Cornsweet Surfaces for Selective Contrast Enhancement

Henrik Lieng¹, Tania Pouli², Erik Reinhard², Jiří Kosinka¹, Neil A. Dodgson¹

¹University of Cambridge, UK

²Technicolor, Rennes, France

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.

Keywords: Contrast enhancement, surface modelling, computational photography

---

Figure 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.

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 (Figure 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 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

Figure 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.

device. Second, edges can be manually defined by the artist, along with the given effect, propagated perpendicularly 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 constructed such that it forms a spatially varying Cornsweet profile for each edge (Figure 2). We therefore term these surfaces *Cornsweet surfaces*.

The novelty of the method and the reason to use spline surfaces is that this allows 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 (Figure 3).

The Cornsweet effect is most suitable, as demonstrated by previous research [5, 6, 7, 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]. Figure 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 aim to enhance, the unsharp mask operation will not typically produce the Cornsweet profile. This can happen where nearby edges are too close to each other or at curved edges.

To reduce these saturation artefacts, the magnitude of the low-pass image, and therefore the Cornsweet profiles, can be adjusted to fit the given image. While such editing is sufficient for single images, it is not suitable for automatic applications and mass editing of image libraries. The unsharp mask is not suitable for such automatic applications due to its image-dependent behaviour. Figure 4 demonstrates this point: the

Figure 3: Previous work aiming to model the Cornsweet profile in image processing is defined by variants of the unsharp mask (middle column). While this operator models the Cornsweet profile well for straight edges adjacent to sufficiently large areas (top), saturation can occur for narrow regions (middle row) and curved edges (bottom). Our method (right) handles all of these cases. Additionally, our method guarantees to level out the Cornsweet effect to the original colour, whereas this is not guaranteed by the unsharp mask, as shown in the scan-line plots (blue: our method, magenta: unsharp mask, green: original image (leftmost image)).

magnitude of the Cornsweet effect is dependent on the difference in luminance of the related image regions. Such behaviour is expected from any convolution-based image operator. In contrast, our method is robust to such image content; thus, our approach can produce constant Cornsweet effects across multiple images.

Finally, our approach is edge-centric and not region-centric as related filter-based methods. That is, to selectively apply Cornsweet-style enhancements with previous methods, the artist needs to provide a segmented version of the input image, implicitly defining the enhancement edges. Our method, on the other hand, also support single edges as input. The latter approach is more sensible for this application as the Cornsweet effect is associated with edges, to enhance the contrast between their adjacent regions, and not the regions themselves.

The discussion above identifies inherent limitations of filter-based methods for the application of Cornsweet-style contrast enhancement. To counter these issues, we have investigated a fundamentally different approach with NURBS surfaces. NURBS are highly flexible, with the potential of being explicitly modelled across the image domain. We can therefore achieve more explicit control over the behaviour of the given effect than previous methods. Our framework is therefore well suited for modelling low frequency changes, such as the Cornsweet profile, on the image plane.

We note that other, more general, approaches also target contrast enhancement. This includes tone reproduction operators that prepare HDR images for display [10], scale-based image processing [11, 12], and edge-aware filtering [13, 14, 9]. In the setting of contrast enhancement, such approaches aim to provide a general enhancement of the image, where the entire

Figure 4: Images produced with our method (top) and the unsharp mask (bottom) with fixed parameters for all images. The scan-line plots show that our method (blue) has a constant behaviour and avoids clipping. The unsharp mask (magenta), on the other hand, increases enhancement magnitudes as the difference between the regions rises.

image signal is altered. Our aim can be viewed as the opposite: to alter the *perceived* contrast across selected image edges, while preserving the original feel of the image.

3. Description of algorithm

The input to our algorithm is a linear RGB image and a set of discrete edges to be enhanced. The edges can be provided by the user or computed automatically (Section 3.2). The maximum spatial extent σ of the countershading profile is determined by the image structure while the magnitude λ of the enhancement can either be computed [8], allowing for fully automated enhancement, or be provided by the user if more precise control is desired.

Our algorithm proceeds in three stages (Figure 5):

1. Based on the input edges, segment the image into a collection of non-overlapping regions bounded by B-spline curves (Section 3.2).
2. Determine an appropriate counter-shading profile within each region, using B-spline patches (Section 3.3) to encode the intensity adjustments (Section 3.4).
3. Apply the countershading to the image (Section 3.5).

