Pacific Graphics 2015
N. J. Mitra, J. Stam, and K. Xu
(Guest Editors)

Volume 34 (2015), Number 7

Contrast-Enhanced Black and White Images

Hua Li [†1] and David Mould[‡2]

[1]University of North Carolina Wilmington [2]Carleton University

Figure 1: Sample results from our method.

Abstract

This paper investigates contrast enhancement as an approach to tone reduction, aiming to convert a photograph to black and white. Using a filter-based approach to strengthen contrast, we avoid making a hard decision about how to assign tones to segmented regions. Our method is inspired by sticks filtering, used to enhance medical images but not previously used in non-photorealistic rendering. We amplify contrast of pixels along the direction of greatest local difference from the mean, strengthening even weak features if they are most prominent. A final thresholding step converts the contrast-enhanced image to black and white. Local smoothing and contrast enhancement balances abstraction and structure preservation; the main advantage of our method is its faithful depiction of image detail. Our method can create a set of effects: line drawing, hatching, and black and white, all having superior details to previous black and white methods.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

The clean, strong appearance of two-tone rendering is favored by many audiences. However, preserving the key image content with a binary palette is challenging. It is difficult to obtain a clear segmentation without user intervention. Color reduction can make the objects indistinct, as different objects may have similar colors.

The most extreme case of palette reduction is to convert a color image to black and white. Over a half century, researchers have proposed many thresholding techniques [SHB07] to generate black and white pictures. Unfortunately, thresholding alone cannot produce a good outcome. A viable black and white method requires combining thresholding with enhancement and smoothing steps in order to both preserve detail and abstract the image.

Non-photorealistic rendering (NPR) researchers have proposed many algorithms for stylizing images into black and white [MG08,XK08,RL10,Win11,WKO12] and creating related effects using a binary palette, e.g., paper cut [XKM07, MZZ10]. Region-based methods [XKM07, MG08, XK08, RL10, MZZ10] and filter-based methods [Win11, WKO12] are two major techniques used in past research. A region-based black and white method relies on the output from segmentation algorithms, which are not completely reliable, and thus the content in an image cannot be satisfactorily preserved in the stylization. Filter-based methods such as those based on eXtended difference-of-Gaussians compendium (XDoG) [RL10,Win11,WKO12,RL13] succeed in capturing a lot of detail and the smooth Edge Tangent Field (ETF) yields a clean, gentle appearance. However, the pro-

† lihua@uncw.edu
‡ mould@scs.carleton.ca

cess deliberately lengthens edges and smooths the ETF, potentially producing artefacts. For example, sharp corners are generally missing, and unwanted edge extensions may depict the content wrongly. The recent NPR book [LR13] describes the problem of rendering using a binary palette as "an underdeveloped topic".

Sticks filtering [CJO98] is an image enhancement technique, successfully used in medical imaging problems to detect and enhance linear structures in low-contrast and noisy images such as those from ultrasonography. Inspired by sticks filtering, we introduce a new filter-based black and white method that attacks the tone reduction problem from the perspective of emphasizing contrast. Region-based methods struggle with thin features that are challenging to segment, but we do not depend on a segmentation. Our method has comparable quality to XDoG-based methods and is better for linear features. Overall, our results have high structural preservation and few artefacts.

Our method operates by first enhancing contrast at each pixel along the orientation of maximal response, eventually thresholding the contrast-enhanced image. Where the image lacks much variation, we locally smooth it. This approach experiences a general tradeoff between detail preservation and abstraction, seen elsewhere in NPR; we have some abstraction owing to the smoothing element, but the method concentrates on detail preservation. Varying the threshold creates images with differing levels of overall blackness; edges are faithfully preserved, and our line drawing variant is extremely robust. We characterize our contributions as follows:

- We introduce an automatic contrast-focused method to produce black and white images from color images.
- We adjust image contrast so as to enable content preservation after thresholding. Low-contrast edges are particularly enhanced. We also employ both global and local smoothness to control the level of abstraction.
- We create line drawing, hatching, and black and white effects. Our line drawing is particularly robust.

The paper discusses previous black and white methods in Section 2 and then describes our algorithm in Section 3, including sticks filtering and proposed variations. Section 4 discusses our results and supplies comparisons to previous methods. Section 5 concludes our paper.

2. Related Work

Image stylization [KK11, RC13, HS13] sometimes targets tone reduction, seeking to convert an image to a cartoon-like or illustrated version. We are interested in extreme tone reduction, converting an image to a binary palette. Halftoning and stippling yield a kind of black and white effect, with black pixels or dots on a white canvas. An entire subfield of NPR treats such effects. While halftoning and stippling use high spatial frequencies, we seek to create black and white effects with large black clusters, in the form of lines and large irregular shapes. Xu et al. presented a method for papercut [XKM07] in 3D space. Buchholz et al. [BBDA10] used a similar graph-cut minimization approach to assign black and white colors to regions for 3D binary stylization. We focus on image-based black and white approaches. There are two major image-based techniques for black and white rendering: region-based methods [XKM07, MG08, XK08, RL10, MZZ10] and filter-based methods [RL10, Win11, WKO12, RL13].

