of the fine fitting. Then we adjust the coordinates of each mesh vertex \mathbf{V}_i according to the displacement value d_i encoded in the UV-map. In more detail, we compute

$$\mathbf{V}_i^{bh} = \mathbf{V}_i + d_i \cdot \mathbf{n}_i (M_i^b \cdot \omega_b + M_i^h \cdot \omega_h) \quad \forall i = [1, \dots, N] \quad (13)$$

where \mathbf{n}_i is the unit normal vector of vertex i and M_i^b, M_i^h is the masking value to select the vertices of beard and hairs respectively. The parameters ω_b and ω_h are manually tuned by the user to attenuate or emphasize the beard and hair effect. Further standard post-processing procedures are employed to smooth some artefacts, especially in the transition regions between skin and hairs, or in the vertices of the morphable model that were not enclosed in the UV-map.

Fig. 10 show the 3D face estimate before and after beard and hairs reconstruction.

5. Experiments

In this section, we report several experiments on 3D full face reconstruction from different ancient coins. At first, we show the results using a fully automatic procedure, then we introduce our interactive tool to improve and customize the final models according to more subjective expectations. The ancient coins are from two major museum collections² and refer to different Roman Emperors in different periods and regions where the coinage may have changed [31–33]. Also, the coin preservation is different among the involved samples.

² the British Museum Department of Coins and Medals and the Museo di Castelvecchio in Verona.

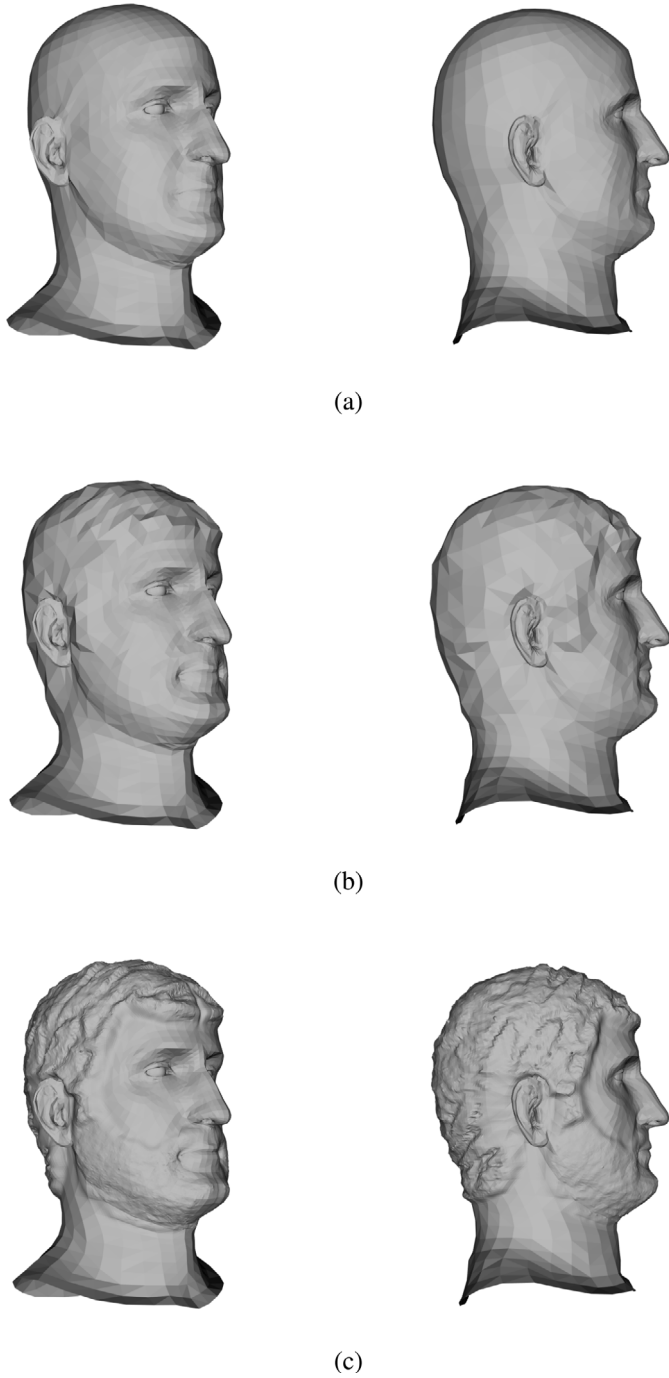


Fig. 10. Hair and beard generation. Bald and shaved model generated by the automatic process (a). The same model with hair and beard applied from the generated displacement map (b). A denser model with hair and beard applied from the same displacement map (c).

5.1. Automatic results

We evaluate the possibility of our method to propose a fully automatic method that provides results exploiting the information encoded on the coin, reducing as much as possible the user intervention. The proposed reconstruction pipeline required the setting of some fixed parameters that are involved in the energy terms described in Section 4.4. We kept the same parameters $\omega_i = 1.0$ for all the experiments. This setting was easy to find since the energy terms are not in conflict one to each other. All the generated 3D models are rendered with a generic