3.1. Notation and B-spline structure

The B-spline patches [15] exist in a three-dimensional space, (x, y, z) , where (x, y) are the image coordinates and z is the intensity adjustment coordinate. Each B-spline patch has its own parameter space, for which we use (u, v) coordinates. There are two B-spline patches for each discrete edge in the image: one on each side. In any given patch, the B-spline curve given by $v = 0$ corresponds to the curve that is fitted onto that discrete edge. These are cubic B-splines with initially open-uniform knot vectors and with control points in image space denoted by $P_i = (x_i, y_i)$. The curves in the v direction are nominally normal

Figure 5: An overview of the stages of our enhancement algorithm.

to the discrete edge and thus define the shape of the counter-shading profiles.

Many variations of the Cornsweet profile have been proposed, including linear [16], parabolic [17], sinusoidal and exponential ramps [4]. The latter three are all monotonic functions with a decreasing gradient, culminating in C^1 continuity between the ramp and constant sections further away from the edge. Under suitably chosen conditions the human visual system finds it difficult to detect such ramps, arguably due to centre-surround processes in the retina [18, 19], and are therefore suitable to increase edge contrast without causing visual artefacts.

Cornsweet profiles can be well-modelled as rational quadratic Bézier curves (knot vector $[0, 0, 0, 1, 1, 1]$, three control points, weights $[1, w, 1]$; $w = 0.8$ in all results presented in this paper). This is implemented in our B-spline patch representation by setting the degree in the v direction to 2. The control points of a B-spline patch are thus $P_{i,j} = (x_{i,j}, y_{i,j}, z_{i,j})$, $j = 1, 2, 3$. The luminance adjustment values, $z_{i,j}$, depend on the enhancement magnitude λ ; they are multiplied by the Y channel of the image, computed in the xyY colour space [20]. Figure 6(d) illustrates the position and function of each control point.

3.2. Boundary detection and spline fitting

Input edges can be defined manually by the user in an image editing application of their choice. We have also investigated automatic edge detection. Specifically, we have experimented with boundary detection algorithms as they locate important image edges (e.g. object boundaries) but ignore less salient edges [21, 22, 23]. We chose Hoiem’s [23] boundary detection method as it proved to be sufficiently robust for our application. Its output is a set of edges connected through junctions. We add the image boundary to this set. If further artistic control is necessary, our system allows the user to refine the

boundary detection by adding and deleting edges through rough scribbles that are refined at a later stage.

We discovered that it is useful to allow the user to mark some edges in the image as *non-enhancing*. These are edges that are input to the image segmentation but which are not themselves enhanced. These edges act as boundaries to any counter-shading from nearby enhanced edges.

Given the input curves, we need to fit splines to them. These define the (x, y) locations of $P_{i,\{1,2\}}$. We use Medioni and Yasumoto method [24] to place our control points. Additionally, we model *sharp* corners with triple knots. A sharp corner is identified by considering the displacement between the B-spline curve and the input curve at the control point. If this displacement is above a threshold, it is regarded as sharp. Note that sharp corners are interpolated, dropping the continuity from C^2 to C^0 at this point.

From these fitted curves, the control meshes for the B-spline surfaces are set up. For each curve, one control mesh is constructed for each side of the curve. Control meshes related to neighbouring curves connected via junction points are then merged. The control point at the junction is set as sharp. Finally, given a curve control point (x_i, y_i) , the following coordinates of a related set of surface control points are defined as:

$$\begin{aligned} (x_{i,1}, y_{i,1}) &= (x_i, y_i), \\ P_{i,2} &= (x_i, y_i, 0), \\ z_{i,\{2,3\}} &= 0. \end{aligned}$$

The remaining coordinates $(x_{i,3}, y_{i,3})$ and $z_{i,1}$ determine the magnitude and extent of the countershading. Their values are defined as explained in the next two sections.

3.3. Finding the profile extent

The $(x_{i,3}, y_{i,3})$ location of $P_{i,3}$ determines the ‘far side’ of that point’s countershading profile (Figure 5(e)). When setting

Figure 6: (a) The discrete edges are converted to B-spline curves. (b) A region is associated with each side of each edge bounded by the curves (red) and the medial axis (green). Four different termination cases related to the first step of finding the profile extents are also shown. (c) A B-spline patch is defined within each region (only two examples are shown for clarity). (d) The Cornsweet profile is represented using a rational quadratic Bézier curve. (e) By changing the weight, w , on the second control point, $P_{i,2}$, the shape of the countershading profile can be controlled. Larger weights attract the curve towards $P_{i,2}$, leading to a steeper profile.

this location, we must ensure that the B-spline patches so defined produce a non-overlapping partition of the image (Figure 6(b)). This challenge is dealt with in three steps (Figure 7): (1) find the maximum possible extent for each P_i , (2) correct for intersections between neighbouring extents, and (3) perform simple thresholding for artefact avoidance. In addition to the control points, P_i , we use the medial axis A defined by the input curves. This contains a set of branches: $A = \{b_k\}$. For notational simplicity, $P_{i,j}$ is regarded as a 2D point in this section.

