Detail and Color Enhancement in Photo Stylization

Rosa Azami
Carleton University
rosaazami@cmail.carleton.ca

David Mould
Carleton University
mould@scs.carleton.ca

Figure 1: We demonstrate a method to enhance the color and details of an abstracted image based on the solution of the Poisson equation. We control the brightness of the abstracted image and emphasize on details where subjects located based on a smooth transition map. The left image is an input, and in right we show our result of color and detail enhancement. The highlighted textures of the grass, antlers and fur on the body of the deer are picked by our algorithm. Our approach in brightening the images showed interesting colorization effects. Original *deer* image ©Peter Pham, CC BY 2.0.

ABSTRACT

Abstraction in non-photorealistic rendering reduces the amount of detail, yet non-essential details can improve visual interest and thus make an image more appealing. In this paper, we propose an automatic system for photo manipulation that brightens an image and alters the detail levels. The process first applies an edge-preserving abstraction process to an input image, then uses the residual to reintroduce and exaggerate details in areas near strong edges. At the same time, image regions further from strong edges are brightened. The final result is a lively mixture of abstraction and enhanced detail.

CCS CONCEPTS

• **Computing methodologies** → **Non-photorealistic rendering**;

KEYWORDS

Image stylization, abstraction, texture indication

CAe'17, July 28-29, 2017, Los Angeles, CA, USA

ACM ISBN 978-1-4503-5080-8/17/07…$15.00
https://doi.org/10.1145/3092912.3092917

ACM Reference format:
Rosa Azami and David Mould. 2017. Detail and Color Enhancement in Photo Stylization. In *Proceedings of CAe'17, Los Angeles, CA, USA, July 28-29, 2017,* 11 pages.
https://doi.org/10.1145/3092912.3092917

1 INTRODUCTION

Photo abstraction, as practiced in image-based artistic rendering, has traces in the traditional arts. Pablo Picasso, Henri Matisse, Hans Hofmann and other abstract expressionists tried to connect to a viewer through simple but effective visual effects. For example, Matisse paintings are characterised by flat shapes, controlled lines, and pointillism. In non-photorealistic rendering, abstraction methods often remove high-frequency details from images. Animal fur or whiskers, bumps on rocks, and facial wrinkles are examples of details at risk of being lost through abstraction. Although removing detail is the essence of abstraction, doing so can also be a barrier concealing the real nature of an image. For example, an aged face may not be distinguishable from a younger face. Moreover, textures and details bring believability and vigor to images, as they do to realistic paintings. We would like to recover and in some cases exaggerate details lost through abstraction. This paper proposes a hybrid detail enhancement and abstraction process, where an input image is first abstracted with edge-aware filtering, and then a subset of the residual is added back to the filtered image; small residual elements and those further from strong edges are suppressed, while extended elements close to strong edges are enhanced. This process is intended to mimic the artistic notion of *indication* [Guptill

2014], where image elements such as texture are suggested with scattered details rather than being represented in full. Within the same process, we brighten the photograph such that the areas of interest will get increased contrast while colors in the areas of less interest will become lighter and less saturated. The combination of selective detail enhancement, image brightening, and image abstraction generates a new stylized image. An example is shown in Figure 1.

Our paper contributions include the following:

- We present an automatic algorithm to locally emphasize image details: details near strong edges are enhanced and those far from edges are reduced.
- We define a system for scoring texture elements, used to establish different levels of texture indication.
- We propose a simple brightening method to be used in conjunction with detail enhancement.

The remainder of this paper is organized as follows. Section 2 reviews related work on image abstractions methods, Poisson equation, and brightness and contrast enhancement. We present our algorithm in section 3. Section 4 discusses the details of our results and compare them with some existing methods in contrast enhancement. Finally, Section 5 concludes this paper and gives possible directions for the future work.

2 RELATED WORK

Many image filters, such as Gaussian convolution or median filtering, remove high-frequency features such as small textures. An audience might prefer to preserve such details, at least in part. Our stylized filter takes inspiration from three different areas of past work. We start by reviewing past work on image abstraction. Next, we give a brief review of Poisson image editing, used for our color shifting method. Last, we discuss some previous approaches to relighting and color shifting a single input image.

2.1 Image Abstraction

Many abstraction methods have concentrated on artistic stylization of 2D content. Broadly speaking, abstraction of photographs can be divided into *stroke-based* techniques [Collomosse and Hall 2002, 2005; Haeberli 1990; Hertzmann 1998] and *filter-based* techniques [Bai and Sapiro 2007; Criminisi et al. 2010; Kyprianidis et al. 2009; Mould 2013; Papari et al. 2007; Paris et al. 2008; Tomasi and Manduchi 1998]. We take a filtering approach in this work.

