Computers & Graphics 42 (2014) 1–13

Contents lists available at ScienceDirect

Computers & Graphics

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

Technical Section

Cornsweet surfaces for selective contrast enhancement

Henrik Lieng [a,*], Tania Pouli [b], Erik Reinhard [b], Jiří Kosinka [a], Neil A. Dodgson [a]

[a] University of Cambridge, UK
[b] Technicolor, Rennes, France

ARTICLE INFO

Article history:
Received 19 December 2013
Received in revised form
31 March 2014
Accepted 5 April 2014
Available online 16 April 2014

Keywords:
Contrast enhancement
Surface modelling
Computational photography

ABSTRACT

A typical goal when enhancing the contrast of images is to increase the *perceived* contrast without altering the original feel of the image. Such contrast enhancement can be achieved by modelling Cornsweet profiles into the image. We demonstrate that previous methods aiming to model Cornsweet profiles for contrast enhancement, often employing the unsharp mask operator, are not robust to image content. To achieve robustness, we propose a fundamentally different vector-centric approach with *Cornsweet surfaces*. Cornsweet surfaces are parametrised 3D surfaces (2D in space, 1D in luminance enhancement) that are extruded or depressed in the luminance dimension to create countershading that respects image structure. In contrast to previous methods, our method is robust against the topology of the edges to be enhanced and the relative luminance across those edges. In user trials, our solution was significantly preferred over the most related contrast enhancement method.

1. Introduction

The contrast of images is commonly enhanced to increase visual appeal or to prepare them for displays with limited capabilities. Extreme contrast enhancements often introduce artefacts and are likely to look objectionable. However, especially when the dynamic range of the medium is restricted, as in print, small and subtle changes are insufficient.

A solution, frequently used by artists to enhance contrast, is countershading adjacent to edges [1]. Countershading usually takes the form of a non-linearly increasing or decreasing luminance ramp, known as a *Cornsweet profile* [2]. The spatial extent of the countershading depends both on the amount of desired enhancement and on the presence and shape of nearby image features (Fig. 1). The artist is able to disguise edge enhancements in the structure of the scene such that they blend in and do not appear as artefacts.

Contrast enhancement is common in image processing. The best current algorithms take account of local intensity variation to improve the enhancement. They do not, however, consider image structure. Human artists *do* take account of structure and can produce countershading that enhances contrast while respecting image content. Inspired by this, we have investigated methods which aim to achieve the same effects in an automated process. The method we describe is a new vector-based algorithm where the edges to be enhanced are assumed to represent key structures in the image. We augment this by allowing additional edges that represent structure but are not themselves enhanced and also by considering local texture.

The purpose of this work is to provide artists with a selective contrast enhancement method that can achieve wider enhancements than those of conventional pixel-based techniques, while automatically taking into account the image structure. Our method can disguise the countershading profiles within image structure. We find, through psychophysical experimentation, that our method achieves selective enhancements that are preferred over the best of the previous methods.

We do *not* suggest a method to generally enhance the appearance of photographs, as with many previous approaches in image processing. Instead, the usage of our framework is targeted towards applications where objects, or regions across specific edges, are emphasised with Cornsweet-inspired contrast-enhancement profiles. We propose two ways of interaction with our system for users to achieve this aim. First, the image can be enhanced automatically, where input edges are found using high-level edge detection algorithms that aim to identify meaningful edges in the image. We envision applications such as smart object enhancement procedures that can be automatically applied as a quick post-processing step on the digital camera device. Second, edges can be manually defined by the artist, along with the given effect, propagated perpendicular to the edges. In addition to luminance contrast enhancements, we also explore adjustment in other channels, such as colour, saturation, and shading.

Our framework (Section 3) takes a single image as input, along with a set of image edges. It generates a set of B-spline surfaces that form a height field over the image, which is used to modify image intensity. The user can control the shape of the modification by manipulating individual control points. The height field is

* Corresponding author.
E-mail address: henrik.lieng@cl.cam.ac.uk (H. Lieng).