Step 1 finds the maximum possible extent c_i for each control point P_i . Let Q_i be the B-spline curve point corresponding to P_i and N_i its unit normal vector. Let r_i be a ray shot from Q_i in the direction of N_i . First, if the starting point Q_i is on the medial axis A , the related branch is deleted from the current set:

$$\text{if } Q_i \in b_k \Rightarrow A_i = A \setminus b_k$$

and the related control point is tagged as $\hat{P}_i$ and the ray as $\hat{r}_i$.

There are five cases which terminate the ray r_i (Figure 6(b)):

- (i) r_i hits the image border,
- (ii) r_i hits the current medial axis A_i ,
- (iii) r_i hits a non-enhancement edge,
- (iv) the magnitude of r_i is equal to the user-defined maximal extent: $\|r_i\| = \sigma$,
- (v) r_i hits the same enhancement B-spline edge Q_i lies on.

Figure 7: Our three-step solution maximises the extents while ensuring that there are no intersections. **Step 1** (a) Any arbitrary set of control points with their associated normal vectors (thin blue lines) can give a complex set of intersections. (b) The extents are bounded by various constraints such as the medial axis (green line). The thick blue lines show rays $\hat{r}_k$ that have started at a medial axis point. The cyan and magenta lines show the uv parameterisation. **Step 2** deals with remaining intersections. (c) First, a ray intersecting with any $\hat{r}_k$ will form triangular patches. (d) The final intersections are solved by (e) pairing rays with low intersection counts. **Step 3** (f) This may lead to extents with inappropriate profiles, (g) which is improved by thresholding.

The extent c_i is then set to either $\|r_i\|$ in the cases of (i)–(iv), or $\|r_i\|/2$ in the case of (v).

Step 2 handles intersections between rays r_i . First, if such an intersection lies on a $\hat{r}_k$ found in Step 1, the corresponding quadrilateral patch degenerates to a triangular patch:

$$P'_{i,3} = \begin{cases} \hat{P}_k + c_k N_k & \text{if } r_i \cap \hat{r}_k \neq \emptyset; \\ P_i + c_i N_i & \text{otherwise.} \end{cases} \quad (1)$$

The remaining intersections are dealt with by truncating the rays (and thus the profile extents c_i) to the furthest point possible, without having any intersection to any other ray. Given a pair of intersecting rays $r' \cap r'' = I$, the number of intersections with other rays is counted, considering only those between Q' (Q'') and I . The pair with the lowest intersection count is truncated to this intersection point. This is performed iteratively until there are no intersections left.

Step 3 After ensuring that there are no self-intersections, neighbouring extents c_i may still be very different, which can cause artefacts (Figure 7(f)). In this final step, we reduce variations in neighbouring extents (Figure 7(g)) such that they remain below a threshold.

3.4. Defining the enhancement magnitude

The coordinates $z_{i,1}$ define the magnitude of the enhancement at the curve. These are set by the input parameter λ and scaled by a texture measure corresponding to visual masking [25]. Finally, the values are clipped and thresholded between neighbours so that the enhancement neither oversaturates nor creates visual staircase artefacts.

As the profiles define the magnitude of the luminance adjustment in logarithmic space, $z_{i,1}$ depends on the input parameter λ and in simple images, a direct assignment would be sufficient. In real scenes though, texture (or lack thereof) in the

image may allow for further amplifications to the enhancement magnitude. Owing to visual masking, stronger enhancements remain tolerable in image areas with higher frequency content. To detect such areas, we rely on a texture map T of the image computed using the method by Bae et al. [26]. Areas of high texture will result in higher values of T , therefore leading to a stronger luminance adjustment, while flat surfaces require no texture-related adjustment. To account for texture-ness, we compute an intermediate $z'_{i,1}$:

$$z'_{i,1} = s \log_{10}(1 + \lambda) \bar{T}_N; \quad T \in [1, 1.2], \quad (2)$$

where $s = \pm 1$ depending on the orientation of the countershading profile for that region. If the average luminance of a region is higher than the region across the curve, $s = 1$ and vice versa. N is the pixel neighbourhood, extending to 40% of the extent of the current profile. Hence, $\bar{T}_N$ is the average texture in the neighbourhood of $(x_{i,1}, y_{i,1})$.