The bilateral filter [Paris et al. 2008; Tomasi and Manduchi 1998] smooths low-contrast regions while preserving high-contrast edges. It involves computing a custom arrangement of weights for each pixel, where the distance of the central pixel to each other pixel is a combination of spatial distance and color space distance. This versatile filter is sometimes a component of abstraction algorithms, such as that of Winnemöller et al. [Winnemöller et al. 2006].

Bilateral filters were not intended to produce an abstraction effect and using them for abstraction by themselves is problematic, as they may preserve isolated features and high-contrast textures. The bilateral filter has difficulty preserving edges of different amplitudes under a single parameter setting. Also, there is a tradeoff between preserving edges and preserving noise or fine details.

Geodesic approaches [Bai and Sapiro 2007; Criminisi et al. 2010], consider a 2D image as a 3D surface and compute weighted distances from a pixel or group of pixels over the manifold. Geodesic filters penalize edges: when crossing an edge, the geodesic distance jumps suddenly but rises slowly thereafter. On the other hand, the geodesic distance rises quickly over textured areas, as numerous small gradients must be crossed [Mould 2013].

Kyprianidis et al. [Kyprianidis et al. 2009] presented a generalization of the Kuwahara filter [Papari et al. 2007]. Their painting-like method preserves shape boundaries but different from conventional edge-preserving filters, the flattening effect occurs along the local feature directions. This effect can cause an unpleasant deviation in extracted high-frequency details.

Mould [Mould 2013] offered an edge-aware filter called *cumulative range geodesic* filtering. In this approach, distance in the image plane is lengthened proportionally to the color distance between the current pixel and the starting pixel. Our proposed system uses Mould's filter as the default filtering process. We intend to enhance textures while preserving sharp edges, and rely on an edge-preserving abstraction method.

2.2 Brightness and Contrast Enhancement

We consider two types of methods for image illumination and contrast enhancement: histogram-based methods [Lee et al. 2013; Pisano et al. 1998; Puff D, Pisano E, Muller K, Johnston R, Hemminger B, Burbeck C, McLelland R 1994; Xu et al. 2014] and filtering-based methods [Bennett and McMillan 2005; Deng 2011; Fattal et al. 2002; Kass and Solomon 2010; Paris et al. 2011].

Histogram-based methods aim to generate an output image having a histogram with a target image. They usually are simple and used mostly to enhance low contrast images. However, there is a trade-off between contrast enhancement and noise amplification [Xu et al. 2014].

More flexibility can be seen in filtering-based methods. Some of these methods are combinations of optimization, segmentation and histogram-based methods [Deng 2011; Yuan and Sun 2012]. Some filter-based methods decompose image into a base layer and a detail layer, with each layer undergoing separate processing before being recombined.

Kass and Solomon [Kass and Solomon 2010] introduced an edge-preserving contrast enhancement technique, suggesting a multi-layer diffusion operator to enhancing image details.

The edge-aware Local Laplacian filter introduced by Paris et al. [Paris et al. 2011] also used the concept of separating base and detail layers. Although they generated high quality images, the exaggerated details in some soft areas like walls could be undesirable. Similarly, our method uses an abstracted image as the base layer in brightening process and a detail map to sharpen the textures.

Naturalness preservation and color enhancement were the main features of previous approaches, while we want to change a natural and uniform image to an artistic one by manipulating the colors. Recent work by Semmo et al. [Semmo et al. 2016] transforms images into an oil painted look. They extract representative colors from the image and assemble a palette, then quantize the input image using that palette. Sharp highlights appearing in the results imitate the

bumped surface of the oil paint on the canvas. A significant drawback of the quantization is that it can introduce hard transitions into previously smooth gradients.

Our color shifting process brightens the colors more strongly at locations further from prominent edges in the image. We solve a Poisson equation to obtain a map over the image, dictating where the residual will be suppressed vs. exaggerated, and telling us where to brighten the image and where to leave the colors unchanged.

2.3 Poisson Equation

The Poisson equation has been used effectively for image synthesis and processing [Elder and Goldberg 2001; Orzan et al. 2007; Pérez et al. 2003]. Elder and Goldberg [Elder and Goldberg 2001] introduced an image editing system in the contour domain. Their system reconstructs the image intensity from brightness at edges. They assumed that in regions other than edges, the intensity function approximately satisfies the Laplace equation $\Delta I(x, y) = 0$, and therefore solving the Laplace equation provides a good reconstruction of the original intensity. Perez et al. [Pérez et al. 2003] also used the Poisson equation as an interpolation mechanism. Their proposed method smoothly imports source image regions into a destination image and could affect the texture, illumination, and color of objects within the selected region. Orzan et al. [Orzan et al. 2007] used a Poisson reconstruction method for image stylization, controlling local detail levels in an input photograph.

Orzan et al. [Orzan et al. 2009] and Jeschke [Jeschke 2016] used the Poisson equation to interpolate colors smoothly on either side of a boundary curve. The Diffusion Curve Image (DCI) introduced by Orzan et al. [Orzan et al. 2009] specifies colors on either side of Bezier curves. Jeschke [Jeschke 2016] improved DCIs by introducing new boundary conditions midway between the diffusion curves.