A region-based method usually relies on optimization, organizing segmented regions into a graph and minimizing an objective function over a labeling of the graph nodes. The objective function contains terms relating to detail preservation and smoothness, and parameters can control the level of abstraction. A labeling assigns black and white to each node in the graph, and as nodes are image segments or pixels, a black and white image emerges from the optimization. Xu and Kaplan used this idea to create artistic thresholding [XK08]. Their objective function was designed to preserve contrast between region boundaries. However, weak edges are not preserved, and the authors comment that their results lack the "exclusive-or" effect, where foreground objects are apparent on top of a background that varies between black and white. Mould and Grant used energy minimization to create black and white [MG08] with a graph-cut based image segmentation; their approach used a base-detail decomposition to balance abstraction and detail preservation, with good treatment of textures and small features. Meng et al. [MZZ10] used facial recognition to trace the key points on a portrait to solve an image binarization problem with sparse cuts. Generally, region-based methods struggle because the segmentation problem is so challenging.

Filter-based methods are usually more effective than region-based methods. Gooch et al. [GRG04] created black and white portraits by adjusting DoGs over different scales. Kang et al. [KLC07] filtered a locally smoothed edge tangent field to create a black and white line drawing. Rosin and Lai introduced minimal Rendering [RL10, RL13], which starts with filter-based line drawing [KLC07] and carefully places a combination of refined lines and blocks with a table-based lookup for the choice of black and white. The XDoG compendium [Win11, WKO12] has great success in generating a family of effects due to its delicate presentation of weak edges and its flexibility in operating over continuous tones. A key to the XDoG method is its use of local smoothness along the edge tangent field to extend short edges as well as to suppress noise. The smoothing is composed of a small-scale average across an edge and a longer-scale average along the edge. We argue that the XDoG-based black and white approaches creates artefacts. For example, sharp corners are hardly seen and unwanted extensions could misrepresent objects.

3. Our Algorithm

We propose a filter-based method inspired by the sticks filter. Our goal is to preserve detail while abstracting the image into black and white. We plan to avoid the region labeling problem. We argue that the central problem in black and white rendering is to express the boundaries of objects. Object boundaries are visible in the original, continuous-tone images because of the contrast across tones. Our basic idea is to enhance this contrast, revealing the objects. If we strengthen an edge by repeatedly adding dark to the darker side and light to the lighter side, eventually, one of the two sides will reach an extreme end of the greyscale spectrum: pure black or white. However, simply enlarging intensity differences amplifies noise. Thus, it is necessary to suppress noise with some sort of smoothing or regularization process. Our method includes several stages: an initial blurring, a conversion from a color image to an enhanced and smoothed greyscale image, a binary thresholding, and final cleanup involving further smoothing and removal of small clusters.

3.1. Sticks Filtering

Sticks filtering is an edge-based image enhancement proposed by Czerwinski et al. [CJO98]. Eramian et al. [EAP07] experimented with variants of sticks filtering on ultrasonographic images, intending to enhance texture; other extensions appear in the survey by Noble and Boukerroui [NB06]. The key benefits of sticks filtering are the ability to obtain nearly optimal lines and boundaries, and to reduce uncorrelated speckle noise.

Sticks filtering tries to determine whether a line of some orientation passes through a pixel. For example, if we consider m discrete orientations, there are $m+1$ hypotheses, where m hypotheses represent possible orientations of linear features and the final hypothesis is the null hypothesis of no line at all. The orientation with the maximum likelihood is chosen. Normally, the determination considers a neighborhood around the pixel. We rasterize the straight line segments, or "sticks", into a square region. Figure 2 shows eight rasterization directions. For a region size of n, at most $2n-2$ orientations can be distinguished. We suggest using $n=5$ or $n=7$.

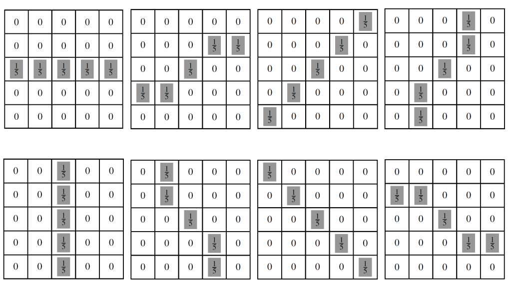

Figure 2: Rasterization of a set of discrete directions, representing 8 line orientations (length $n=5$).

Sticks filtering uses the maximum result from the m oriented filters as the output for the current pixel. Our method takes two ideas from sticks filtering: first, measuring an image quantity along a set of discrete directions, and second, using the maximum response to decide what to do at the central pixel. In our case, we plan to enhance contrast. We compute the difference among all the directions around the region. The chosen orientation, with maximum difference in intensity from the region, aligns with the locally strongest feature. Unlike previous black and white methods, which often explicitly search for edges or boundaries, we let the contrast from the original image guide the algorithm: contrast is increased along the direction of maximum response, amplifying local tonal differences.