To avoid clipping where the enhancements would exceed the maximum or minimum allowed luminance of the output format, $z'_{i,1}$ is reduced if necessary by considering the luminance values in the neighbourhood N . We perform a check to keep pixel values within the dynamic range $[\mathcal{Y}_{min}, \mathcal{Y}_{max}]$ of the input image:

$$Z_i = \max_{(x,y) \in N} \log x(x, y);$$

$$z_{i,1} = \begin{cases} \log \mathcal{Y}_{max} - Z_i & \text{if } Z_i + z'_{i,1} > \log \mathcal{Y}_{max}; \\ z'_{i,1} & \text{otherwise.} \end{cases}$$

For negative values of $z'_{i,1}$ a similar calculation is performed to avoid results below $\log \mathcal{Y}_{min}$. If an input edge is not connected to other image edges or if two neighbouring edges require countershading profiles of opposite orientations, the $z_{i,1}$ at control points at the ends of such edges are set to zero.

Finally, we reduce variations in neighbouring magnitudes such that they remain below a threshold, set to 13% of the dynamic range. This threshold was found empirically.

3.5. Constructing the full luminance profile

Now that the control meshes for the B-spline surfaces are defined, the full luminance profile can be constructed. This is achieved by rendering the surfaces and projecting their depth values. This could be efficiently done by using the depth image produced by an off-screen rendering process with OpenGL or any other renderer. Since we did not have any requirements on performance, we used a MATLAB implementation with a naïve recursive evaluation of the surfaces.

The resulting, smooth discretised surface will not correspond exactly with the input edge. There will be a small number of pixels that are incorrectly categorised. That is: a pixel on the ‘wrong’ side of the edge may be adjusted, when it should have been left untouched; or a pixel on the ‘correct’ side of the edge may be left untouched when it should have been adjusted. Pixels on the wrong side of the curve simply have their adjustment set to zero. They are identified by considering whether the vector between the pixel and its closest curve pixel has the same

direction as the normal vector of that curve point. Untouched pixels on the correct side of the curve are inpainted by bilinear interpolation to the closest drawn pixels. This is sufficient because only a small number of pixels are affected and the bilinear interpolation gives a result very close to the accurate value.

The output luminance $Y_o(x, y)$ is computed for each pixel by modifying the original luminance Y using the discretised spline profile D and an anti-aliasing mask M :

$$Y_o(x, y) = (\exp(D(x, y)) \times M(x, y)) \times Y(x, y), \quad (3)$$

where $\times$ represents a component-wise multiplication. The image is then transformed from xyY to RGB to produce the final output.

Simple averaging is employed as the anti-aliasing method defining the mask M . More specifically, pixels close to the input edges are blurred with a Gaussian filter (filter size 3; $\sigma = 1$). An accurate anti-aliasing mask will not contribute much for this problem. The reason for this is that the alpha matte is effectively turned into an averaging filter, as difference on the edges in the enhancement image (approximately λ), are much larger than the original differences encoded in the alpha matte. When this alpha matte is scaled to fit the new Cornsweet edges, these original differences turn negligible. Finally, we note that aliasing is a general problem of applying Cornsweet profiles in photographs, as they are defined across hard edges. Thus, our solution smooths these hard edges, as edges in photographs are typically not hard transitions.

4. Results and comparisons

Our algorithm was implemented in MATLAB and is therefore not optimised for computational efficiency. Consequently, 1MP images currently require a few minutes of processing time. We envisage however that in a programming environment better suited to spline operations, the time would decrease significantly.

In Section 2 we established intrinsic limitations of the unsharp mask. In this section, we show that the related methods aiming to model the Cornsweet effect in photographs, all of which employ the unsharp mask, also face the same problems. The most related method, as it only requires a single input image, is the approach by Trentacoste et al. [8] (see Figure 8 for comparisons). The scan-line plots in Figure 8 show that the method of Trentacoste et al. saturate where neighbour edges are too close and when edges are curved. Moreover, their profiles do not always level out to the original luminance, whereas our approach does. Our psychophysical experiments (Section 5) indicate that our contributions *do* matter for a wide range of photographs.

4.1. User selectivity

As noted in Section 2, our method is better suited in terms of user selectivity compared with the unsharp mask (see Figure 10 for comparison). Since the Cornsweet effect is associated with edges, and not regions, such effects can be achieved directly as our method is edge-centric. Edge selectivity can be emulated

Figure 8: (c) Results of our method with a limited extent $\sigma_{\max}$ with the value of λ computed according to Trentacoste et al. [8]. Column (d) shows the corresponding results using the edge-aware enhancement of Trentacoste et al. (e) The scan-line plots (blue curves: our method, magenta: Trentacoste et al., green: input image) are plotted in luminance between the red arrows in (a).

with the unsharp mask by blurring out the hard edges in the segmented image, as demonstrated in Figure 10. However, this is not as convenient as directly selecting the edges and does not achieve full selectivity.