Like previous methods, we solve the Poisson equation, using it for both brightening and detail enhancement. We obtain a smooth map over the image considering two main boundary conditions: edges of the image and the ridges of the distance transform map obtained from edges. Jeschke [Jeschke 2016] also used the term "ridges" to refer to curves midway between the diffusion curves.

The resulting map provides a smooth gradient between the edges and the ridges far from edges. We intend to keep the brightness close to edges intact; such regions mostly introduce main textures of the subjects. However, areas far from edges do not carry significant high-frequency details, so detail can be suppressed here.

3 PROPOSED METHOD

The smooth map obtained from the Poisson solution can be used to govern detail and brightness adjustment. We aim to generate a stylized photograph by recombining the high-frequency information while avoiding unfavorable overenhancement; Wang et al. [Wang et al. 2013] mention that extreme high-frequency details may result in an unnatural-looking photograph. In this section we present our solution for an automatic but controllable detail and brightness enhancement process.

Our algorithm has several steps. First, we extract residuals by subtracting the filtered image from the input image. To determine which detail elements should be emphasized and which eliminated, we use two post-processing steps; sticks filtering [Czerwinski et al. 1999] and scoring connected components. Next, we solve a Poisson equation to obtain a map that governs the magnitude of brightening and the exaggeration of the residuals. Finally, we combine the selected residuals with the brightened filtered image. Figure 2 shows the pipeline of our method with an example image.

Given an image I, we apply the cumulative range geodesic filter [Mould 2013]. We subtract the filtered image from the input image, yielding the residual map $R(x, y)$. The signed residuals then undergo additional processing before being recombined with the output image.

3.1 Sticks Filtering

Extended linear structures, like tree branches or facial wrinkles, are visually important yet may be suppressed even by an edge-preserving filter. In seeking to recover them, the variations in intensity of the residual pixels nearby may make it difficult to distinguish thin features from sparse and disconnected features. To enhance thin and connected line structures, we use the sticks filter, originally designed to reduce speckle noise and preserve linear structures [Czerwinski et al. 1999].

We use the sticks directions as a decision process. We create a sticks map $T(x, y)$ using equation 1:

$$T(x, y) = \max_{1 \le s \le i} \Delta(\mu(s_i) - \hat{I}_n(x, y)) \tag{1}$$

$$\hat{I}_n(x, y) = \frac{\sum_{(r,s) \in n \times n} I(r, s)}{n^2} \tag{2}$$

$$\mu(s_i) = \frac{\sum_{(j,k) \in s_i} I(j, k)}{n} \tag{3}$$

The sticks direction $T(x, y)$ determines the maximal contrast response, defined as the maximum difference between the average intensity $\mu(s_i)$ of a stick and the average intensity $\hat{I}_n(x, y)$ of the neighboring pixels. Parameter n is the length of each stick; there are i sticks. In our approach we used $n = 5$ for $i = 8$ sticks orientations.

3.2 Scoring Connected Components

Textured areas can be characterized by aggregations of high-frequency details, as opposed to sparse or low-frequency details. We estimate whether a given detail is part of a texture by assigning it a score based on the residual content in its neighbourhood.

We convert the grayscale residual map to a binary image and find its connected components [Serra 1998]. The residual image is thresholded, setting to zero all $R(x, y)$ such that $|R(x, y)| < 5$; the nonzero $R(x, y)$ are then considered foreground pixels in the later connected components calculations. Positive and negative components are scored separately. For each connected component C, we get a territory (everything within d pixels of any pixel in C; we used $d = 10$) and we sum the magnitudes of all residual pixels with the same sign as pixels in C. Thus, large components, with a large territory, as well as smaller components close to areas with large residuals, will receive large scores.

The scores will determine which portions of the residual are emphasized: we will sort the components in order of decreasing score, and the top portion will be exaggerated while the rest are discarded. We normally report the scoring threshold s_t as a relative position in the sorted list; e.g., we might discard all but the top 20%, using $s_t = 0.2$.

