Graphical Models 76 (2014) 202–213

Contents lists available at ScienceDirect

Graphical Models

journal homepage: www.elsevier.com/locate/gmod

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

ARTICLE INFO

Article history:
Received 19 July 2013
Received in revised form 2 January 2014
Accepted 12 February 2014
Available online 28 February 2014

Keywords:
Bas-relief
Non-photorealistic rendering
Shape-from-shading
Photometric stereo

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, we give experimental results showing how incorporating non-photorealistic rendering (NPR) into our previous framework enables us to both suppress unwanted detail, and yet also emphasize important features. We have consider an alternative approach to recovering relief shape, based on photometric stereo instead of SFS for surface orientation estimation. This can effectively reduce the computational time.

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 has considered automatic bas-relief generation from 3D scenes [1–4] and 2D images [5–7]. As images are much cheaper and easier to capture than 3D models, starting from the latter 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 a technique based on 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 stylized rather than faithful. They enhance salient shapes, simplify texture features, and suppress unwanted detail in other areas. Motivated by the way in which sculptors abstract appearance, in this paper, we show how similar effects can be achieved by use of a non-photorealistic rending (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.

☆ This paper has been recommended for acceptance by Siome Goldenstein and Peter Lindstrom.

* Corresponding author.
E-mail 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).

http://dx.doi.org/10.1016/j.gmod.2014.02.002
1524-0703/

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 relit the image to a standard direction (rather than 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, which as we later show, 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 [1–4] start with a *depth-map* of a 3D scene, and selectively compress depths to create a bas-relief surface. Fewer works [5–7] 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, when applied to portrait photographs, the results 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, noting 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 thus in general 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 with similar visual appearance to a 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 show that by incorporating non-photorealistic rendering into our previous method, this 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. watercolor, 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 watercolor 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 emphasizing 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, an unexplored idea.

3. Framework

The framework of our method is shown in Fig. 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 box in Fig. 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