3.2. Contrast Adjustment

We enhance contrast iteratively, repeatedly applying Equation 1. Some amount of grey is added to or subtracted from $I(x,y)$ to enhance the initial tendency of the pixel. The sign of $\triangle I$ depends on the difference between the brightness of the current pixel $I(x,y)$ and that of the local region. If the corresponding pixel is brighter than the average of its neighbours, $\mathrm{sign}(I(x,y)-\bar{I}(x,y))=1$; otherwise, $\mathrm{sign}(I(x,y)-\bar{I}(x,y))=-1$.

$$O(x,y)=I(x,y)+\triangle I \tag{1}$$

$$\triangle I=\mathrm{sign}(I(x,y)-\bar{I}(x,y))\times\delta I_{\max}\times p_d \tag{2}$$

$$\delta I_{\max}=\max_{s_1,s_2,\cdots,s_m}(|\mu(s_i)-\bar{I}(x,y)|), \tag{3}$$

where $\mu(s_i)=\frac{\sum_{\in s_i} I(j,k)}{n}$, the average intensity along stick s_i.

$$\bar{I}(x,y)=\frac{\sum_{\in n\times n\ \mathrm{neighborhood}} I(r,s)}{n\times n} \tag{4}$$

The quantity $\delta I_{\max}$ is computed from the maximal contrast response. Figure 3 shows the chosen sticks at five different locations for the scenario of a dark object placed on a light background. For example, at location 1, the masks cover the edge of the grey object. The horizontal line (in red) shows the direction of maximum difference between the average intensity (μ) of a stick and the average intensity ($\bar{I}(x,y)$) of the local area. Here, the chosen horizontal direction follows the edge of the rectangular object. The pixel at location 1 is lighter than its neighbours, so the sign is equal to 1: this pixel tends towards white, hence will be adjusted to become even lighter. The intent is to strengthen the initial contrast, increasing the difference between the pixel and the local area. Location 5 is on a uniform area and the choice

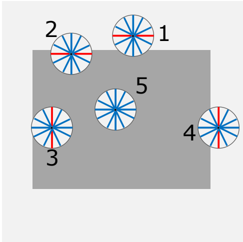

Figure 3: A darker object on a light background. Labels $1,\cdots,5$ mark the locations of filtering. The blue bars are the sticks and the red bars are the chosen orientation of the line.

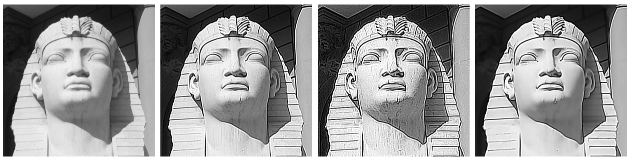

(a) Original (b) (c) (d)

Figure 4: Image enhanced after filtering. (b) $p_{d_1}=1, p_{d_2}=1, p_{d_3}=1$. (c) $p_{d_1}=2, p_{d_2}=2, p_{d_3}=2$. (d) $p_{d_1}=2, p_{d_2}=1, p_{d_3}=1$. Both (b) and (c) are unsmoothed; (d) globally and locally smoothed.

of direction is unimportant. The parameter p_d controls the magnitude of contrast enhancement, shown in Equation 2: larger p_d produces a higher contrast.

In our implementation, we iterate the enhancement process three times. Figure 4 shows the effect of contrast enhancement. Each iteration uses a separate $p_{d_1}, p_{d_2}, p_{d_3}$ to adjust the contrast. Figure 4 (c) uses a stronger enhancement than Figure 4 (b). Both (b) and (c) clearly show the structure of the image, but (c) presents more weak edges as well as more noise. Thresholding the enhanced image yields many unwanted elements: hence, we add a smoothing step, described next.

3.3. Smoothing

Unconstrained filtering as described above enhances noise excessively. A proper black and white style demands a certain abstraction, here accomplished with smoothing. Also, we want to avoid distortion: we would like to be as faithful as possible to the image content, leaving further corrections to a final post-processing stage.

Global Smoothing: Before starting contrast enhancement, we apply a Gaussian filter to attenuate high frequencies. We compute $G\prime(x,y)=\sum(I*G)(x,y)$ where $*$ denotes discrete convolution and $G(x,y)=\frac{1}{2\pi\sigma^2}e^{-\frac{x^2+y^2}{2\sigma^2}}$. The window size $W_G=5$ unless otherwise stated. The image $G\prime$ then proceeds to later processing. This initial Gaussian smoothing provides global control over the abstraction.

Local Smoothing: We interleave local smoothing operations with contrast enhancement. To determine where local

Figure 5: Comparison. Left: Mould and Grant's method [MG08]; middle: Our method without local smoothness; left: Our method with local smoothing using $T_{\sigma_1}=5$, $T_{\sigma_2}=20$, and $T=150$.

smoothing should be applied, we compute the standard deviation of intensity along each direction and find the maximum. Where the maximal standard deviation is large, the pixel is in a region with a lot of texture or edges; we should not aggressively smooth such areas. We employ two threshold values T_{σ_1} and T_{σ_2}, with $T_{\sigma_2} >> T_{\sigma_1}$. Our recommended process uses three iterations. For the first two iterations, if $\sigma_{\max} > T_{\sigma_1}$, we use p_{d_1} and p_{d_2} to enhance contrast. In the third iteration, we enhance contrast (by p_{d_3}) where $\sigma_{\max} > T_{\sigma_2}$; elsewhere, pixels are assigned the local intensity $\bar{I}(x,y)$. This process first enlarges features and then smooths areas with little intensity variation, such as the face area or the background sky. We use values chosen empirically: $T_{\sigma_1}=5$ and $T_{\sigma_2}=20$ unless otherwise stated.