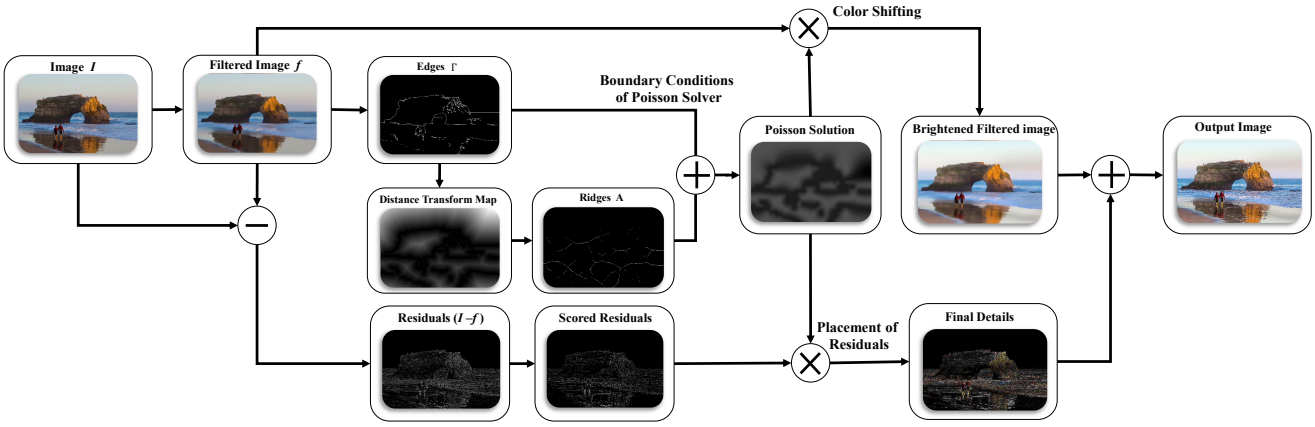

Figure 2: Overview of our method. Original *rock arch* image ©Vadim Kurland, CC BY 2.0.

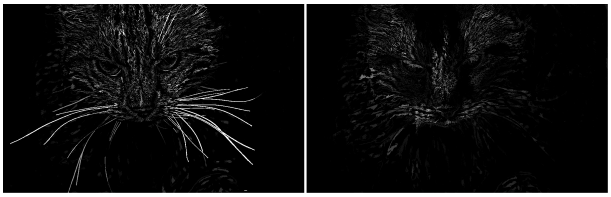

Figure 3: Scored connected components. Left: positive residuals; right: negative residuals.

Figure 3 shows an example of the scored connected components of the positive and negative signed residuals. Using the signed residuals allows us to better control the size of the connected components: individual elements are scored separately, rather than foreground and background being mixed together in a single component as might be the case when absolute values are used.

3.3 Poisson Solver

We plan to preserve the colors close to edges and brighten the image far away from the edges. We will use a Poisson solver to give us a map of where each pixel lies relative to the edges. A possible alternative is to compute the distance from edges directly, but this quantity would not be smooth and could be difficult to normalize properly.

We solve a Poisson equation F over a 2D domain Ω, with Dirichlet boundary conditions at edges Γ set to zero and a fixed value v on the ridges A. We extract edges Γ using the classical *Canny edge detector* [Canny 1986] and the ridges A are effectively the medial axis of the distance transform, that is, the pixels equidistant between edges. Thus, we have the following overdetermined linear system:

$$\begin{aligned} \Delta \mathrm{F(p)} &= 0 \qquad p \in \Omega \\ \mathrm{F(p)} &= 0 \qquad p \in \Gamma \\ \mathrm{F(p)} &= v \qquad p \in A \end{aligned} \tag{4}$$

The solution to the above will be a collection of smooth hills, with minimum height at the edges themselves (where the height is zero) and maximum at the ridges between the edges (where the height is v, set to 50 in all images unless mentioned otherwise).

We will use the map P obtained from solving the above system of equations both to tell us where to brighten the image and where the detail should be maximally enhanced. Brightening is greater where $P(x, y)$ is larger; details are enhanced most where $P(x, y)$ is low and least where $P(x, y)$ is greatest.

3.4 Image Brightening

The next step is to brighten the filtered image. Brightening the filtered image furthers our goals in two ways. First, the smooth transition of the Poisson map varies the brightness across the image and highlights subjects in focus. Second, a shift in colors presents that generates a gradient of colors that results in interesting effects. Our color shifting process is defined as follows:

$$I_{\text{bright}}(x, y) = (1 + b \times \Phi(x, y) \times \gamma) \times I_{\text{f}}(x, y) \tag{5}$$

$$\Phi(x, y) = \min(1 - I(x, y).\text{red}, 1 - I(x, y).\text{green}, 1 - I(x, y).\text{blue}) \tag{6}$$

where γ is a normalized value based on the Poisson solution at each pixel position and b is a user-chosen value that scales the level of brightening. We control the over-brightening of the image based on the values of the color channels for each pixel. In equation 6, $\Phi(x, y)$ is the minimum of $1 - c$ over all channels, where c is the normalized color for each channel. In equation 5, we attenuate the colors based on $\Phi(x, y)$. This operation is core of the color shifting process.

3.5 Detail Enhancement

When we recombine the details with the brightened filtered image, we want to keep the details with higher scores while avoiding to bring back the details with lower scores. We keep the details above a threshold score unchanged, while the details with lower scores will be reduced or eliminated.