In addition to selecting Cornsweet edges, the artist can also add non-enhancing edges. These edges serve as constraints to the spatial extent of the enhancement profiles (Section 3) and provide a way to locally adjust the extent σ . In Figure 9, this application is demonstrated by ensuring that enhancement profiles do not cross shadow boundaries. Moreover, Figure 9 also demonstrates the difference between general image processing methods and our approach: general methods aim to *generally* alter the input image, while our method aims to apply the Cornsweet profile to a selective set of edges, without altering the original look of the image. Conversely, the type of edge selectivity available with our method is not available in these general approaches.

4.2. Reference-guided enhancement

Our algorithm is well suited for reference guided enhancements since each control point can be independently scaled. This is demonstrated by the textureiness measure which increases enhancement magnitudes in areas of high textureiness. Other uses of such guided enhancements include the use of depth maps [5, 6, 27] and contrast recovery from HDR tonemapping [7].

In Figure 11 the textureiness measure is complemented with an additional depth difference measure, where adjacent regions of large difference in depth values are emphasised. This effectively increases enhancements between foreground and background objects. Note that Luft et al. [5] also target depth-based contrast enhancement, inspired by neo-impressionism. However, as they employ the unsharp mask to solve this problem, they face the same problems as the standard unsharp mask operator.

Krawczyk et al. [7] employed Cornsweet-style countershading to bring back lost contrast from HDR tone mapping. As with the related methods aiming to model Cornsweet profiles they employ the unsharp mask. Thus, saturation artefacts can occur (Figure 12). Such artefacts are typically not as strong as those produced by Luft et al. and Trentacoste et al. since the size of the Gaussian kernel is adjusted for each level of the image scale space, according to the reference image. Additionally, clipping in luminance was avoided with a post-processing step; however, the correct Cornsweet profiles were not recovered, as illustrated in Figure 12. We note that we experimented with a similar scale-space approach for single images, but found defining robust kernel sizes for each level without any reference image challenging.

Finally, modifying the reference image can be of value for artistic applications. In Figure 13, bright strokes are added to the image defining the textureiness measure, thus increasing the

Figure 9: Comparison with general state-of-the-art image processing methods (e,f,g). While such methods provide a general enhancement of the image, altering the overall look, our method (h) *selectively* enhances a given set of edges (green strokes in (b)) with the Cornsweet effect while preserving the original feel of the image. Furthermore, the selective nature of our method cannot be matched with pixel-based methods. For example, non-enhancing edges (b-red strokes) ensure that enhancement profiles do not cross shadow boundaries, as shown in (c) by the traced normal vectors in (red: negative profiles, yellow: positive profiles, blue: corner points and junctions) and the magnitude of the enhancement of our method (d).

Figure 10: User selectivity is achieved with our method (top) by directly selecting the Cornsweet edges (top-left). In comparison, to achieve selectivity with the unsharp mask (bottom), image segments must be defined (bottom-left). Edge selectivity can be emulated by blurring hard edges in the segmented image. However, the rightmost images, showing the magnitude of the enhancements, demonstrate that full edge selectivity is difficult to achieve with the unsharp mask.

magnitude of the enhancement profiles at those locations. This is therefore a way to locally edit the magnitude, λ , of the enhancements.

4.3. Other applications

Our Cornsweet surfaces can also be applied in settings other than luminance enhancement. In Figure 14 and the supplementary webpage we demonstrate other artistic applications with modification in colour, saturation, chroma, grey-scale, and shading channels. Additionally, artistic filters typically add additional frequencies to the image. Thus, high magnitudes of en-

hancements can be applied without being objectionable due to visual masking (Figure 14(d)):

While such edge-centric enhancement is challenging with pixel-based methods, similar approaches in vector-centric image editing have been explored. Most similar is the approach by Orzan et al. [28]. Their method diffuses colours out from selected curves in the image. Our method differs from the approach by Orzan et al. in two ways. First, their approach is aimed towards fully vectorising the image, whereas our method supplements the input raster image with spline surfaces. Our method is therefore more suited towards applications where a selected effect (see Figure 14 for examples) is modelled into the image from a smaller set of edges. Second, our approach enables support for editing gradients out from curves, via the weight w . Thus, our method provides the artist with an additional degree of freedom not supported by the method of Orzan et al. which only supports colours at curves.