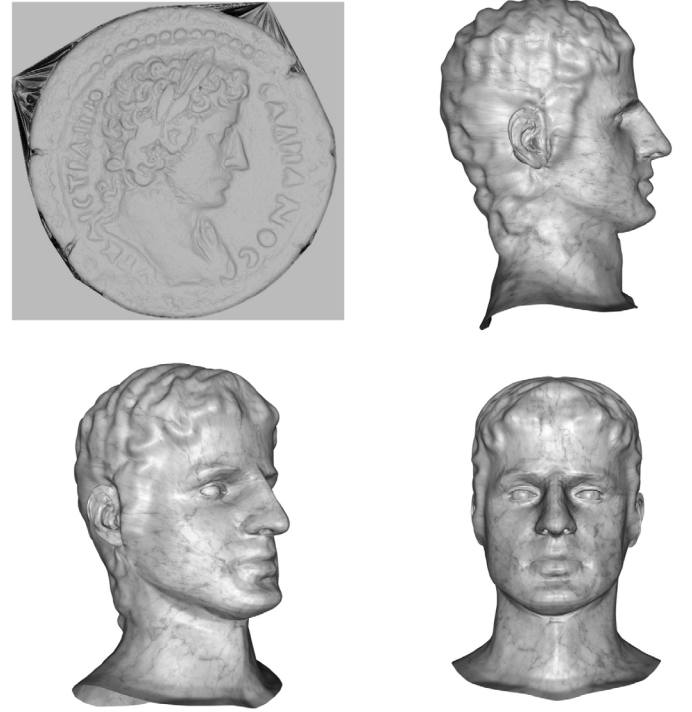


Fig. 11. Museum of Castelvecchio, Verona, Inv. 71976. Bronze, Roman provincial coin of Hadrian, mint of Alexandria in Egypt.: AD 117/118 Catalogue №2502, AE, 34 mm, 19.63 g.

marble texture. For each example, we show the 3D scan of the original coin and three views of the 3D model in the round.

At first, we focused on three different coins of the same Roman Emperor: Hadrian (AD 117-138). Fig. 11 shows an example of a young Hadrian (around 41 years old) portrayed on a coin of Alexandria in Egypt. The specimen is well-preserved and provides accurate details also in the coin relief. The reconstruction is quite accurate, especially along the face profile. The portrait is beardless, and the hair is well-defined.

A second version of Hadrian in a more mature age (around 57 years old) is shown in Fig. 12. This coin too comes from Alexandria in the province of Egypt. The coin is in poorer condition and the portrait is slightly damaged in the beard region. The result is satisfactory and coherent with the emperor's facial features that appear on the coin.

A third version of Hadrian is shown in Fig. 13. This time the coin was minted in Rome, only a few years after the one in Fig. 12, when the emperor was in his early sixties. The coin is well-preserved with excellent portrait details. The model is a credible reconstruction, both in the face profile and in the beard and hair.

These three models generated from three different coins (produced in two different mints) of the same emperor show great similarities between each other, even if the level of preservation is variable. Another three models are shown in Fig. 14, Fig. 15, and Fig. 16 for the emperors Antonio Pius (AD 138-161), Marcus Aurelius and Carausius respectively: the first two were generated from the 3D scans of bronze coins struck in Asia Minor (modern Turkey), the third from a silver coin minted in Britannia (modern UK). Although the coins' condition is not always good, all the results provide a valid reconstruction of 3D models in the round of the emperors' portrait as seen on the coins. The model that depicts Carausius is particularly noteworthy, not only because of the peculiar look of the emperor compared to the other two, but especially because there is no extant sculpture in the round that portrays him.

The method has been implemented in Python using PyTorch library. The optimization procedure is solved with the available numerical

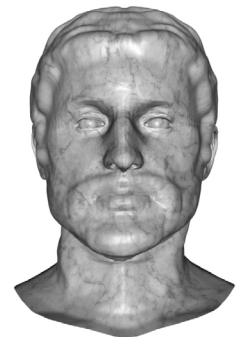
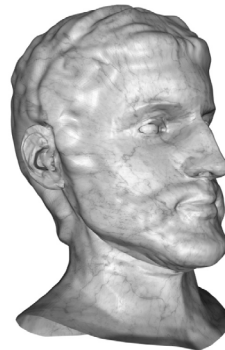
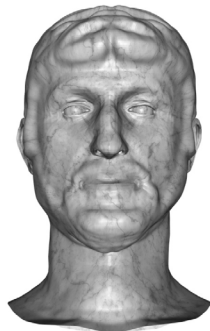
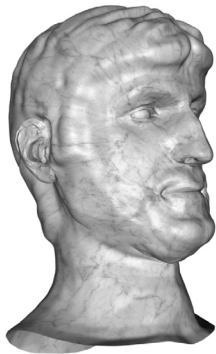
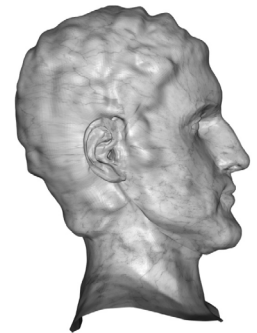
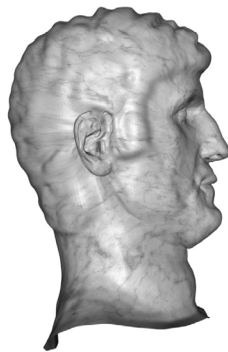


Fig. 12. Museum of Castelvechio, Verona, Inv. 71996. Bronze, Roman provincial coin of Hadrian, mint of Alexandria in Egypt: AD 133/134. Catalogue №2545, AE, 35 mm, 25.46 g.

Fig. 14. Museum of Castelvechio, Verona, Inv. 70341. Bronze, roman provincial coin of Antoninus Pius, mint of Pergamum in the province of Asia: c. AD 144-158. Catalogue №1452, AE, 21 mm, 5.8 g.