Let $T(x, y)$ be the detail map obtained directly from the sticks decision process. We will compute a final detail map $D(x, y)$ governed by a scoring factor $0 < s(x, y) < 1$, output from the sticks decision process $T(x, y)$, a weight $\lambda(x, y)$ based on the normalized

Poisson distances $\hat{P}(x, y)$, and a user-selected constant k, controlling the extent of the detail exaggeration. The detail map $D(x, y)$ is computed as follows:

$$D(x, y) = k \times s(x, y) \times \lambda(x, y) \times T(x, y) \times R(x, y) \tag{7}$$

$$s(x, y) = \frac{(|\text{Score(x,y)} - s_t| - s_t)^2}{{s_t}^2} \tag{8}$$

$$\lambda(x, y) = (m - (m - 1)\hat{P}(x, y)) \tag{9}$$

Parameter $s(x, y)$ in equation 8 is a parabolic downward function that ranges between 0 and 1; its domain varies based on the scoring threshold s_t. If the score of a pixel (x, y) exceeds s_t its details will be preserved completely, otherwise the amount of the details to be preserved depends on the value of the $s(x, y)$.

The parameter $\lambda(x, y)$ ranges from a minimum of 1 far from the edges to a maximum of m on the edges. We used $m = 3$ for our results. Figure 4 shows different detail maps for $m = 3$, $m = 4$, $m = 5$, and $m = 10$; we can see that the magnitude of $D(x, y)$ increases with larger m. Note that both m and k can increase the details, but they work differently: m increases details locally while k acts globally. We set $k = 3$ for Figure 4.

Finally, we combine the brightened image and the detail map to produce the output:

$$I_{\text{out}}(x, y) = I_{\text{bright}}(x, y) + D(x, y) \tag{10}$$

Figure 5 shows the effects of filtering and then the outcomes from brightening the filtered image and from computing the detail map. The final image displays the outcome of merging the brightened base image and the detail layer. The filtering process removed the whiskers, eyebrows and many small hairs on the face of the cat. Next, the filtered image was brightened by a factor of $b = 2$, with the largest effect on the cat's chest, left shoulder, forehead, and eyes. We set the exaggeration parameter and the score threshold to $k = 3$ and $s_t = 0.5$, respectively. The brightened abstracted image and the detail map are merged in the last step, reintroducing the whiskers and other details on the cat.

4 RESULTS AND DISCUSSION

Our stylization method took advantage of the smooth Poisson solution map to brighten up the images with automatic control over the regions of interest. We created brighter-looking images while focusing the attention on subjects.

The non-uniform illumination made a shift in colors and created more illustrative and expressive results than the original colors. Changes in brightness are present in all images. An interesting illusion of motion appeared in some results, like the water in *swans*, *whale*, and *Venice* in Figure 6. The parameters selected depend on image properties and on the intended result. Moreover, independent parameters allow some manipulation of the results. Figure 6 shows some representative stylized images from our method. Parameters were set on a per-image basis. The brightness parameter used for the *rock arch*, *old lady*, *eye*, *swans* and the *rock formation* was $b = 1$; for all others, we set $b = 2$. The detail exaggeration parameter set to $k = 3$ for all images. The score threshold value was $s_t = 0.33$ for the *Venice*, $s_t = 0.92$ for *rock formation*, and $s_t = 0.5$ for the rest.

In *whale* we can see a small shift in color of the surrounding water, making the whale more distinct from the background. Similarly, the brightness on the face of the *old lady* shows a cheerful face that emphasizes details, such as the wrinkles. In the *eye* example, we created a more chromatic iris, yet the skin color around the eye remained normal. Also, the thin vessels on sclera and the eyelashes were preserved well by our detail enhancement process.

In our technique, the more details the original image contains, the more flexibility we have to manipulate them. The *Venice*, *skull rocks*, and *old lady* images contain a significant amount of details. However, we did not add any details to the flattened regions so as to keep the abstract background. The brightening effect and detail enhancement through the Poisson solution are visible in all results. Abstraction has its largest magnitude on the ridges where the Poisson solution has higher values; details are enhanced on edges, where the Poisson solution has its minimum. Note that the abstraction may not be very apparent near ridges, depending on the actual image content; for example, in a pure blue sky, there is no difference between the abstracted (smoothed) image and the original. To alter the balance between abstraction and detail exaggeration, we can vary the threshold of the connected component scores. The most abstracted image was the *rock formation*, using $s_t = 0.9$ and hence keeping only a few details. We found the combination of the smooth *skull rocks* and the sharp features of the shrubs to be quite striking. In *waterfall*, the tree, surface of the rocks, and green textures are emphasized, making the content of the image more clear.

Figure 7 shows results from the benchmark image set [Mould and Rosin 2016]. In the *barn*, the details of the barn, trees and some ground textures are enhanced. The beard and wrinkles of the *Yemeni* face are emphasized. The seared trees of the *desert* are exposed in front of a brighter background, and the fur and whiskers of the *cat* are enhanced.

The original Flickr images are illustrated in Figures 8; we also used some images from the NPRgeneral benchmark set [Mould and Rosin 2016].