The impact of smoothing is shown in the thresholded images in Figure 5. With smoothing, the shadow line on the face is clean, while the unsmoothed version is jagged. Figure 4 (d) shows the output before thresholding.

Post-processing: Like previous black and white methods, we have a final post-processing step. Here, we use Potrace [Sel15] on the thresholded image, further smoothing and removing small isolated regions. The final output is vectorized.

In summary, our algorithm converts an input color image to a greyscale image, blurred using a $W_G \times W_G$ Gaussian filter. We then compute the standard deviation and mean of intensity along each direction. When the maximal standard deviation exceeds a threshold, we enhance contrast along the maximal direction. If the maximal standard deviation does not exceed a second threshold, the pixel is smoothed by replacing with its neighborhood average. The preceding operates over three iterations. During the first two iterations we enhance contrast as much as possible, using a lower threshold so as to preserve structural detail. In the last iteration, we attempt to increase smoothness, and so use a larger threshold (a second threshold), refining some areas. Then we threshold the enhanced greyscale image. In a final step, small regions in the binary image are removed and the output is vectorized and further smoothed using Potrace.

Figure 6 shows a comparison with earlier black and white methods. Our results demonstrate high-quality structural

(a) Original image (b) XDoG [Win11, WKO12] (c) Mould and Grant [MG08] (d) Our method

(e) Original image (f) XDoG [Win11, WKO12] (g) Mould and Grant [MG08] (h) Our method

Figure 6: Comparison. (d) $p_{d_1} = 2, p_{d_2} = 2, p_{d_3} = 1$, $T_{\sigma_1} = 5$, $T_{\sigma_2} = 20$, and $T = 100$; (h) $p_{d_1} = 2, p_{d_2} = 1, p_{d_3} = 1$, $T_{\sigma_1} = 5$, $T_{\sigma_2} = 20$, and $T = 90$.

Figure 7: Thick edges using larger filters, here $W_G = 7$ and $N = 15$.

preservation. Another result, in Figure 7, shows how larger-scale filtering can generate results with thicker edges and a higher level of abstraction.

3.4. Effects

The contrast enhancement provides a great deal of structural detail. Based on the procedure above, we demonstrate some variations including line drawing and hatching.

3.4.1. Overall Tone Adjustment

Our algorithm maintains the structure very well, emphasizing edges both as region boundaries and as local details. The final black and white image is obtained by thresholding, and varying the threshold level is an obvious means of controlling the overall image darkness. Figures 8 (d-f) use the same set of parameters but different thresholds T. A standard result uses $T = 128$; an effect closer to line drawing has $T = 0$. Both these thresholds show the portrait very well, much better than the results from Xu and Kaplan [XK08]. With a very large threshold, most of the image is black: the result is perhaps unappealing, but the content is still clear. Another example of threshold variation appears in Figure 9.

3.4.2. Line Drawing

Line drawings can be produced by setting $T = 0$: after iterating the contrast adjustment a few times, the pixels along the edges will have converged to black or white. A line drawing could be viewed as the limiting case of this convergence, all tones having been pushed to the two extremes. To create line drawings with few iterations, we set the enhancement parameters to large values: $p_{d_1} = 4, p_{d_2} = 2$, and $p_{d_3} = 2$, and the smoothing parameters to small values: $T_{\sigma_1} = 4$ and $T_{\sigma_2} = 4$. This will produce attractive line drawings with clean and long edges, as shown in the examples in 8(f) and 9(f).

3.4.3. Hatching

One form of hatching uses patterns of strokes, such as sets of parallel lines, to indicate material or shading. Here, we introduce directional patterns to regions with specific intensity levels. Figure 10 and Figure 11 show hatching exam-

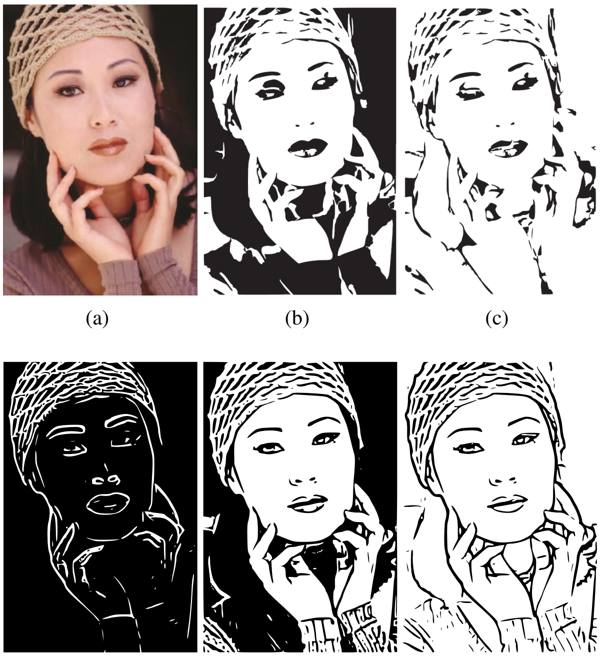

(a) (b) (c)

(d) $T = 220$ (e) $T = 128$ (f) $T = 0$

Figure 8: Tone adjustment. Top: tone control by Xu and Kaplan [XK08]; Bottom: our method; left to right, $T = 200, T = 128, T = 0$. Lower centre and left use $p_{d_1} = 2, p_{d_2} = 1, p_{d_3} = 1$.