http://dx.doi.org/10.1016/j.cag.2014.04.001
0097-8493/

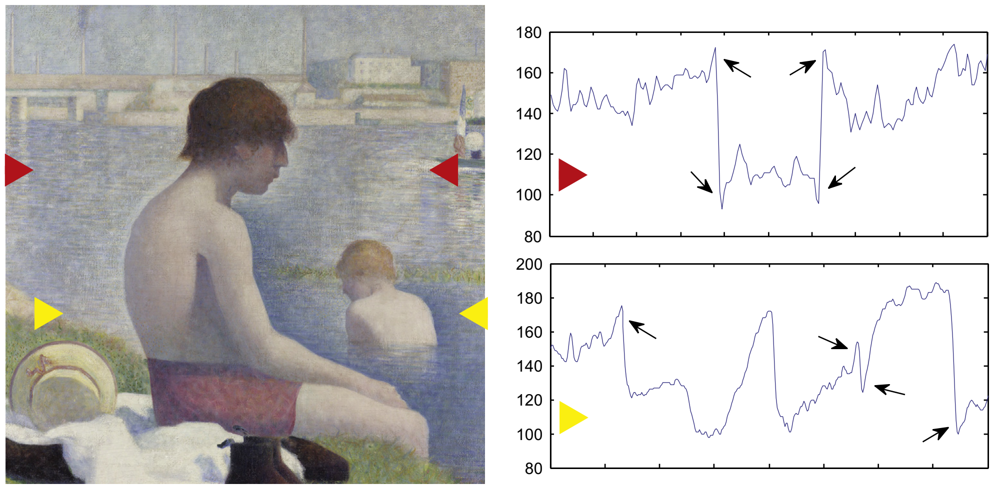

Fig. 1. An example of countershading in Seurat's *Bathers at Asnières*. Two horizontal profiles are shown. The Cornsweet profiles are indicated by black arrows. Each profile is 2000 pixels wide taken horizontally between two coloured triangles. The image was blurred by a 5-pixel radius Gaussian blur to smooth out the brushstrokes. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

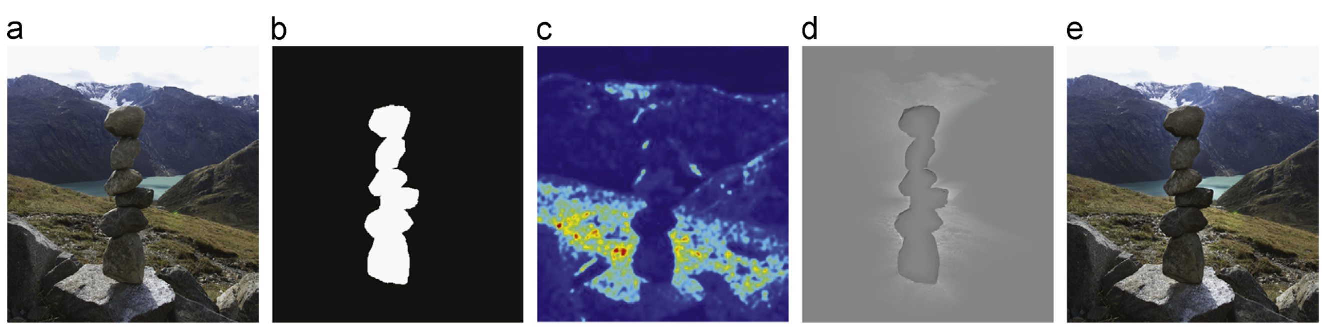

Fig. 2. (a) input image, (b)–(d) intermediate steps, (e) final result. Our selective contrast enhancement respects image structure by considering only object boundary edges (b), while information such as the amount of texture in the image (c) guides the strength of the adjustment. This produces an enhancement map (d) which is applied to the image.

constructed such that it forms a spatially varying Cornsweet profile for each edge (Fig. 2). We therefore term these surfaces as *Cornsweet surfaces*.

The novelty of the method and the reason to use spline surfaces are that these allow us to take other nearby edges into account: the profile is as wide as nearby edges allow. The importance of this approach lies in the observations that a wider Cornsweet profile will be less noticeable for the same amount of enhancement than a narrower profile and that Cornsweet profiles should not cross existing edges, as this may lead to disturbing saturation artefacts. An enhanced image is constructed by transforming the height at each pixel into a luminance multiplier, which is applied to the input image.

In summary, we make the following contributions:

- A novel image parametrisation approach based on B-spline surfaces.
- Selective countershading guided by image-content and artistic input allowing for strong contrast enhancements that better respect the edge structure.
- Psychophysical experimentation evaluating the visual quality of our results.

2. Related work

Unsharp masking is filtering an image by subtracting a low-pass version of itself. It induces different effects, depending on the width of the filter kernel. While small kernels increase sharpness and fine detail [3], they do not significantly alter perception of contrast [4]. Large kernels, however, *do* alter the appearance of contrast and can produce Cornsweet profiles (Fig. 3).

The Cornsweet effect is most suitable, as demonstrated by previous research [5–8], when the effect is only produced from a selected set of edges in the image; that is, when the low-pass version of the image for the unsharp mask is defined by an alternative image. This was first demonstrated by Luft et al. [5] using the input image's depth map (assuming such a map is defined). Thus, the contrast enhancement of the image is *not* applied generally across the image, but *selectively* according to depth layers. This method was further improved for temporal coherency in the setting of synthetic rendering [6]. If such depth maps are not defined, colour segmentations can be used. Trentacoste et al. demonstrated this by using the weighted least squares (WLS) smoothing filter [9] to produce a piece-wise linear colour segmentation of the input image [8]. In the setting of restoring lost contrast of HDR tone mapping, Krawczyk et al. used the unsharp mask on the scale space of the image [7]. The parameters of the unsharp mask were set separately for each level to best match the contrast of the HDR image.

The drawback of the unsharp mask for this particular problem is that it tends to introduce visible artefacts, which are strongly dependent on the interplay between the chosen parameter settings and the image content [8]. Fig. 3 demonstrates this point on various types of edges. For straight edges, where the convolution filter is only covering two image regions, the Cornsweet profile is modelled accurately by the unsharp mask. However, if the filter locally covers additional regions not associated with the edge we