4.1 Effect of Parameters

Our method has three user-chosen parameters with direct influence over the results: the brightening factor b, the detail exaggeration factor k, and the score threshold s_t.

We illustrated the effect of the brightening parameter b in Figure 11 for $b = 1$, and $b = 2$. The ceiling of the cave, the surface of the rocks and water are the regions most affected by brightening. However, the darkest area (lower left) does not experience significant changes, owing to the multiplicative brightening process: in general, areas with intensity near zero will not be brightened.

Results from different score thresholds are shown in Figure 9. The exaggeration parameter was set to $k = 3$ for all results. Choosing a higher threshold removes details with lower scores. From left to right, the first image is the original image, and the remaining images are enhanced with thresholds $s_t = 0.9$, $s_t = 0.75$, $s_t = 0.33$ and $s_t = 0.16$, respectively. In this case, more aesthetic images result from higher thresholds. By reintroducing only a few details, we can reveal the abstraction while emphasizing prominent details of the image.

Figure 10 shows results from different detail exaggeration parameters. Small details on the surface of the stones that are not

Figure 4: Effect of parameter m in $\lambda(x, y)$. Left to right: detail maps for $m = 3$, $m = 4$, $m = 5$, and $m = 10$.

Figure 5: Example steps of the enhancement process. Left to right: original image; filtered image; result from brightening the filtered image; detail map; and the final result from merging the brightened image and the detail map. Original *Angry cat* ©Peter Pham, CC BY 2.0.

pronounced for $k = 3$ become visible for $k = 5$ and $k = 7$, emphasizing the roughness of the stone surface. Conversely, while at $k = 3$ the stylized cat image is a good mix of abstraction and detail, for higher k the details on the face and forehead become too severe. In extreme cases ($k = 7$), overenhancement of the details manifests visually as noise; noise-like elements appear in the background around the cat's whiskers. In images such as the cat with large-scale coherent texture, more gentle detail exaggeration is warranted.

4.2 Comparisons

Our method is akin to the work by Bae et al. [Bae et al. 2006], who pioneered the base-detail decomposition that we also follow. Our approach can generate a more colorful and brighter image. In order to capture the detail layer they applied the *bilateral* filter but we chose the *cumulative range geodesic* filter since we want to partly abstract the image. They magnify high-frequency detail according to an estimate of textureness; our scores can be considered a textureness estimate as well, but we add a second factor to govern the detail amplification, which is the map of distance to a strong edge, enabling us to focus the detail towards the key image elements.

We compare the performance of our algorithm with results from state-of-the-art contrast enhancement methods in Figure 12. From left to right we see the original image of the *Stanford Memorial Church*, the LDR result from Lee et al.'s method [Lee et al. 2013], gradient domain HDR compression from Fattal et al.'s approach [Fattal et al. 2002], and the result from our algorithm.

The result from Lee et al. preserves the details well in the brighter area, while in dark regions on the ceiling, it loses visibility. The colors are dark and stuffy and it does not convey the original colors. The result by Fattal et al. is brighter and more colorful than that from Lee et al. The details are not emphasized everywhere, as we can see, for example, on the small and large arcs.

In our result, the contrast is enhanced naturally and the details are well preserved almost everywhere. The reader is encouraged to zoom in to better see the differences. The colors are much brighter than either the originals or the result from Lee et al. and the dark hidden regions showed up with a good amount of details. Also, the glasses, fence, and stairs are better emphasized than in Fattal et al.'s result. Note that in order to achieve a uniform level of detail enhancement, we could omit the local variation of detail enhancement.

Finally, we compare our result with the local Laplacian filters of Paris et al. [Paris et al. 2011]; see Figure 13. Although both local Laplacian filtering and the proposed algorithm brightened the original image, our algorithm provides better visibility in some regions. For instance, in the right corner of the room, there is a plant and a statue on a long table in front of it, which is unclear in the local Laplacian result. Also, the over-enhanced result from local Laplacian filtering exaggerates the noise on the wall as is apparent in the close-up images in the bottom row.

4.3 Timing

The computation time varies linearly with the number of pixels. Our algorithm takes about 3 minutes to process a 1024×640 image on a laptop Intel(R) Core(TM) i7-4510U with a 2.6 GHz CPU and 16.0 GB of RAM. Filtering and solving the linear system of the Poisson equation takes the majority of the time. The cumulative range geodesic filter takes less than a minute to process an image with a mask size of 500 and 22 seconds to process with a mask size 200. Solving the linear system is the most time consuming part of our algorithm, requiring almost 2 minutes. Our implementation is not optimized and there is a great deal of room to improve the speed. We used a naive edge detection in our work, while a cleaned up edge detection can decrease the number of edge pixels and solve the linear system faster.