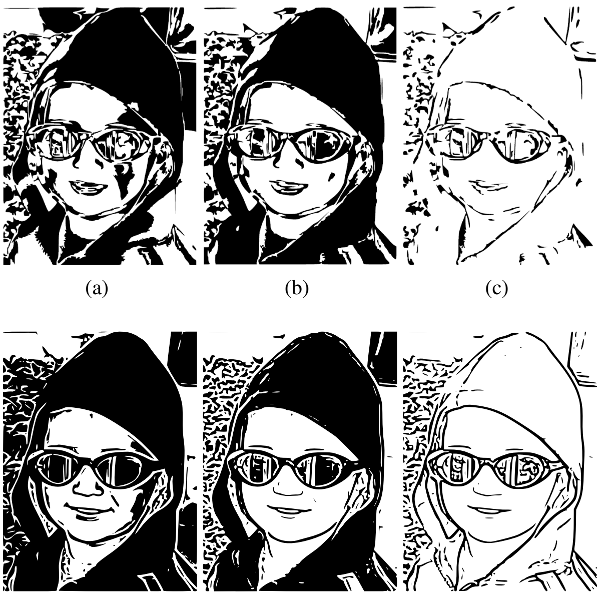

(a) (b) (c)

(d) $T = 200$ (e) $T = 128$ (f) $T = 0$

Figure 9: Tone adjustment. Top: tone control by Xu and Kaplan [XK08]; Bottom: our method; left to right, $T = 200, T = 128, T = 0$. Lower centre and left use $p_{d_1} = 2, p_{d_2} = 2, p_{d_3} = 1$, $T_{\sigma_1} = 5$, and $T_{\sigma_2} = 20$.

Figure 10: Hatching. Left: Original image; centre and right: two different hatching effects.

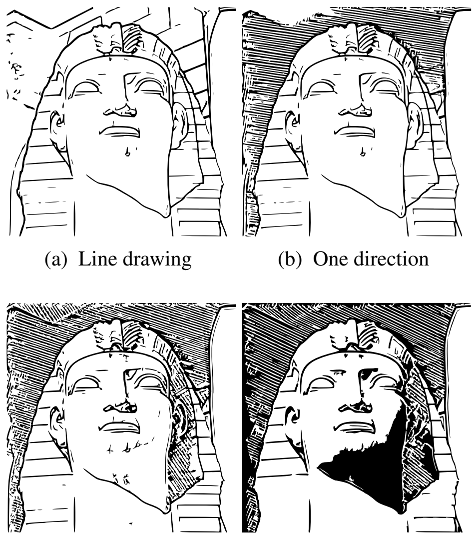

(a) Line drawing (b) One direction

(c) A few directions (d) With black adjustment

Figure 11: Varied effects from line drawing. Clockwise from upper left: plain line drawing; hatching with one direction; hatching with two directions; hatching with tone adjustment

ples with fake contrast. Figure 11 (b) introduces one directional pattern to a dark region with an average intensity $\bar{I}(x,y)$ below 60. The pattern is formed by propagating the adjustment of the current pixel along a pre-specified direction. The adjustment is done in-place, akin to error diffusion, so that adjustments made in one step influence later pixels in raster order. The interaction between image edges and the pre-specified direction produces alternating black and white lines, as sometimes a positive adjustment is propagated, sometimes negative. The method can use multiple directions: Figure 10 uses one direction for intensities below 60 and a second direction when the average intensity $\bar{I}(x,y)$ is between 60 and 100. However, the effectiveness of this technique depends hugely on the choice of parameters, which must be customized for the input image, and considerable experimentation may be needed.

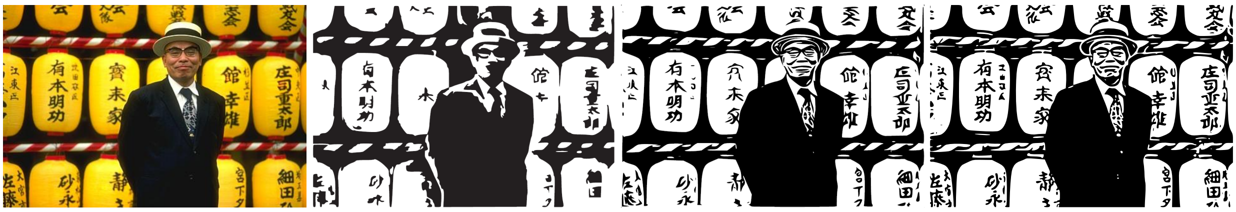

(a) Original image (b) Artistic thresholding [XK08] (c) XDoG [Win11, WKO12] (d) Our black and white

Figure 12: Comparison. (d) Our method uses $p_{d_1} = 1, p_{d_2} = 1, p_{d_3} = 1$, and $T = 90$.

(a) Original image (b) Mould and Grant [MG08] (c) XDoG [Win11, WKO12] (d) Our black and white

Figure 13: Comparison. Image (c) courtesy of Holger Winnemöller.

(a) (b) (c) (d) (e) (f)

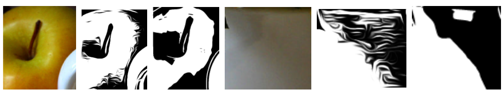

(g) (h) (i) (j) (k) (l)

Figure 14: Zoomed-in comparisons with XDoG. Top row: a zoomed-view of Figure 12. Bottom row: a zoomed-view of Figure 13. (a), (d), (g), and (j) are original images; (b), (e), (h), and (k) are results from XDoG; (c), (f), (i), and (l) are results from our method.

