Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/58830/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Wu, J. ORCID: <https://orcid.org/0000-0001-5123-9861>, Martin, R. R., Rosin, P. L. ORCID: <https://orcid.org/0000-0002-4965-3884>, Sun, X. ORCID: <https://orcid.org/0000-0002-6114-0766>, Lai, Y. ORCID: <https://orcid.org/0000-0002-2094-5680>, Liu, Y. - H. and Wallraven, C. 2014. Use of non-photorealistic rendering and photometric stereo in making bas-reliefs from photographs. Graphical Models 76 (4) , pp. 202-213. 10.1016/j.gmod.2014.02.002 file

Publishers page: <http://dx.doi.org/10.1016/j.gmod.2014.02.002>
<<http://dx.doi.org/10.1016/j.gmod.2014.02.002>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Use of Non-Photorealistic Rendering and Photometric Stereo in Making Bas-reliefs from Photographs

J. Wu^a, R.R. Martin^a, P.L. Rosin^a, X.-F. Sun^a, Y.-K. Lai^a, Y.-H. Liu^b, C. Wallraven^c

^a*School of Computer Science & Informatics, Cardiff University, Cardiff CF23 6PS, UK*

^b*Department of Computer Science, Aberystwyth University, Aberystwyth SY23 3DB, UK*

^c*Brain and Cognitive Engineering, Korea University, Seoul, South Korea*

Abstract

Automatic bas-relief generation from 2D photographs potentially has applications to coinage, commemorative medals and souvenirs. However current methods are not yet ready for real use in industry due to insufficient artistic effect, noticeable distortion, and unbalanced contrast. We previously proposed a shape-from-shading (SFS) based method to automatically generate bas-reliefs from single frontal photographs of human faces; however, suppression of unwanted details remained a problem. Here, our experimental results show how incorporating non-photorealistic rendering (NPR) into our previous framework enables us to both suppress unwanted detail, and yet also emphasise important detail. We have also experimented with an alternative approach to recovering relief shape, using photometric stereo instead of SFS for surface orientation estimation. This can effectively reduce the computational time.

Keywords: bas-relief, non-photorealistic rendering, shape-from-shading, photometric stereo.

Email addresses: J.Wu@cs.cf.ac.uk (J. Wu), ralph@cs.cf.ac.uk (R.R. Martin), Paul.Rosin@cs.cf.ac.uk (P.L. Rosin), Xianfang.Sun@cs.cf.ac.uk (X.-F. Sun), Yukun.Lai@cs.cf.ac.uk (Y.-K. Lai), yyl@aber.ac.uk (Y.-H. Liu), wallraven@korea.ac.kr (C. Wallraven)

1. Introduction

Bas-reliefs have been used for centuries in art and design, especially as portraits on coins. Manual production of bas-reliefs requires considerable artistic skill, and is also time consuming. Recent research in computer aided design has considered automatic bas-relief generation from 3D scenes [1, 2, 3, 4] and 2D images [5, 6, 7]. As images are much cheaper and easier to capture than 3D models, starting from the former potentially has much wider application. However, it is difficult, if not impossible, to formalize the knowledge and experience of sculptors who create bas-reliefs by hand, posing great challenges in automating the process by computer. It is even more challenging when starting from 2D images due to the extra 3D shape recovery step, which adds further complication.

In our previous work [7], we reported a method to generate bas-reliefs from frontal human face photographs using the technique of shape from shading (SFS). Our two-step pipeline combined techniques from neural networks, image relighting, and SFS. The results preserved salient features reasonably well, but were affected by noise arising from unwanted details. When sculptors create bas-relief portraiture by hand, the result is stylised rather than faithful. They enhance salient shapes, simplify texture features, and suppress unwanted detail in other areas. Motivated by how sculptors abstract appearance, in this paper, we show this can be achieved by use of a non-photorealistic rendering (NPR) method [8, 9] specifically designed for abstraction of face images. The addition of this abstraction step significantly improves quality compared to our previous results.