Figure 6: Stylization effects. From upper left: *rock arch, whale, swans, old lady, eye, skull rocks, Venice, rock formation* and *waterfall.*

Figure 7: Brightened and detail-enhanced images from benchmark set [Mould and Rosin 2016] . Left to right: *barn, Yemeni, desert* and *cat.*

4.4 Limitations

We experienced some failure cases. The most common failure case for our method is the poor color shifting in regions far away from edges, such as the sky. The sky in *Venice* in Figure 14 is an example of this case. This deficiency could be improved by choosing a larger value for v in Poisson solver. We demonstrated an improved outcome for the *Venice* sky in Figure 6 using $v = 50$, while using $v = 20$ in Figure 14. Another possible direction to remedy this problem would be using a flexible value for v, such that the value of v on the ridge changes with respect to its distance transform value.

Another cause of failure is preserving too much texture; this depends on the choice of the input image. Images with many close edge points preserve most of the texture from the original image. We want the output result to partially contain the background abstracted image, so that the detail map shows the effect of enhancement. The rocks, ocean and grass of *Etretat* in Figure 14 show this issue; while overall it looks nice, the abstraction effect can hardly be seen. We may be able to address this issue with more sophisticated edge detection that can eliminate the redundant edges.

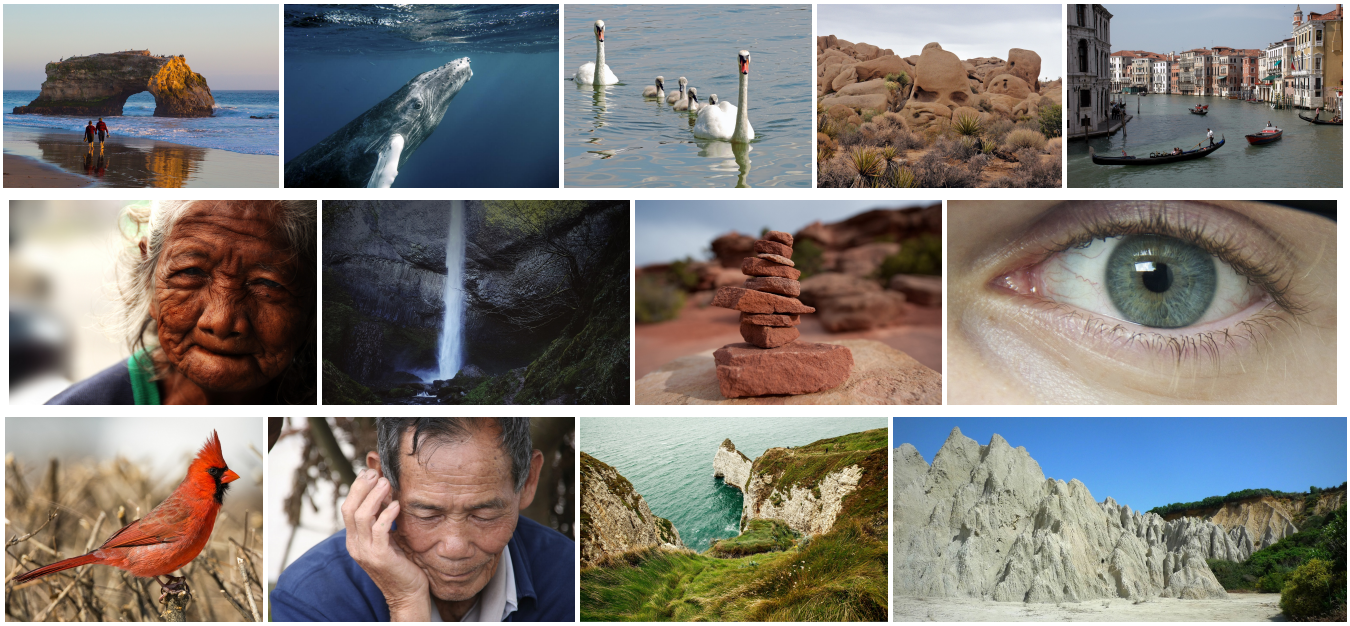

Figure 8: Original Flickr images used in our paper. From upper left: *rock arch* ©Vadim Kurland, CC BY 2.0, *whale* ©Christopher Michel, CC BY 2.0, *swans* ©Dun.can, CC BY 2.0, *skull rocks* by Joshua Tree National Park (public domain), *Venice* ©Pug Girl, CC BY 2.0, *old lady* ©Lemuel Cantos, CC BY 2.0, *waterfall* ©Loren Kerns, CC BY 2.0, *stones* by CanyonlandsNPS (public domain), *eye* ©Evan Schaaf, CC BY-SA 2.0, *northern cardinal* ©Bill Damon, CC BY 2.0, *Chinese* ©Anja Disseldorp, CC BY 2.0, *Etretat* ©Cristian Bortes, CC BY 2.0 and *rock formation* ©Henry Burrows, CC BY-SA 2.0.