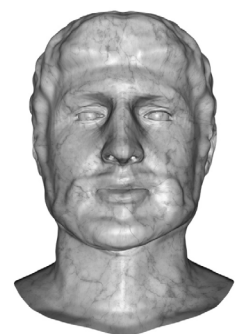
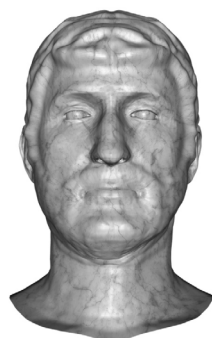
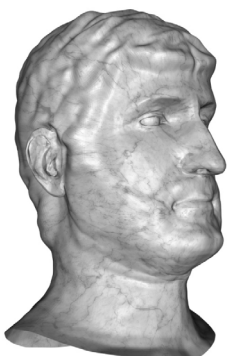
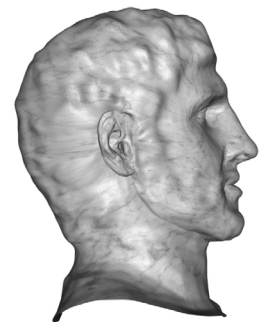
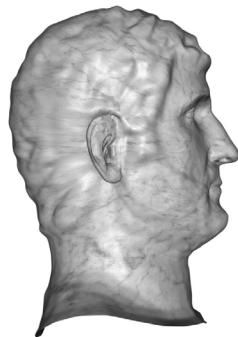


Fig. 13. British Museum, London, Inv. 1872,0709.568. Brass sestertius of Hadrian, mint of Rome: AD 134-138. BMC 1723, AE 33 mm, 25.48 g.

Fig. 15. Museum of Castelvechio, Verona, Inv. 72399. Bronze, Roman provincial coin of Marcus Aurelius Caesar, mint of Amastris in Pontus and Bithynia: AD 144-161. Catalogue №1321, AE, 18 mm, 3.78 g.

solvers using automatic differentiation. Experiments are evaluated on a PC Intel® Core™ i7-10750H CPU @ 2.60 GHz 6 core with GPU NVIDIA

GTX 1650 4 GB. The computation time is less than 10 min for the entire reconstruction pipeline.

Table 1

Ablation study results: each row represents a reconstruction task, performed on the Hadrian coin (Fig. 13), in which a single ω_r is manually set to 0. The last row represents a regular reconstruction. Each column reports the error E_r measured for the correspondent region r , and the highest value is highlighted. One can notice that where the ω_r is null, the error grows significantly in the correspondent region, leading to a less accurate reconstruction, as shown in Fig. 17.

$\omega_r = 0$	E_{eye}	$E_{\text{front-neck}}$	E_{jaw}	$E_{\text{upper-lip}}$	$E_{\text{nose\&forehead}}$	E_{hair}	$E_{\text{back-neck}}$	$E_{\text{smoothing}}$
Eye	1.18e-4	2.60e-7	1.53e-7	2.88e-8	1.19e-7	3.19e-9	7.76e-8	2.40e-6
Front-neck	1.17e-10	2.48e-4	1.06e-7	1.31e-8	1.20e-7	1.14e-9	7.16e-8	2.41e-6
Jaw	2.74e-10	2.78e-7	4.79e-5	4.61e-8	1.49e-7	5.86e-9	7.97e-8	2.47e-6
Upper-lip	2.58e-10	2.96e-7	2.13e-7	4.07e-6	1.51e-7	5.47e-9	7.56e-8	2.48e-6
Nose&forehead	1.31e-10	2.94e-7	1.85e-7	3.64e-8	1.36e-5	2.43e-10	7.33e-8	2.48e-6
Hair	3.55e-10	2.58e-7	2.07e-7	3.54e-8	1.58e-7	3.06e-6	7.39e-8	2.49e-6
Back-neck	2.73e-10	2.79e-7	2.57e-7	9.82e-8	1.57e-7	4.51e-9	2.61e-5	2.45e-6
-	3.02e-10	2.75e-7	2.10e-7	4.84e-8	1.70e-7	6.30e-9	7.82e-8	2.48e-6



Fig. 16. British Museum, London, Inv. R.3506. Silver denarius of Carausius, uncertain mint in Britannia: AD 287. RIC 596: AR, 19 mm, 5.10 g.

5.2. Ablation study

We propose an ablation study to evaluate the importance of the involved energy terms. We focus on the face sculpture reconstruction of Hadrian coin in Fig. 13. We propose several automatic generative procedures where in turn an energy term is neglected, i.e., its weight is set to zero. In particular, we consider the following weights of face contours ω_r with $r \in \{\text{front-neck, jaw, upper-lip, nose\&forehead, back-neck}\}$, the weight for hair ω_{hair} , and the weight for eyes ω_{eye} . In Table 1 we report the error costs of each energy term for each reconstruction case. Note that, as expected, when a weight is set to zero its corresponding error term increases significantly. We show the reconstructed results in Fig. 17. It is visually clear that reconstructed face sculpture changes drastically especially when facial contours are not considered. For instance an incoherent nose reconstruction is visible in Fig. 17.f, and a wrong chin and mouth is depicted in Fig. 17.d.