4. Results and Discussion

In this section, we show some additional results and comparisons with previous methods. We begin with XDoG. The XDoG-based method for black and white stylization produces extremely high-quality structural preservation. Our method is comparable, and for linear features, superior. Figures 12 to 16 give some comparisons with the XDoG-based method. Figure 14 shows a close-up.

(a) XDoG (b) Ours (c) Ours

Figure 15: Left: XDoG [Win11, WKO12]. Middle: our method with $p_{d_1} = 1.5, p_{d_2} = 1, p_{d_3} = 1, T = 50$. Left: $p_{d_1} = 2, p_{d_2} = 2, p_{d_3} = 1, T = 45$.

In Figure 12 (d), our result clearly and correctly shows the characters on the lanterns. XDoG omits some strokes, smooths out the corners, and extends short strokes to longer strokes, producing incomplete or deformed characters. The top row of Figure 14 shows more detail. In Figure 15 (c), our result shows small details such as the texture of the hat and highlights in the eyes. We can also adjust the abstraction level and make it close to the abstraction of the XDoG-based method. See Figure 15 (b), where we use a large mask for smoothing ($W_G = 7$, $N = 11$). Figure 16 shows how XDoG sometimes misses key structures such as the silhouette of the arm and uses multiple tones to reproduce the content. Our method can clearly and delicately represent the detail.

In addition, XDoG applied to black and white has difficulty representing areas with gradual tone changes. In such areas, the region boundaries are uncertain, and spurious fea-

(a) Original (b) XDoG (c) XDoG (d) Ours (e) Ours

Figure 16: Comparisons with XDoG. Left to right: original image; XDoG with high abstraction and multiple tones; XDoG with low abstraction and multiple tones; our line drawing ($p_{d_1} = 4, p_{d_2} = 2, p_{d_3} = 1, T = 0$); our black and white ($p_{d_1} = 2, p_{d_2} = 2, p_{d_3} = 1, T = 50$).

Figure 17: Above: source (©Flickr.com user Carsten Ullrich); artistic thresholding [XK08]. Below: our result ($p_{d_1} = 2, p_{d_2} = 2, p_{d_3} = 1, T = 128$).

tures appear, possibly due to the flow-based extension. Figure 13 (c) and the close-up view in Figure 14 (h) and (k) show this problem: the surface of the apple and the shadow do not reveal the structure well. Our method can show the regions with solid black because our local smoothness unifies the shape, visible even in binary tones. Similar artefacts appear around the neck of the man in Figure 15 (a). Our results in Figure 15 (b) and (c) have solid region boundaries around the neck.

XDoG is a powerful technique, and we think its smoothing step is highly suitable for promoting abstraction. In cases with unclear but flowing features, such as hair, XDoG may be preferable because of its ability to extend edges. However, our method is more faithful to the original structure, particularly in cases of texture, sharp corners, and crossing edges.

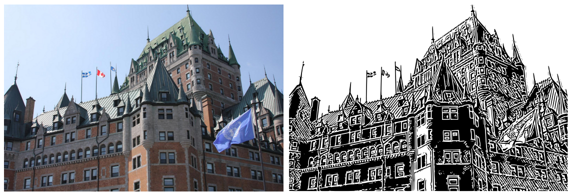

Figure 18: Line drawing with tone adjustment. Left: original image; right: our method.

We previously gave a visual comparison with results from Mould and Grant's method [MG08] in Figure 6. Both Figures 6 (d) and (h) capture the hair very well. The quality of the eyes in our result is superior to Mould and Grant's. Subfigures (g) and (h) offer an even clearer comparison: fine structure in the beard is preserved by our method, but is mostly missing from Mould and Grant's result.

Our method preserves texture, demonstrated in Figure 17. The original image has considerable texture, including leaves, tree trunks and branches, and the distant buildings. Our method delicately preserves all of them, unlike the result from artistic thresholding [XK08]. Figure 18 shows a building converted into a black and white image, with fine architectural details and clear overall structure.

As noted by Xu and Kaplan [XK08], region-based black and white methods cannot easily provide an "exclusive-or" effect. In our method, maintaining the contrast along the edges does create an "exclusive-or" effect. Figure 19 shows a comparison between the result from Xu and Kaplan and ours. Our method clarifies the edges around the tree trunk and the twigs, nicely separating the tree from the background building. Our result also shows the vertical texture on the wall and the detail on the door, providing a certain degree of texture indication, entirely missing from Xu and Kaplan's result. Another example of the "exclusive-or" effect is shown in Figure 20. Consider the lower part of the longest stem in the image. The stem occludes some leaves, but both the foreground stem and the background leaves are black. We surround the stem's silhouettes at that point with white edges. In this way, we can clearly see both the foreground and background objects.