SFS is a standard technique for recovering 3D shape from a single image of an object by making use of the shading information in the image in conjunction with a known or estimated lighting direction and assumptions about material properties. An alternative approach is to use photometric stereo [10], which requires two or more images taken under different (known) lighting directions, from a single fixed viewpoint. In our previous work [7], we relight the image to a standard direction (instead of estimating the original lighting). The same image relighting technique also enables us to produce multiple images, allowing the use of photometric stereo instead of SFS for shape recovery. We later show how this is more efficient.

In summary, we show how NPR can be combined into our previous method [7] and how photometric stereo can be used instead of SFS. The resulting contributions of this paper are as follows. Firstly, and specifically, the

addition of an NPR step to our earlier pipeline effectively removes unwanted clutter and detail, while enhancing salient features, significantly improving the visual quality of the results. Secondly, and more generally, the output of NPR techniques has previously been intended for direct human viewing. This paper shows that NPR can be used in a different way, applying it as an *intermediate step* in a modeling pipeline. Thirdly, the use of photometric stereo effectively reduces the computational time in the bas relief production pipeline, compared to use of SFS.

2. Related Work

A survey of automatic relief generation can be found in [11]. The majority of recent bas-relief generation methods [1, 2, 3, 4] start with a *depth-map* of a 3D scene, and selectively compress depths to create a bas-relief surface. Relatively fewer works [5, 6, 7] have considered using 2D *images* as input.

A two-level approach considering low frequency shape and high frequency detail was used in [6] to restore brick and stone reliefs from images taken as rubbings. The authors also demonstrated that their method could restore relief effects from objects in photographs. However due to the assumption of brick and stone relief priors, the method is best suited for objects made of homogeneous materials with relatively little texture and low albedo. As the authors note, results when applied to portrait photographs fail to correctly model locally concave (e.g. eyes) and convex (e.g. nose) areas.

The approach in [5] aimed to create relief surfaces that approximate desired images under known directional lighting. They applied a modified SFS method with height constraints. The authors noted that the integrability constraint enforced by SFS constrains the radiance for each element of a recovered surface. To use this observation, they associated each pixel with not just one, but several, surface elements. Relaxing the integrability constraint leads to reduced control over global shape, which in turn makes the appearance of the final bas-relief sensitive to changes in viewing direction and illumination. Such reliefs are unsuited to applications such as coinage, which may be viewed under varying conditions.

Our previous work [7] observed that images of real bas-reliefs *do not* approximate images of the corresponding 3D objects. The goal of a bas-relief generation algorithm should not be to create a relief whose image matches an input image—this is an unattainable and undesirable goal, given the very different geometry—but to make a bas-relief that looks like a high quality

bas-relief produced by a human sculptor. Based on this observation, in [7] we proposed a two-step approach to recover bas-reliefs from face photographs. An off-line process learns a model of the mapping from a face photograph to an image of a corresponding bas-relief. Given a new face photograph, it is first mapped to an appropriate *image* of a corresponding bas-relief. Subsequently, the bas-relief *surface* is recovered using feature preserving SFS [12]. The results produced have a global geometry consistent with expectations (i.e. a flattened version of the original shape), giving them a stable appearance under differing viewing direction and illumination. However, while preserving salient features, the results are affected by noise and other unwanted detail. In this paper, we incorporate non-photorealistic rendering into our previous method, which both reduces noise and excess detail, and at the same time enhances salient features.

Non-photorealistic rendering (NPR) uses a variety of techniques to render 3D models, 2D images or video in many different ways. For instance, different media can be simulated (e.g. watercolour, charcoal, crayon), as well as different painterly styles (e.g. impressionistic, pointillist, cubist) [13]. Most NPR work has focused on the production of art-like rendering [14]. However, a few papers have considered applications of NPR. For instance, [15] integrated NPR into an interior design CAD system to render rooms in a watercolour style intended to be more visually appealing to customers than photorealistic renderings. Several methods have been developed to generate exploded views of 3D models, illustrating components of complex objects and their manner of assembly [16]. In other technical illustration, NPR rendering is used to provide a clearer representation of shape, structure, and material than a traditional photorealistic rendering [17]: NPR techniques provide a level of abstraction, removing unwanted clutter and detail while retaining or emphasising salient features. However, in all of these applications, the NPR output is the final goal, intended to be viewed directly by the user. In contrast, we use NPR output as an *intermediate* result which is input to subsequent processing steps.