5.3. Interactive modelling

In the previous section, we have shown that automatic procedure is quite satisfactory. However, in some circumstances it will be interesting to modify the automatic results to provide a more personalized version

of the 3D face considering further priors and information that can be evaluated by experts. To this aim, we developed an interactive application that provides the user some semantic parameters to change the results. The idea is to create a post-processing interface, where users can manipulate aspects of the morphable model that cannot be handled explicitly by the automatic process. In practice, this application performs further optimization steps on the estimation of the shape morphable model parameters $\hat{\beta}^*$ defined from the automatic procedure by exploiting a new energy function composed of new energy terms. In more details, we propose the following constraints:

Eye distance. Since the coin shows only one profile of the face, the distance between the two eyes cannot be inferred. Given the vertices of the left V_{le} and right eye V_{re} of the morphable model, we introduce the following energy term:

$$E_{\text{dist}} = \left| \|V_{\text{le}} - V_{\text{re}}\| - l_{\text{dist}} \right| \quad (14)$$

where l_{dist} is a free parameter.

Eye depth. The depth of eyes is very important to characterize the human face. We allow the tuning of this morphological aspect. Given the z component V_e^z of eye coordinates (the same for both eye), we define the following energy term:

$$E_{\text{depth}} = \left| V_e^z - l_{\text{depth}} \right| \quad (15)$$

where l_{depth} is a free parameter.

Head width. The entire width of the face can be modified. Let be σ_x the variance of V_x coordinates. We define the following energy term:

$$E_{\text{width}} = \left| \sigma_x - l_{\text{width}} \right| \quad (16)$$

where l_{width} is a free parameter. Therefore, the final modified 3D face is obtained by optimizing the shape parameters $\hat{\beta}^{\text{final}}$ according to the following total energy function:

$$E_{\text{tot}} = E_{\text{dist}} + E_{\text{depth}} + E_{\text{width}} \quad (17)$$

The parameters l_{dist} , l_{depth} , and l_{width} are manually set by the user using the user interface.

Fig. 18 shows examples of the interactive application in action on the emperor Hadrian. Slide bars for the tuning of parameters are on either sides. The coin from which the profile was extracted is displayed above alongside the 3D model in the round generated by the automatic procedure. Then we show a further version, in which the head width and the distance between the eyes have been modified. This function is extremely handy and user-friendly and it shows how efficiently a portrait can be modified by altering its main morphological features. Fig. 19 reports a further interactive experience on emperor Marcus Aurelius Caesar.

The code for both 3D face sculpture reconstruction and interactive refinement is public available.³