4.4. Discussion

The Cornsweet surfaces we produce, via the NURBS framework, can be modelled in both spatial and luminance domains in other ways than presented in this article. For example, we decided to model the sign of a given surface by comparing the difference in luminance of the two given image regions. Thus, we have aimed to model the Cornsweet illusion, which states that the luminance of the brighter region should be increased, and vice versa. However, in some situations other aspects can be incorporated to such modelling decisions. In Figure 11, for example, the leftmost cup is modelled with a negative outgoing profile, while the other three cups are associated with positive outgoing profiles.

Figure 11: Reference-guided enhancement with depth. Object boundaries of large differences in depth are emphasised. Note that the enhancements are amplified to demonstrate the difference. The scan-line plots (blue curves: our method, magenta: Luft et al., green: input image) are plotted in luminance between the red arrows shown in the original image.

In this example, some would argue that spatial or depth coherence, in addition to following the Cornsweet profile for each edge separately, should dictate the sign of the surfaces. Given that decisions can be defined algorithmically, such algorithmic changes are trivially achieved with our framework, since each surface is explicitly modelled. In the setting of depth coherency (Figure 11), the sign of the enhancements can be dictated by depth values, and not luminance, similar to the approach by Luft et al.

Cornsweet surfaces, modelled with NURBS, are defined with C^0 creases along points associated with triple knots; for example, sharp corners are modelled with such triple knots. While this can be a visible artefact in flat shaded images, such as Figure 5, we have not noticed any related artefacts in photographs, as visual masking masks these creases. We note that alternative surface representations, such as subdivision surfaces, can provide smooth surface representations for meshes of both quads and triangles. Subdivision surfaces can therefore be employed to avoid such C^0 creases.

4.5. Limitation

Our method is not suitable for edges and regions where the Cornsweet effect is negligible. Such edges are found adjacent to narrow regions. Dooley and Greenfield demonstrated that perceived contrast is *not* altered with profiles of less than 0.2 visual degrees [17]. Thus, edges surrounding objects spanning only a

Figure 12: Modelling Cornsweet profiles with an HDR image. Our enhancement is computed in the original HDR image, which can then be tonemapped. The scan-line plots (blue curves: our method, magenta: Krawczyk et al., green: tonemapped image without Cornsweet enhancement) show that our method more accurately models the Cornsweet profile.

Figure 13: A bright stroke is added to the image defining the texture measure (left), resulting in a bright halo behind Buddha's head (right).

few pixels in the image, such as leaves and distant trees, cannot be enhanced with the Cornsweet effect. Approaches aiming to model the Cornsweet effect, such as our method, are most effective where sufficiently large regions are well defined adjacent to relatively hard edges, as the Cornsweet profile can be properly defined in such regions.

5. Experimental evaluation

In Section 2 and 4 we demonstrated that the unsharp mask and related methods create saturation and other artefacts if edges are not sufficiently straight or adjacent edges are too close. Our solution explicitly deals with these problems and does not create such artefacts. In this section we ask whether this contribution matters for general photographs. To answer this question, we conducted a psychophysical experiment comparing the relative merit of our work and the method by Trentacoste et al. [8]. We only compared with Trentacoste et al. since their method is the only approach aiming to model the Cornsweet effect from single images.

5.1. Design

We designed a 2-alternative forced choice (2AFC) experiment comparing the two methods at similar levels of enhance-

Figure 14: Enhancements of a flower (a) in saturation (b) and grey-scale (c) channels. (d) Extreme enhancements can be of value in the setting of artistic filtering. For comparisons with images without enhancement for (c,d), see the supplementary webpage. All of these images were produced with the weight w set to 0.4.

ment. For each trial, participants were asked to select the *best enhanced image in their opinion* between corresponding results from both methods. The unprocessed image was also shown as a reference to enable participants to determine the extent and quality of each enhancement.

Ten different images were processed with each method, using 4 values of enhancement strength $\lambda = (0.5, 1.0, 1.5, 2.0)$ for our method, resulting in 40 trials per participant. The images used were selected to depict a wide variety of objects and scenes, including both indoor and outdoor settings. We opted against including portraits in this comparison as different criteria may be used when assessing imagery with people, which could bias the results. We believe however that the selection of images (see the supplementary material) sufficiently demonstrates the flexibility of our enhancement technique.

To select comparable enhancement strengths between the two methods, a series of images with varying λ values were produced for Trentacoste et al.’s method and their global contrast computed. Matching pairs for each λ were selected such that they minimised the difference between global contrast within the two methods for each image. The enhancement extent was fixed to 4.2 degrees of visual angle (assuming a 0.5m viewing distance), which corresponded to $\sigma = 250px$ in our method.