3. Framework

The framework of our method is shown in Figure 1, with the red dashed-line boxes indicating where NPR is incorporated into our previous pipeline [7]. For the alternative approach where the relief shape is recovered using photometric stereo, the pipeline remains the same except that the blue dashed

Figure 1: Framework for generating bas-reliefs from photographs: (a) offline learning of map from face images to relief images, (b) online relief generation process. Red dashed boxes indicate NPR steps added to our previous pipeline. Blue dashed box indicates where to replace SFS with photometric stereo.

box in Figure 1 is replaced by photometric stereo processing.

There are two major components in the framework. An offline process is used to *learn* the relationship between a frontal monochrome image of a 3D human face (after NPR processing for abstraction) and a corresponding frontal image of a 3D bas-relief of that face. Then, given a new face photo (converted to monochrome if necessary), we first use NPR to abstract its appearance, and then apply the rest of the pipeline to get the final bas-relief surface. A detailed description of the NPR method we use [8, 9] is given in Section 4. The alternative approach of using photometric stereo is given in Section 5. We now briefly summarise the other steps for completeness; full details can be found in [7].

3.1. Mapping Face Images to Bas-Relief Images

During offline training, computer generated 2D frontal images of a reference 3D face model and a corresponding 3D bas-relief model are produced, using the same reflectance model and lighting conditions in each case; the bas relief model is generated algorithmically from the reference model [4]. Having rendered the 3D face model, NPR is used to abstract the resulting

Figure 2: Two pairs of training images. Left: model image and bas-relief image from lighting direction $(1,1,1)$. Right: training images from lighting direction $(-1,1,1)$.

image. A feedforward neural network [18] is used to learn the mapping between the abstracted 3D model image and the relief image. The mappings are learned under the assumption that the intensity of each foreground pixel in the bas-relief image is determined by the intensities in a neighborhood around the same pixel in the corresponding 3D model image.

In fact, learning is done twice, for two pairs of training images with lighting directions $(1, 1, 1)$ (from above right) and $(-1, 1, 1)$ (from above left), as shown in Figure 2; this helps later to overcome problems such as those due to shadows.

3.2. Image Relighting

Given a new face photograph, we relight it from the same two lighting directions as above, using the *quotient image* technique [19] and a bootstrap image set from Yale Face Database B [20]; images are coarsely aligned using the tip of the nose and the centers of the eyes. Figure 3(a-c) shows an example of applying image relighting using this bootstrap set. The original photograph was taken under ambient light, then processed using our NPR method which clearly cleaned the image and enhanced the features. The mid-tone used in NPR is adjusted to that of the bootstrap set to ensure comparability of image intensities. Results after relighting are shown in Figure 3(c); they exhibit realistic lighting effects with few artifacts.

3.3. Generating Bas-relief Images

After relighting, the appropriate learned function is applied to each relit image, according to the lighting direction used; this generates an image which estimates what the corresponding bas-relief should look like. See Figure 3(d).

Figure 3: An example of generated bas-relief images. (a) Original image; (b) NPR processed; (c) Images relit from directions $(\pm 1, 1, 1)$; (d) Corresponding bas-relief images.

Figure 4: An example of average surface.

3.4. Bas-relief Surface Recovery and Combination

To determine the geometry of the corresponding relief surface, we apply a modified version of a structure-preserving shape-from-shading method [12] to each bas-relief image. A small linear scaling is finally applied to each of the recovered surfaces to achieve the exact relief height desired (we used a target image width to depth ratio of 50:1). The pair of relief surfaces is averaged to get the final output as shown in Figure 4; such averaging helps to reduce artifacts due to shadows, highlights and lighting direction.

4. NPR method

We now explain in more detail the new NPR process used to abstract facial appearance. It closely follows that used in [8, 9] for producing pictures, but omits the stage that performs region extraction. To handle features at multiple resolutions, a three level pyramid is created from the input image, and NPR is applied at each level, the results being combined by averaging. The resulting image is rather flat in tone, with lines of various strengths and scales. A simple 50 : 50 average is made of this NPR image and the original input image to provide a result with more continuous tone variation: see Figure 3(b). The main stages of the NPR pipeline are illustrated in Figure 5 and explained next; full details can be found in [8, 9].