³ <https://github.com/GeoVr-organization/RESP3D-generator>

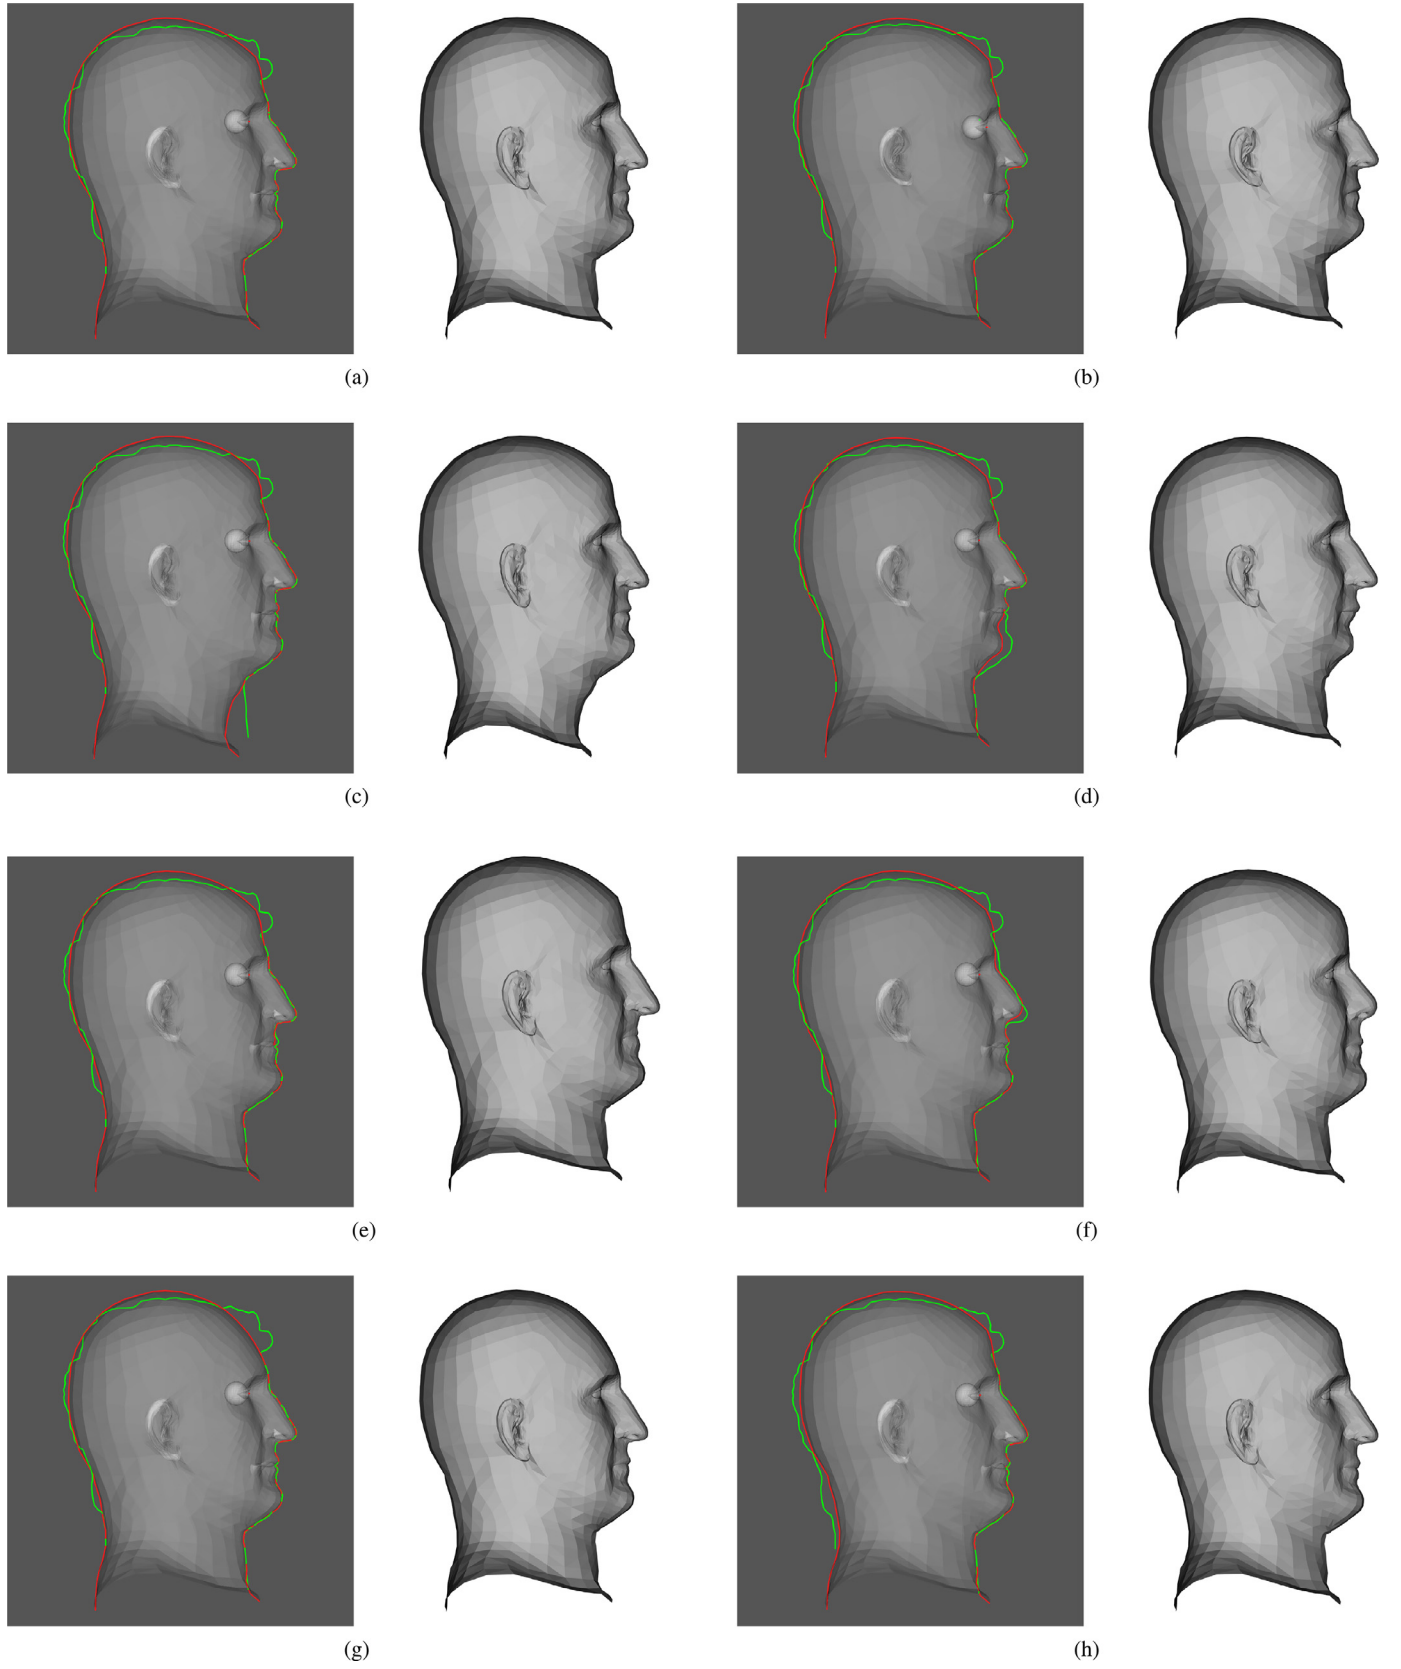


Fig. 17. Ablation study graphical results. Each pair of images represents a reconstruction task, performed on the Hadrian coin (Fig. 13), in which a single ω_r is manually set to 0: on the left, the coin and mesh line profile are highlighted (green and red respectively); on the right, the correspondent full mesh. (a) the regular reconstruction, (b) $\omega_{eye} = 0$, (c) $\omega_{front-neck} = 0$, (d) $\omega_{jaw} = 0$, (e) $\omega_{upper-lip} = 0$, (f) $\omega_{nose\&forehead} = 0$, (g) $\omega_{hair} = 0$ and (h) $\omega_{back-neck} = 0$.