The experiment was designed in MATLAB using the Psychophysics Toolbox and performed on a laptop display with the resolution of 1366×768 . 17 participants (7 female, 10 male, age mean 43.6, st. dev. 20.0) took part in the experiment, all reporting normal or corrected to normal vision, and normal colour vision. Participants were positioned so they could comfortably view and use the laptop (approximately 0.5m viewing distance), but viewing position and distance were not precisely controlled as we were interested in measuring viewer preferences rather than perceptual phenomena. The order of stimuli as well as the position (left, right) of each method were randomised.

Figure 15: Detailed results for our experiment. The stacked bar charts show how often each method was selected, grouped by λ (left) and by image (right).

5.2. Analysis

In a total of 680 trials, our method was selected 464 times, leading to a normalised mean of 0.68 (st. dev. 0.11). The results were analysed with a one-way ANOVA, finding that our method led to a significantly preferred enhancement overall ($F(1, 33) = 95.99, p < 0.001$).

To further explore the results, we performed a multi-way ANOVA on the enhancement strength λ , the image and the participant, as well as their interactions. We found that the selection of λ led to a significant effect in terms of enhancement preference ($F(3, 679) = 4.88, p < 0.005$), although post-hoc analysis indicated that this was only the case between $\lambda = 0.5$ and $\lambda = 2.0$. Interestingly, although for all λ values our method was selected more, increasing the strength of the enhancement seemed to progressively reduce the preference towards our method. The detailed results are shown in Figure 15(a).

Similarly, significant differences in participant choices were found between the different images ($F(9, 679) = 37.48, p < 0.001$). Specifically, we found that our method performed significantly better in images where contrast was already high. In such cases, our adaptive adjustment allowed for the contrast to be enhanced where possible, while preserving local details. Compared against high contrast regions in the same images enhanced with the method by Trentacoste et al. which were likely to become over or under exposed, effectively losing local detail (an example is shown in Figure 8). Detailed results can be seen in Figure 15(b).

Although a significant effect was found for participant number, post-hoc tests indicated that most participants responded similarly with only few exceptions. Detailed results for each participant can be seen in Figure 16(a) and (b). No interactions were observed between participants and enhancement strength λ , suggesting that the preference for softer enhancements is universal. In contrast, the selection of image led to a significant interaction both with participants and λ values ($F(27, 679) = 2.94, p < 0.001$ and $F(144, 679) = 1.97, p < 0.001$ respectively), suggesting that the preferred amount of enhancement varies per image.

Overall, our analysis shows that our adaptive enhancement method leads to contrast enhancements that are preferred by viewers. We have found that participants generally prefer softer enhancements, which aligns with the visual effect obtained with

Figure 16: Per-participant results. The left bar charts show how often the two methods were selected by each participant, while the right plot shows groupings for the participants as determined by post-hoc analysis. Participants highlighted in blue were significantly different from each other, but not from the main group. As such, they were not considered as outliers. Note that the dashed line in red indicates the overall mean.

our method: only salient edges are enhanced with the Cornsweet profile being gently embedded into the image.

6. Conclusions and future work

Although contrast enhancement is a common task in editing photographs or digital imagery, little control is available to the user regarding the placement and form of the enhancements. We have presented an algorithm for selective enhancement of edges, which uniquely respects image structure.

Rather than enhancing all edges in the image, we enhance a selective set of edges, chosen either via user input or high-level edge-detection algorithms. Our vector-centric approach creates countershading profiles that not only follow the selected edges but also respect structure in the image, such as T-junctions and occlusions. This allows us to achieve more natural results that emulate the countershading used by artists, without sacrificing contrast compared to traditional unsharp masking-based solutions.

Our adaptive solution ensures that image regions are maximally enhanced without becoming saturated, therefore preserving local detail that could otherwise be lost. Through a psychophysical experiment, we find that the enhancement achieved by our method is preferred in a range of scenes and parameter levels, confirming the utility of our technique.

Our experiment suggests that the preferred adjustment varies per image. For example, it might be that countershadowings between two regions of similar colours are more preferred over two regions of very different colours (or vice versa). An interesting avenue for future research is whether such countershadowing can be learned.

Our framework is currently limited to treating single images. Extension to video footage could be an interesting avenue for future research.

References

- [1] Ratliff F. Paul Signac and color in neo-impressionism. Rockefeller Univ. Press; 1993.
- [2] Cornsweet TN. Visual Perception. Academic Press; 1970.