Figure 5: The NPR pipeline. Each dashed-line box is similar, and corresponds to one of the three pyramid levels. The results are averaged with the input image to get the final results shown on the right.

4.1. Line Extraction

Following [21], lines are extracted by applying a difference of Gaussians (DoG) filter. To reduce fragmentation, the shape and orientation of the filter kernel is adaptively determined using local image characteristics. The DoG is applied in a direction orthogonal to the intensity gradient, as determined by local edge flow; this is estimated by bilateral filtering of edge tangents.

4.2. Deleting connected components

The extracted lines can potentially retain clutter. To remove unwanted detail, an opening and closing set-morphological operation [22] is applied. Unlike standard mathematical morphology, fixed structuring elements are replaced by connected components in the image. For example, closing sets each pixel to the lowest threshold value at which it belongs to a connected background component with an area greater than or equal to the given threshold.

4.3. Eye Correction

The human visual system is sensitive to faces, and the eye regions are especially salient. We therefore take special care to avoid errors around the eyes: fragmentation of lines there can lead to their removal at the filtering stage; superposition of lines can also lead to intensity inversions (e.g. black lines drawn as mid-gray). Such errors are mitigated by performing eye detection [23], extracting the lines from the masked area and using a lower threshold for the opening and closing filtering. These more detailed lines are directly overlaid on the previous NPR result.

5. Photometric stereo

Photometric stereo, originally proposed by Woodham [10], estimates surface normals on an object from multiple images taken from a fixed viewpoint but under different lighting conditions. As in shape from shading, photometric stereo uses the radiance values recorded at each image. However, the use of two or more images under different lighting rather than a single image makes normal estimation in photometric stereo a well defined problem, in principle providing improved shape recovery accuracy.

Under the assumptions of Lambertian reflectance, uniform albedo, and known distant lighting with unit intensity, the aim of photometric stereo is to find at each pixel a surface normal $\mathbf{n}$ satisfying a set of linear equations:

$$I_i = \mathbf{n} \cdot \mathbf{l}_i, \quad i = 1, \dots, k \quad (1)$$

where $\mathbf{l}_i$ is the lighting direction used when capturing image i , and I_i is the intensity of the pixel in image i . A unique solution of Equation (1) exists given $k = 3$ linearly independent lighting directions [24]. In our specific application, we can control the lighting via image relighting; organising the lighting directions to be orthogonal yields the following simple solution to Equation (1):

$$\mathbf{n} = \begin{bmatrix} n_1 \\ n_2 \\ n_3 \end{bmatrix} = [\mathbf{l}_1 \quad \mathbf{l}_2 \quad \mathbf{l}_3] \begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} \quad (2)$$

where each $\mathbf{l}_i$ is also written as a column vector, and a matrix product used to find the components of $\mathbf{n}$.

The three orthogonal lighting directions we use are found by starting off with a set of orthogonal unit vectors along the $\mathbf{x}$, $-\mathbf{y}$, and $\mathbf{z}$ axes, Then

Figure 6: Results of relighting and mapping. From left to right: NPR processed image, three relit images and corresponding bas-relief images.

rotating them by -45° about the $\mathbf{y}$ axis, followed by -60° about the $\mathbf{x}$ axis. This results in normalized lighting directions of $(\frac{\sqrt{2}}{2}, \frac{\sqrt{6}}{4}, \frac{\sqrt{2}}{4})$, $(-\frac{\sqrt{2}}{2}, \frac{\sqrt{6}}{4}, \frac{\sqrt{2}}{4})$, and $(0, -\frac{1}{2}, \frac{\sqrt{3}}{2})$ corresponding to three lights: above the head and to the left, above the head and to the right, and centrally below the chin. An example showing three relit images, and the corresponding bas-relief images from the same three directions, is given in Figure 6.