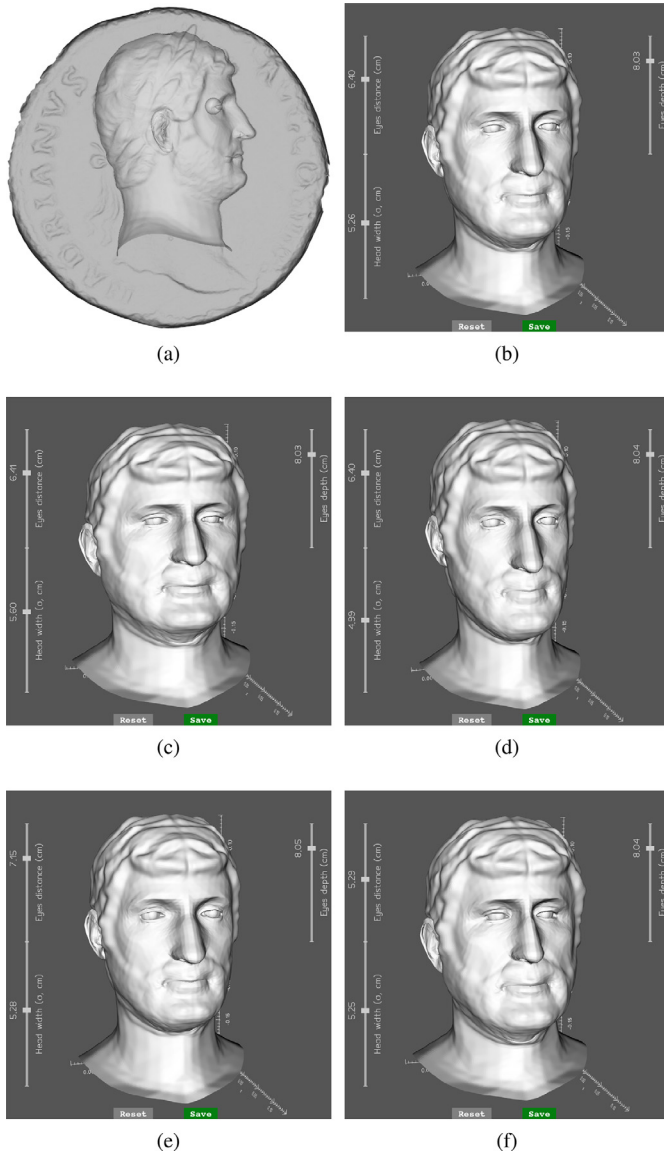


Fig. 18. (a) The automatic result compared to the input coin, for the emperor Hadrian. Starting from automatic result (b), users can change some aspects of the model such as the head width ((c) larger, (d) smaller) or the interocular distance ((e) larger, (f) smaller), by editing specific sliders.

6. Discussion and conclusions

In this paper, we proposed a new method for the estimation of a 3D model of a portrait in the round derived from the relief profile of a portrait stamped on ancient coins. This work is very relevant to the study of the image-making process of historical figures in the ancient world, especially of Roman emperors, whose portraits were produced serially and disseminated on a global scale in different formats and on different media. We designed a well-defined 3D reconstruction pipeline based on the automation of the entire workflow and the full integration between the data recorded from the ancient artefacts (i.e., the coins) and the ones based on a morphable model of the human face. Our work includes several models generated through this process which document the versatility of the proposed method applied to different coins from different regions and periods. We also exploited a refined version of the automatic procedure's output by adding an interactive application that enables the user to modify the 3D model by adopting a semantic modelling approach. We have shown that this tool can

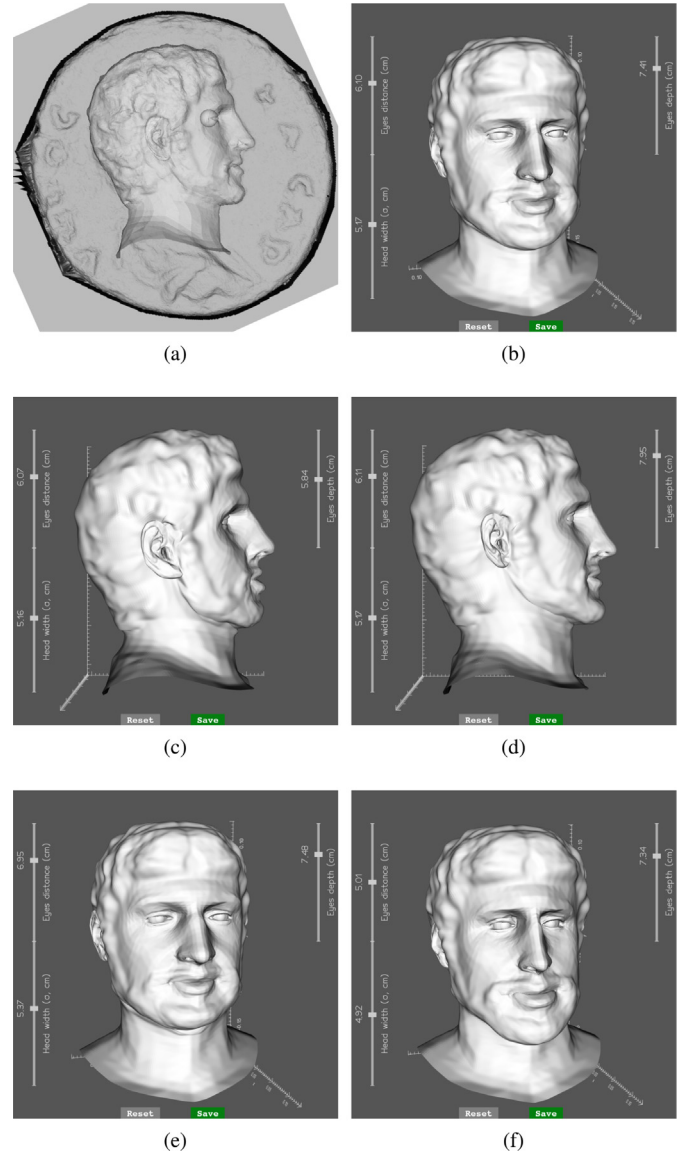


Fig. 19. (a) The automatic result compared to the input coin, for the emperor Marcus Aurelius Caesar. Starting from automatic result (b), users can change some aspects of the model such as the eyes' depth ((c) inner, (d) outer) or the interocular distance ((e) larger, (f) smaller), by editing specific sliders.

very effectively produce alternative versions of the automated model that maintain its original facial morphology, avoiding the creation of degenerate shapes.