Coherent line drawing [KLC07] is commonly used in NPR as an initial stage to capture edges. Smoothing along the Edge Tangent Field (ETF) yields high-quality edges, butcan produce unnatural swirling artefacts and rounded corners. Our line drawing effect has comparable quality and detail to coherent line drawing, but also preserves sharp corners, which are commonly lost in coherent line drawing. Figure 21 shows a comparison. Our method has a very close quality to that of coherent line drawing. Figure 22 shows a close-up view of Figure 21. The corners of the windows are rounded in Figure 22 (b) and are quite sharp in our result

(a) Original

(b) Artistic thresholding [XK08]

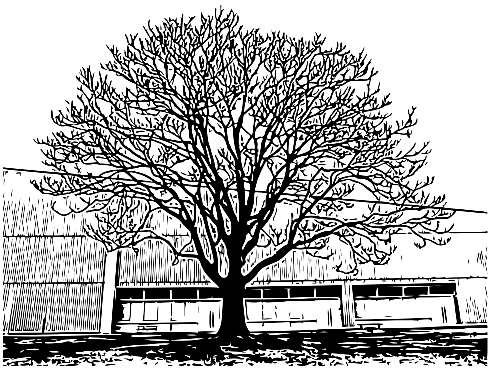

(c) Ours

Figure 19: An example of the "exclusive-or" effect. Above: original image; artistic thresholding [XK08]. Below: our result ($p_{d_1} = 1, p_{d_2} = 1, p_{d_3} = 1, T = 130$).

Figure 20: The long black stem is separated from the black leaves by the silhouette in white. Left: source (©Flickr.com user liz west); right: our result.

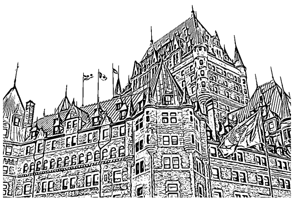

(a)

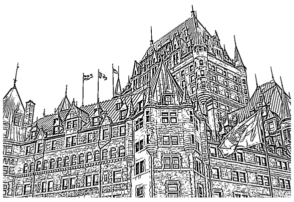

(b)

Figure 21: Line drawing comparison. The original image appears in Figure 18. (a) Coherent line drawing [KLC07]; (b) Our line drawing without using Potrace.

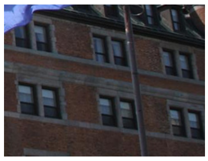

(a)

(b)

(c)

Figure 22: A zoomed-in view of Figure 21. Look at the corners of the windows. (a) Original image; (b) Coherent line drawing [KLC07] with rounded corners; (c) Our result, having a more natural appearance with sharp corners.

Figure 23: Black and white from our method. Left: original; right: processed with $p_{d_1} = 2, p_{d_2} = 1, p_{d_3} = 1$ and $T = 90$.

(c) – a subtle effect, but apparent on examination. More aggressive smoothing, used in some results, further degrades corners and line crossings. Using the maximum instead of the average allows us to avoid this problem.

Figure 23 provides an additional example using a binary palette. The black and white example shows the roof details, the distant building, the windows, and the people on the street very convincingly. Figure 1 shows more results.

Limitations. Our black and white algorithm demands multiple iterations and calculations for each of several possible orientations, and hence is slower than the XDoG method. Like other previous black and white methods, our method also demands the configuration of a set of parameters. While the parameters for the line drawing effect are very stable and robust, and the suggested parameters for black and white images rarely need modification, the hatching, unfortunately, is fragile and its parameters need careful tuning. Also, hatching is not effective in some cases because the fake hatching lines may cross the original edges, harming the content. A general issue with all variants is that excessive contrast enhancement can produce undesired double contours.

The thresholding operation can introduce spurious edges within smooth gradients. Sometimes this helps to convey shape, as when the new edge is an isophote on a shaded object, but other times the new edge is a distraction. Arguably the main limitation of our method is its focus on faithful structure preservation: in the tradeoff between abstraction and detail, it favors detail. We hope that the method can be a basis for future work that produces greater abstraction.

5. Conclusions

We propose a contrast-preserving method to create black and white effects. The main purpose is to preserve structural content. In future we plan to explore post-processing styles, such as hand-made drawing and hatching. Applying media to the existing black and white or line drawing effects such as painting the edges in ink or watercolor would be interesting. If possible, semantic information may be useful to further simplify the detail with thick strokes or user-defined strokes. We expect other researchers could use our algorithm as an initial stage to create other styles. Hatching needs further exploration.

Previous researchers have explored color harmony, but black and white images cannot easily benefit from such research. However, contrast harmony might be a related idea; excessive contrast can be visually distracting, while insufficient contrast makes an image dull. How best to characterize and obtain a pleasing contrast could be a fruitful future direction.

References

[BBDA10] BUCHHOLZ B., BOUBEKEUR T., DECARLO D., ALEXA M.: Binary shading using appearance and geometry. *Computer Graphics Forum 29*, 6 (2010), 1981–1992. URL: http://dx.doi.org/10.1111/j.1467-8659.2010.01712.x, doi:10.1111/j.1467-8659.2010.01712.x. 2

[CJO98] CZERWINSKI R., JONES D., O'BRIEN W.D. J.: Line and boundary detection in speckle images. *Image Processing, IEEE Transactions on 7*, 12 (Dec 1998), 1700–1714. doi:10.1109/83.730381. 2, 3

[EAP07] ERAMIAN M. G., ADAMS G. P., PIERSON R. A.: Enhancing ultrasound texture differences for developing an in vivo "virtual histology" approach to bovine ovarian imaging. *Reproduction, Fertility and Development 19*, 8 (2007), 910–924. URL: http://www.publish.csiro.au/nid/44/paper/RD06167.htm. 3