After finding surface normals using Equation (2), we apply the same postprocessing as in our previous work [7] to further enhance the salient features, and integrate the field of recovered surface normals to obtain the relief surface. A small linear scaling is then applied to achieve the exact height desired.

6. Experimental results

We now present and discuss various results obtained using these approaches. Evaluation must rely on visual inspection or a user study, as ground truth for quantitative evaluation does not exist. Desirable qualities of the result are that the face should be smooth yet recognisable, salient features should be distinct and well-preserved, and unwanted details should be suppressed. The geometry of the generated bas-relief should also be appropriate, i.e. convex and concave in the same areas as in the original 3D shape, so that the relief’s appearance is stable under changes of viewing and illumination directions. We use height maps of the generated bas-reliefs and renderings in different views to help reveal their true geometries.

We first present results where NPR is used in conjunction with SFS or photometric stereo for shape recovery, and compare them with those obtained using our previous method [7]. Secondly, we evaluate the robustness of our new method under changes of facial expression, head pose, and illumination direction in the input photograph. We finally compare our results with those generated by other approaches in the literature. The generated bas-relief

Figure 7: Comparative results. From left to right: original image, image processed using NPR, bas-relief surface from our previous method [7], new result using NPR and SFS, and new result using NPR and photometric stereo.

surfaces are rendered with gold or silver effect, and with lighting direction $(0, 0.5, 1)$ unless otherwise specified.

6.1. Effects of use of NPR and photometric stereo

We first consider the effects of adding NPR processing to our pipeline, using SFS. As we will see, it makes a clear improvement, so we also use NPR when we alternatively consider use of photometric stereo for relief surface recovery. In each case, the results are compared with those from our previous method, without NPR, using SFS. Figure 7 shows input images, input images after NPR processing, and the bas-reliefs generated. The third column shows that bas-reliefs generated by our previous method are noisy

Figure 8: A generated bas-relief rendered under different lighting, and from different viewing directions.

with excessive small detail and clutter. By abstracting the appearance, the addition of NPR clearly improves the results, no matter whether the reliefs are recovered using SFS or photometric stereo. The bas-reliefs generated by both shape recovery approaches (fourth and fifth columns in Figure 7) are smooth with well preserved facial features, and a rounded and appealing appearance. Moreover, they are easily recognizable.

To determine whether SFS or photometric stereo gives preferable results, we performed a user study. 26 participants unfamiliar with our research were shown the four input images in the left column of Figure 7 as reference images, and the results in the fourth (shape-from-shading) and fifth (photometric stereo) columns (to avoid bias, photometric stereo and SFS output were shown in a random order). The users were asked which results they preferred and why. In 84% of cases the results recovered using SFS were preferred, with users expressing a preference for the sharper features and slightly greater detail. In the remaining 16% of cases, the participants thought SFS over-emphasized the features and disliked the thicker outlines, preferring the photometric stereo results. Because of the majority preference for the results based on SFS, the remaining experiments, apart from timing the methods, used SFS for shape recovery.

6.2. Relief geometry

Figure 8 analyses a sample generated bas-relief surface (recovered using SFS). First, we show a depth map of the surface, which indicates its overall geometry. The regions of the face have correct shape on a large scale (i.e. do not suffer from convex-concave inversions), which ensures that the relief retains a good appearance under changes of viewing and lighting directions—see Figure 8(b) and (c). Many bas-reliefs, such as those on coins, are constrained to have constant albedo, and hence any albedo changes, as well as any texture, have to be represented as shape variations. As noted in our earlier paper, a beneficial side-effect of SFS is that it changes shading into geometry. Figure 8(d) shows that, as well as producing shading for eyebrows and lips, texture on the shirt is changed into geometry, and the spectacles are handled well, although not represented in the training data, demonstrating the robustness of our approach to a wide range of facial appearance.