Limitations and future work. The main limitation of the proposed method is the need for manual extraction of landmarks and face contours on the coins. Automatic procedure are not trivial since most of the available methods are based on 2D textured images for frontal view. Conversely, in our work, the task is more challenging due to the absence of colour information and the lack of numerous samples. Moreover, in many coins the state of preservation is reduced, and therefore the facial landmarks are not easy to be detected from the relief, especially in the presence of dense beard. The most recent and promising deep learning methods can be exploited, but a large dataset of annotated samples is required for the learning procedure. To this aim, a dataset of synthetic coins can be generated using recent methods on the relief simulation from 3D models such as [34].

The interactive post-processing phase can also be improved, since only few morphological features have been involved in the current

version. Many further options can be considered to allow the user to engage more directly with the model, especially scholars specialized in the study of the ancient world who aim to undertake focused research on portraiture. Finally, a better evaluation of the proposed work can be carried out by designing a proper user study that assesses both automatic results and interactive procedures.

CRedit authorship contribution statement

Umberto Castellani: Writing – review & editing, Writing – original draft, Supervision. **Riccardo Bartolomioli:** Software, Investigation. **Giacomo Marchioro:** Validation, Software, Formal analysis, Data curation. **Dario Calomino:** Writing – original draft, Validation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Umberto Castellani reports financial support was provided by University of Verona. Dario Calomino reports financial support was provided by University of Verona. Riccardo Bartolomioli reports financial support was provided by University of Verona. Giacomo Marchioro reports financial support was provided by University of Verona. Umberto Castellani reports a relationship with University of Verona that includes: employment. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors are unable or have chosen not to specify which data has been used.

Acknowledgements

This paper is published within the project RESP (The Roman Emperor Seen From the Provinces), which has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 101002763). The authors wish to thank Richard Abdy, curator in the Department of Coins and Medals at the British Museum, and Antonella Arzone, curator in the Museum of Castelvecchio in Verona, for allowing the 3D scanning of coins in the respective collections. The scans in the British Museum were performed by Mike Donnelly, Paul Wilson and Mark Williams from the WMG at Warwick University (using a Nikon ModelMaker H120 scan head mounted on a portable MCx25+ scanning arm); in the Castelvecchio Museum they were carried out by Giacomo Marchioro (with a conoscopic holography microprofilometer developed at the OpDaTeCH laboratory of the University of Verona lead by Claudia Daffara).