- [3] Neyenssac F. Contrast enhancement using the Laplacian-of-a-Gaussian filter. *CVGIP: Graph Models Image Process* 1993;55(6):447–63.
- [4] Kingdom F, Moulden B. Border effects on brightness: A review. *Spatial Vision* 1988;3(4):225–62.
- [5] Luft T, Colditz C, Deussen O. Image enhancement by unsharp masking the depth buffer. *ACM Trans Graphics* 2006;25(3):1206–13. URL: <http://doi.acm.org/10.1145/1141911.1142016>. doi:10.1145/1141911.1142016.
- [6] Ritschel T, Smith K, Ihrke M, Grosch T, Myszkowski K, Seidel HP. 3D unsharp masking for scene coherent enhancement. *ACM Trans Graphics* 2008;27(3):art. 90.
- [7] Krawczyk G, Myszkowski K, Seidel HP. Contrast restoration by adaptive countershading. *Comp Graph Forum* 2007;26(3):581–90.
- [8] Trentacoste M, Mantiuk R, Heidrich W, Dufrot F. Unsharp masking, countershading and halos: Enhancements or artifacts? *Comp Graph Forum* 2012;31(2):555–64.
- [9] Farbman Z, Fattal R, Lischinski D, Szeliski R. Edge-preserving decompositions for multi-scale tone and detail manipulation. In: *ACM Trans. Graphics*; vol. 27. ACM; 2008, p. 67.
- [10] Reinhard E, Heidrich W, Debevec P, Pattanaik S, Ward G, Myszkowski K. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. 2nd ed.; Morgan Kaufmann; 2010.
- [11] Fattal R. Edge-avoiding wavelets and their applications. In: *ACM Trans. Graphics*; vol. 28. ACM; 2009, p. 22.
- [12] Paris S, Hasinoff SW, Kautz J. Local laplacian filters: edge-aware image processing with a laplacian pyramid. *ACM Trans Graphics* 2011;30(4):68:1–68:12.
- [13] Tomasi C, Manduchi R. Bilateral filtering for gray and color images. In: *ICCV. IEEE*; 1998, p. 839–46.
- [14] Durand F, Dorsey J. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans Graphics* 2002;21(3):257–66.
- [15] Prautzsch H, Boehm W, Paluszny M. *Bézier and B-spline Techniques*. Berlin: Springer; 2002.
- [16] Wachtler T, Wehrhahn C. The Craik-O’Brien-Cornsweet illusion in colour: Quantitative characterisation and comparison with luminance. *Perception* 1997;26:1423–30.
- [17] Dooley RP, Greenfield MI. Measurement of edge-induced visual contrast and a spatial-frequency interaction of the Cornsweet illusion. *J Opt Soc America* 1977;67(6):761–5.
- [18] Kuffler SW. Discharge patterns and functional organization of mammalian retina. *J Neurophysiology* 1953;16:37–68.
- [19] Marr D. *Vision*. San Francisco: WH Freeman; 1982.
- [20] Billmeyer F, Saltzman M. *Principles of color technology*. 3rd ed. ed.; New York: Wiley; 2000. ISBN 047119459X.
- [21] Ren X, Fowlkes CC, Malik J. Figure/Ground assignment in natural images. In: *ECCV*. 2006, p. 614–27.
- [22] Leichter I, Lindenbaum M. Boundary ownership by lifting to 2.1D. In: *ICCV*. 2009, p. 9–16.
- [23] Hoiem D, Efros AA, Hebert M. Recovering occlusion boundaries from an image. *Int J Comput Vision* 2011;91(3):328–46. doi:10.1007/s11263-010-0400-4.
- [24] Medioni G, Yasumoto Y. Corner detection and curve representation using cubic B-splines. *Comp Vis, Graphics and Image Proc* 1987;39(3):267 – 78. doi:10.1016/S0734-189X(87)80181-0.
- [25] Ferwerda JA, Pattanaik SN, Shirley P, Greenberg DP. A model of visual masking for computer graphics. In: *SIGGRAPH*. 1997, p. 143–52.
- [26] Bae S, Paris S, Durand F. Two-scale tone management for photographic look. In: *SIGGRAPH*. ISBN 1-59593-364-6; 2006, p. 637–45. doi:10.1145/1179352.1141935.
- [27] Ihrke M, Ritschel T, Smith K, Grosch T, Myszkowski K, Seidel HP. A perceptual evaluation of 3D unsharp masking. In: *Human Vision and Electronic Imaging XIV. Proc. SPIE* 7240; 2009, p. 72400R.
- [28] Orzan A, Bousseau A, Winnemöller H, Barla P, Thollot J, Salesin D. Diffusion curves: a vector representation for smooth-shaded images. *ACM Trans Graphics* 2008;27(3):92. doi:10.1145/1360612.1360691.