6.3. Robustness

Our previous work [7] explored the robustness of our method to changes of input head pose. We next examine whether the addition of NPR maintains or even improves this robustness. We used nine photographs of the same person under frontal lighting in varying poses, as laid out in Figures 9(a) and (b). Photos are from the Extended Yale Face Database B [20]. Pose 0 is the frontal pose, poses 1–5 were about 12° from pose 0, and poses 6–8 about 24° . The training set and the bootstrap data were all frontal faces as previously described. Figure 9 shows that both our previous method and the new method generate bas-reliefs which still have overall correct shapes and maintain most of the prominent features, which we attribute to the stability of the relighting method when only slight head pose change occurs. The smoothing term in SFS and the final averaging of two surfaces also helps to reduce artifacts. In fact, the artifacts are more obvious using our previous method, especially when the pose change is significant. It seems that the NPR process, which emphasizes the salient features and smooths other areas helps to increase the stability of the relighting step, and thus further improves the robustness.

We also evaluated the robustness of our method under changes of facial expression, using photos showing faces with six different expressions: neutral, happy, surprise, sad, fear, and anger. The photos were provided by the CVPR group in the Department of Computer Science at the University of York. The results are shown in Figure 10. Even though the face has exaggerated

Figure 9: Results on faces with varying pose. Top: original photo for pose 0. Reliefs generated for varying poses are shown in accordance with input pose. (Photos from the Extended Yale Face Database B [20])

expressions, our new method generates plausible bas-reliefs with few artifacts and clearly recognisable expressions. The main noticeable artifacts occur at the tip of nose, and arise due to our assumption of Lambertian reflectance without considering the specular component in the input image. Mallick et al. give a PDE approach [25, 26] for removing specularity in images, which could potentially help to solve this problem.

We used another set of photographs, also from the Extended Yale Face Database B [20], to evaluate the robustness of our method to changes of illumination direction; results are shown in Figure 11. The generated bas-reliefs have generally similar appearance, although they are not exactly the same—relighting has nullified most of the differences in the input images. However, aside from the small artifacts mentioned earlier, other unwanted

Figure 10: Faces with different expressions. Left to right: neutral, happy, surprise, sad, fear, and anger. (Photos from CVPR group, Department of Computer Science, University of York)

Figure 11: Faces illuminated from different directions. Azimuth and elevation of the lighting are (left to right): $(0^\circ, 0^\circ)$, $(0^\circ, +20^\circ)$, $(-15^\circ, +20^\circ)$, $(+15^\circ, +20^\circ)$, $(-25^\circ, 0^\circ)$, and $(+25^\circ, 0^\circ)$ respectively. (Photos from the Extended Yale Face Database B [20])

artifacts appear wherever there are large shadows in the input (see the two right columns of Figure 11), as the quotient image technique cannot handle shadows. This limits the applicability of our method.

Finally, we applied our method to photographs of other face-like objects, as a further assessment of robustness. Figure 12 shows results for a dog, a cat and a car. These bas-reliefs are less plausible compared to results on human faces. The dog bas-relief has poor overall geometry because of the changing colour of its fur. The cat bas-relief is reasonable, although the fur is somewhat oversimplified. Our method performs less well on the car because of depth discontinuities in the front grill and the presence of sharp edges,

Figure 12: Results on other objects.

Figure 13: Comparison with previous work. From left to right: Alexa's work [5], Li's work [6], our previous work [7], and our new result with NPR.

which are not well handled by SFS and not well supported by the training data either.

6.4. Comparison to other approaches

In Figure 13, we compare our results with those generated using previously reported techniques for producing a bas-relief from a single image [5, 6] (improvements over our own previous work [7] were discussed earlier). Given a photograph, the goal of [5] is quite different—to create a relief approximating the input photograph under known directional illumination, while the aim of [6] is to restore brick and stone reliefs by treating the photograph as a rubbing image. Their results in Figure 13 meet their respective goals, but would have drawbacks for applications such as coin making. As experts from The Royal Mint note, when creating bas-reliefs for coinage, it is a desirable goal to simplify the photograph and emphasize key features. Simplification is important for both manufacturing reasons, and to ensure the coin's robustness to wear. The result produced using the techniques in [6] (second left in Figure 13) has many small details which would be difficult to manufacture,