References

- [1] Fittschen K. The portraits of Roman emperors and their families: Controversial positions and unsolved problems. 2011.
- [2] Pintus R, Pal K, Yang Y, Weyrich T, Gobbetti E, Rushmeier H. A survey of geometric analysis in cultural heritage. *Comput Graph Forum* 2016;35(1):4–31.
- [3] Elkhuisen WS, Callewaert TWJ, Leonhardt E, Vandivere A, Song Y, Pont SC, et al. Comparison of three 3D scanning techniques for paintings, as applied to Vermeer's 'Girl with a Pearl Earring'. *Herit Sci* 2019;7:2050–7445.
- [4] Duda RO, Hart PE, Stork DG. Pattern classification. 2nd ed.. Wiley; 2001.
- [5] Botsch M, Kobbelt L, Pauly M, Alliez P, Levy B. Polygon mesh processing. CRC Press; 2010. URL: <https://books.google.it/books?id=AuXqBgAAQBAJ>.
- [6] Sipiran I, Lazo P, Lopez C, Jimenez M, Bagewadi N, Bustos B, et al. SHREC 2021: Retrieval of cultural heritage objects. *Comput Graph* 2021.
- [7] Yazli N, Baka E, Thalmann N, Kaplanidi D, Partarakis N, Karuzaki E, et al. Modeling craftspeople for cultural heritage: A case study. *Comput Animat Virtual Worlds* 2022;33:373–89. <http://dx.doi.org/10.1002/cav.2075>.
- [8] Egger B, Smith WAP, Tewari A, Wuhrer S, Zollhoefer M, Beeler T, et al. 3D morphable face models—Past, present, and future. *ACM Trans Graph* 2020;39(5).
- [9] Li T, Bolkart T, Black MJ, Li H, Romero J. Learning a model of facial shape and expression from 4D scans. *ACM Trans Graph, (Proc. SIGGRAPH Asia)* 2017;36(6):194:1–194:17. <http://dx.doi.org/10.1145/3130800.3130813>.
- [10] Maim J, Haegler S, Yersin B, Mueller P, Thalmann D, Gool LV. Populating ancient Pompeii with crowds of virtual Romans. In: Arnold D, Niccolucci F, Chalmers A, editors. VAST: International symposium on virtual reality, archaeology and intelligent cultural heritage. The Eurographics Association; 2007, p. 109–16. <http://dx.doi.org/10.2312/VAST/VAST07/109-116>.
- [11] Abate AF, Nappi M, Ricciardi S, Sabatino G. FACES: 3D Facial reConstruction from anciEnt Skulls using content based image retrieval. *J Vis Lang Comput* 2004;15:373–89.
- [12] Wilkinson CM, Saleem SN, Liu CYJ, Roughley M. Revealing the face of rameses II through computed tomography, digital 3D facial reconstruction and computer-generated imagery. *J Archaeol Sci* 2023;160:105884.
- [13] Xu W, Akama R, Tanaka HT. 3D face modeling from ancient Kabuki drawings. In: Mudge M, Ryan N, Scopigno R, editors. The 6th international symposium on virtual reality, archaeology and cultural heritage VAST. The Eurographics Association; 2005, p. 59–65. <http://dx.doi.org/10.2312/VAST/VAST05/059-065>.
- [14] Blanz V, Vetter T. A morphable model for the synthesis of 3D faces. In: Proceedings of the 26th annual conference on computer graphics and interactive techniques. USA: ACM Press/Addison-Wesley Publishing Co.; 1999, p. 187–94.
- [15] Lanitis A, Stylianou G. e-Restoration of faces appearing in cultural heritage artefacts. In: 2009 15th international conference on virtual systems and multimedia. 2009, p. 15–20.
- [16] Xu H, Zhao Z, Cao Y, Chen C, Ge H, Liu Z. 3D face reconstruction using a spectral-based graph convolution encoder. In: Companion proceedings of the ACM on web conference 2024. New York, NY, USA: Association for Computing Machinery; 2024, p. 633–6.
- [17] Huber P, Hu G, Tena R, Mortazavian P, Koppen P, Christmas W, et al. A multiresolution 3D morphable face model and fitting framework. In: 11th international joint conference on computer vision, imaging and computer graphics theory and applications. 2016, p. 79–86.
- [18] Ferrari C, Lisanti G, Berretti S, Bimbo AD. A dictionary learning-based 3D morphable shape model. *IEEE Trans Multimed* 2017;19(12):2666–79.
- [19] Ploumpis S, Wang H, Pears N, Smith WAP, Zafeiriou S. Combining 3D morphable models: A large scale face-and-head model. In: 2019 IEEE/CVF conference on computer vision and pattern recognition. 2019, p. 10926–35. <http://dx.doi.org/10.1109/CVPR.2019.01119>.
- [20] Galteri L, Ferrari C, Lisanti G, Berretti S, Del Bimbo A. Deep 3D morphable model refinement via progressive growing of conditional Generative Adversarial Networks. *Comput Vis Image Underst* 2019;185:31–42.
- [21] Hu G, Yan F, Kittler J, Christmas W, Chan CH, Feng Z, et al. Efficient 3D morphable face model fitting. *Pattern Recognit* 2017;67:366–79.
- [22] Sariyanidi E, Zampella CJ, Schultz RT, Tunç B. Inequality-constrained 3D morphable face model fitting. *IEEE Trans Pattern Anal Mach Intell* 2024;46(2):1305–18. <http://dx.doi.org/10.1109/TPAMI.2023.3334948>.
- [23] Bas A, Smith WAP, Bolkart T, Wuhrer S. Fitting a 3D morphable model to edges: A comparison between hard and soft correspondences. In: Chen C-S, Lu J, Ma K-K, editors. Computer vision – ACCV 2016 workshops. Cham: Springer International Publishing; 2017, p. 377–91.
- [24] Jackson AS, Bulat A, Argyriou V, Tzimiropoulos G. Large pose 3D face reconstruction from a single image via direct volumetric CNN regression. *Int Conf Comput Vis* 2017.
- [25] Bai H, Kang D, Zhang H, Pan J, Bao L. FFHQ-UV: Normalized facial UV-texture dataset for 3D face reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2023, p. 362–71.
- [26] Calomino D, Bologna F, Wilson P, Donnelly M, Williams M. Imaging Hadrian in Britain between coinage and sculpture: A new digital approach to the study of Roman imperial portraiture. *Britannia* 2023;54:251–74. <http://dx.doi.org/10.1017/S00068113X23000387>.
- [27] Burnett A. Coinage in the Roman world. 1986.
- [28] Gaburro N, Marchioro G, Daffara C. A versatile optical profilometer based on conoscopic holography sensors for acquisition of specular and diffusive surfaces in artworks. In: Optics for arts, architecture, and archaeology VI, vol. 10331. SPIE; 2017, p. 48–56.
- [29] Horta R, Teixeira S, Nascimento R, Silva A, Amarante JM. The freestyle facial artery perforator flap for reconstruction of simultaneous periorbital and cheek defects. *J Craniofac Surg* 2016;27:1.
- [30] Lévy B, Zhang HR. Spectral mesh processing. In: ACM SIGGRAPH 2010 courses. New York, NY, USA: Association for Computing Machinery; 2010, p. 1–312. <http://dx.doi.org/10.1145/1837101.1837109>.
- [31] Mattingly H, Carson RAG, Hill PV. British Museum Coins of the Roman Empire. 1–6, London, UK; 1923–1962.
- [32] Mattingly H, Sydenham EA. The Roman Imperial Coinage. 1–10, London, UK; 1923–1994.
- [33] Arzone A, Cappiotti F. Sylloge Nummorum Graecorum Italia. Civici Musei d'Arte, Verona. Roma, IT; 2017.
- [34] Ji Z, Feng W, Sun X, Qin F, Wang Y, Zhang YW, et al. ReliefNet: Fast bas-relief generation from 3D scenes. *Comput Aided Des* 2021;130:102928.