[GRG04] GOOCH B., REINHARD E., GOOCH A.: Human facial illustrations: Creation and psychophysical evaluation. *ACM Trans. Graph. 23*, 1 (Jan. 2004), 27–44. URL: http://doi.acm.org/10.1145/966131.966133, doi:10.1145/966131.966133. 2

[HS13] HALL P., SONG Y.-Z.: Simple art as abstractions of photographs. In *Proceedings of the Symposium on Computational Aesthetics* (New York, NY, USA, 2013), CAE '13, ACM, pp. 77–85. URL: http://0-doi.acm.org.libcat.uncw.edu/10.1145/2487276.2487288, doi:10.1145/2487276.2487288. 2

[KK11] KYPRIANIDIS J. E., KANG H.: Image and video abstraction by coherence-enhancing filtering. *Computer Graphics Forum 30*, 2 (2011), 593–602. URL: http://dx.doi.org/10.1111/j.1467-8659.2011.01882.x, doi:10.1111/j.1467-8659.2011.01882.x. 2

[KLC07] KANG H., LEE S., CHUI C. K.: Coherent line drawing. In *Proceedings of the 5th International Symposium on Non-photorealistic Animation and Rendering* (New York, NY, USA, 2007), NPAR '07, ACM, pp. 43–50. URL: http://doi.acm.org/10.1145/1274871.1274878, doi:10.1145/1274871.1274878. 2, 8, 9

[LR13] LAI Y.-K., ROSIN P. L.: *Image and Video based Artistic Stylisation*. Springer, London, Heidelberg, 2013, ch. Non-photorealistic Rendering with Reduced Colour Palettes. 2

[MG08] MOULD D., GRANT K.: Stylized black and white images from photographs. In *Proceedings of the 6th International Symposium on Non-Photorealistic Animation and Rendering 2008, Annecy, France, June 9-11, 2008* (2008), pp. 49–58. URL: http://doi.acm.org/10.1145/1377980.1377991, doi:10.1145/1377980.1377991. 1, 2, 4, 5, 7, 8

[MZZ10] MENG M., ZHAO M., ZHU S.-C.: Artistic papercut of human portraits. In *Proceedings of the International Conference on Multimedia* (New York, NY, USA, 2010), MM '10, ACM, pp. 931–934. URL: http://0-doi.acm.org.libcat.uncw.edu/10.1145/1873951.1874116, doi:10.1145/1873951.1874116. 1, 2

[NB06] NOBLE J., BOUKERROUI D.: Ultrasound image segmentation: a survey. *Medical Imaging, IEEE Transactions on 25*, 8 (Aug 2006), 987–1010. doi:10.1109/TMI.2006.877092. 3

[RC13] ROSIN P., COLLOMOSSE J.: *Image and Video based Artistic Stylisation*. Springer, London, Heidelberg, 2013. 2

[RL10] ROSIN P. L., LAI Y.-K.: Towards artistic minimal rendering. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (New York, NY, USA, 2010), NPAR '10, ACM, pp. 119–127. URL: http://doi.acm.org/10.1145/1809939.1809953, doi:10.1145/1809939.1809953. 1, 2

[RL13] ROSIN P. L., LAI Y.-K.: Artistic minimal rendering with lines and blocks. *Graph. Models 75*, 4 (July 2013), 208–229. URL: http://dx.doi.org/10.1016/j.gmod.2013.03.004, doi:10.1016/j.gmod.2013.03.004. 1, 2

[Sel15] SELINGER P.: Potrace. http://potrace.sourceforge.net/, 2001-2015. 4

[SHB07] SONKA M., HLAVAC V., BOYLE R.: *Image Processing, Analysis, and Machine Vision*. Cengage Learning; 3 edition (March 19, 2007), 2007. 1

[Win11] WINNEMÖLLER H.: XDoG: Advanced image stylization with extended difference-of-gaussians. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering* (New York, NY, USA, 2011), NPAR '11, ACM, pp. 147–156. URL: http://doi.acm.org/10.1145/2024676.2024700, doi:10.1145/2024676.2024700. 1, 2, 5, 7

[WKO12] WINNEMÖLLER H., KYPRIANIDIS J. E., OLSEN S. C.: XDoG: An extended difference-of-gaussians compendium including advanced image stylization. *Computers & Graphics 36*, 6 (2012), 740–753. doi:10.1016/j.cag.2012.03.004. 1, 2, 5, 7

[XK08] XU J., KAPLAN C. S.: Artistic thresholding. In *Proceedings of the 6th International Symposium on Non-photorealistic Animation and Rendering* (New York, NY, USA, 2008), NPAR '08, ACM, pp. 39–47. URL: http://doi.acm.org/10.1145/1377980.1377990, doi:10.1145/1377980.1377990. 1, 2, 5, 6, 7, 8, 9

[XKM07] XU J., KAPLAN C. S., MI X.: Computer-generated papercutting. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications* (Washington, DC, USA, 2007), PG '07, IEEE Computer Society, pp. 343–350. URL: http://dx.doi.org/10.1109/PG.2007.15, doi:10.1109/PG.2007.15. 1, 2