as indeed does the result from our previous approach [7] (second right in Figure 13). The features in the result from [6] lack depth, so potentially would not be durable in certain applications. The result produced using techniques in [5] (Figure 13 left, with uncropped background which may affect comparison with the other two methods) approximates the input image under a fixed directional lighting. This would be unsuitable for applications such as coin making, as reliefs on coins are viewed under various lighting conditions, from various directions; such requirements can be compared to those for e.g. architectural reliefs which may always be viewed from below, but are nevertheless subject to changing lighting. Again, approximating the appearance without simplification may lead to manufacturing difficulties. Our result (Figure 13 right) was rendered using physically based rendering (using the pbrt toolkit) [27] with silver material and infinite skylight-day lighting (then converted to grey for comparison). Compared to the other approaches, the abstracted appearance with emphasized features has the greatest commonality with the kind of handcrafted bas-reliefs found on coins.

6.5. Timings

Our prototype implementation was based on MATLAB 7.9, using a computer with a 2.13GHz Intel Xeon CPU. Approximate times taken by each step of our previous method are given in [7]: for images of size 701×841 , the offline neural network training takes 3 hours, then given a new photograph, the steps used to produce a bas-relief take about 4.5 minutes. Addition of the NPR process takes about an extra 30 seconds, increasing the time to 5 minutes; about 4 minutes are due to the SFS step. In comparison, the photometric stereo approach has the advantage of speed: the photometric stereo step takes just 6s, so the whole relief processing process for a new image takes about 1 minute. These are due to the straightforward, non-iterative photometric stereo computation, unlike that needed for SFS. Ultimately, the user can trade off speed for perceptual quality of results.

7. Conclusions

Bas-reliefs of human faces are of particular interest in art and design. We previously proposed a method [7] to automatically generate bas-reliefs from frontal photographs of human faces. Here, we show that incorporating NPR allows abstraction of appearance in a similar way to that performed by sculptors, suppressing small details and clutter while emphasising salient

facial features. Experimental results show that incorporating NPR significantly improves the results. These are robust to limited changes of head position, facial expression, lighting direction, and viewing direction. We have also demonstrated that photometric stereo offers a faster alternative to shape-from-shading for bas-relief surface recovery. Expert opinion from The Royal Mint notes that our work “has good results with potential to assist artists to create bas-reliefs automatically”.

While the addition of NPR has significantly enhanced the quality of the results from our bas-relief generation framework, it could undoubtedly be taken further. As mentioned before, a pre-processing step, such as a PDE approach proposed by Mallick et al. [25, 26], could remove specularities in the input images which causes artifacts in the generated bas-reliefs. Taking into account facial albedo information during the shape recovery step, and using more sophisticated reflectance models, could make further improvements. The function learning process could be based on more than one training image, and training images could be based on real face models. An enlarged bootstrap set for image relighting might better span the space of facial albedos. Ultimately, practical applications demand extension of our method to faces seen in profile, and to a wider class of objects.

NPR has previously been used as an end in itself. This paper has also drawn attention to the potential usefulness of NPR as an *intermediate* step in a processing pipeline, an idea which may be of benefit to various other applications.

Acknowledgments

This work was supported by the One Wales Research Institute of Visual Computing (RIVIC) funded by a grant from the Welsh Government. We are grateful to Marc Alexa, Jinhui Yu, and Zhuwen Li for providing us their results for comparison, and to other researchers who have made various data publicly available. We also wish to thank Delcam and The Royal Mint for their support and useful discussions.

References

References

- [1] W. Song, A. Belyaev, H. Seidel, Automatic generation of bas-reliefs from 3d shapes, in: Proceedings of IEEE International Conference on Shape

Modeling and Applications, 2007, pp. 211–214.

- [2] T. Weyrich, J. Deng, C. Barnes, S. Rusinkiewicz, A. Finkelstein, Digital bas-relief from 3d scenes, *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2007* 26 (3), artical No. 32.
- [3] J. Kerber, A. Belyaev, H. Seidel, Feature preserving depth compression of range images, in: *Proceedings of the 23rd Spring Conference on Computer Graphics*, 2007, pp. 110–114.
- [4] X. Sun, P. L. Rosin, R. R. Martin, F. C. Langbein, Bas-relief generation using adaptive histogram equalization, *IEEE Transactions on Visualization and Computer Graphics* 15 (4) (2009) 642–653.
- [5] M. Alexa, W. Matusik, Reliefs as images, *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2010* 29 (4), artical No. 60.
- [6] Z. Li, S. Wang, J. Yu, K.-L. Ma, Restoration of brick and stone relief from single rubbing images, *IEEE Transactions on Visualization and Computer Graphics* 18 (2) (2011) 177–187.
- [7] J. Wu, R. Martin, P. Rosin, X.-F. Sun, F. Langbein, Y.-K. Lai, A. Marshall, Y.-H. Liu, Making bas-reliefs from photographs of human faces, *Computer-Aided Design* 45 (3) (2013) 671–682.
- [8] P. Rosin, Y. Lai, Towards artistic minimal rendering, in: *Non-Photorealistic Animation and Rendering*, 2010, pp. 119–127.
- [9] P. Rosin, Y. Lai, Artistic minimal rendering with lines and blocks, *Graphical Models* 75 (4) (2013) 208–229.
- [10] R. Woodham, Photometric method for determining surface orientation from multiple images, *Optical Engineerings* 19 (1) (1980) 139–144.
- [11] J. Kerber, M. Wang, J. Chang, J. Zhang, A. Belyaev, H.-P. Seidel, Computer assisted relief generation a survey, *Computer Graphics Forum* 31 (8) (2012) 2363–2377.
- [12] R. Huang, W. Smith, Structure-preserving regularisation constraints for shape-from-shading, in: *International Conference on Computer Analysis of Images and Patterns*, 2009, pp. 865–872.

- [13] P. Rosin, J. Collomosse, *Image and Video-Based Artistic Stylisation*, Springer, 2013.
- [14] W. Geng, *The Algorithms and Principles of Non-photorealistic Graphics: Artistic Rendering and Cartoon Animation*, Springer, 2010.
- [15] T. Luft, F. Kobs, W. Zinser, O. Deussen, Watercolor illustrations of cad data, in: *International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, 2008, pp. 57–63.
- [16] M. Tatzgern, D. Kalkofen, D. Schmalstieg, Multi-perspective compact explosion diagrams, *Computers & Graphics* 35 (1) (2011) 135–147.
- [17] A. Gooch, B. Gooch, P. Shirley, E. Cohen, A non-photorealistic lighting model for automatic technical illustration, in: *SIGGRAPH*, 1998, pp. 447–452.
- [18] T. Fine, *Feedforward neural network methodology*, Springer Verlag, 1999.
- [19] T. Riklin-Raviv, A. Shashua, The quotient image: Class based re-rendering and recognition with varying illuminations, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (2) (2001) 129–139.
- [20] A. Georghiades, P. Belhumeur, D. Kriegman, From few to many: Illumination cone models for face recognition under variable lighting and pose, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (6) (2001) 643–660.
- [21] H. Kang, S. Lee, C. Chui, Coherent line drawing, in: *Non-Photorealistic Animation and Rendering*, 2007, pp. 43–50.
- [22] A. Meijster, M. Wilkinson, A comparison of algorithms for connected set openings and closings, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (4) (2002) 484–494.
- [23] P. Viola, M. Jones, Rapid object detection using a boosted cascade of simple features, in: *Proc. IEEE Computer Vision and Pattern Recognition*, Vol. I, 2001, pp. 511–518.

- [24] T. Okatani, K. Deguchi, On uniqueness of solutions of the three-light-source photometric stereo: Conditions on illumination configuration and surface reflectance, *Computer Vision and Image Understanding* 81 (2001) 211–226.
- [25] S. Mallick, T. Zickler, P. Belhumeur, D. Kriegman, Dichromatic separation: specular removal and editing, in: *SIGGRAPH 2006 Sketches*, 2006, artical No. 166.
- [26] S. Mallick, T. Zickler, D. Kriegman, P. Belhumeur, Specularity removal in images and videos: A pde approach, in: *European Conf. Computer Vision*, Vol. 1, 2006, pp. 550–563.
- [27] M. Pharr, G. Humphreys, *Physically Based Rendering - From Theory to Implementation* (second edition), Elsevier, 2010.