Figure 9: Detail enhancement with different thresholding of the sorted connected component scores. Left to right: original, $s_t = 0.9$, $s_t = 0.75$, $s_t = 0.33$ and $s_t = 0.16$. Original *cow* image ©Alastair Campbell, CC BY-SA 2.0.

Also, a coarse transition of colors between adjacent pixels is perceived as noise in some images. This could be due to clamping of the colors in different channels. To address this issue we need a smooth transition in shifting the colors after augmenting the detail map; also, using a different color space may help with this limitation. In Figure 14 the noisy line of the fingers and the line above the shoulder of the old man shows this problem.

5 CONCLUSIONS

In this paper, we presented an automated image manipulation technique intended to give an appealing blend of detail and abstraction. In regions of the image near strong edges, assumed to represent relevant detail, we exaggerate detail, and in regions more distant from strong edges, we allow the abstraction effect to dominate. Also, we brighten the image far from edges, preserving the original image colors near edges. The net effect is a cheerful image in which some details are emphasized while others are suppressed.

We note a few possible directions for future work. In this paper, we operated in the RGB colorspace; treating color differently, including moving to a perceptually uniform colorspace to perform our calculations, is likely to be beneficial. We assumed that strong edges correspond to salient detail, but this assumption sometimes

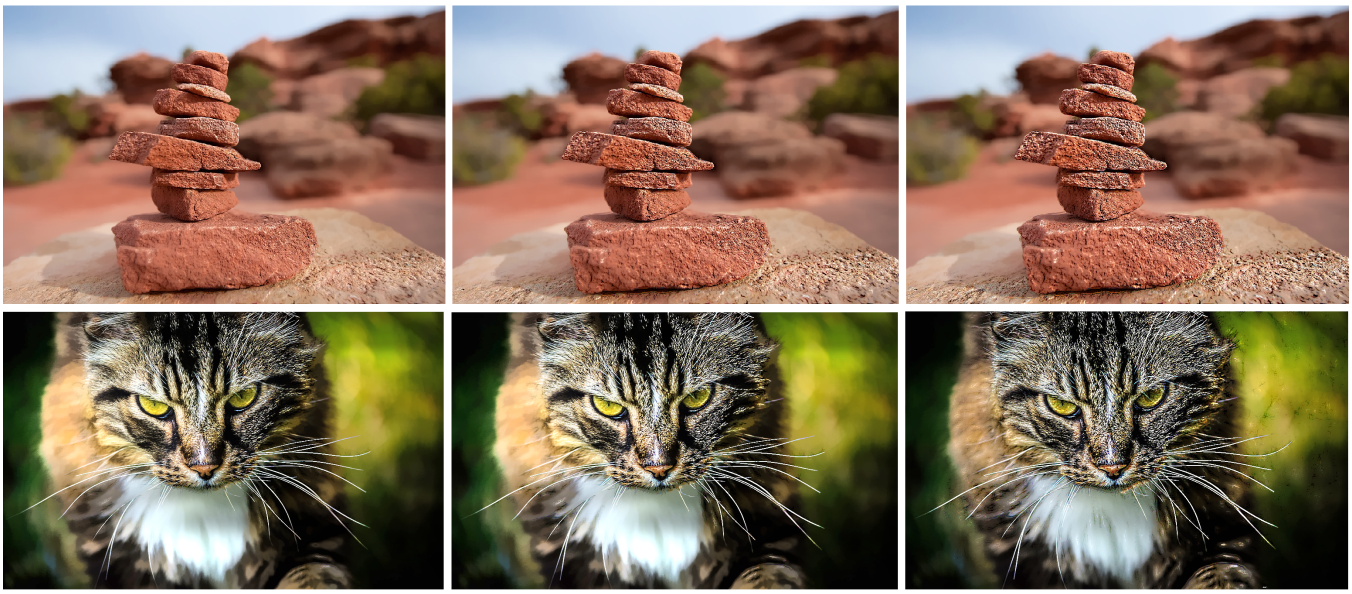

Figure 10: Different exaggeration parameters. Left to right: $k = 3$, $k = 5$, $k = 7$. Original *stones* image by CanyonlandsNPS (public domain) and *angry cat* image ©Peter Pham, CC BY 2.0.

Figure 11: Different brightness. Our results with b=1 (left) and b=2 (right). Original *Oparara* image (not shown) taken from benchmark set [Mould and Rosin 2016].

fails, and making direct use of a salience estimate is likely to be more helpful. Our ad-hoc scoring system is adequate, but a more elaborate scheme that better keeps track of context is likely possible. Finally, we made use of the sticks filter to repair linear features, but other shapes – circular features, say, or corners – might also be of interest.

ACKNOWLEDGMENTS

We thank GIGL members for their useful feedback on our work, and thanks to the anonymous reviewers for many insightful comments. Funding for this work was provided by NSERC and by the government of Ontario. We used many images from Flickr under a Creative Commons license. Thanks to the numerous photographers who provided material: Peter Pham (deer, angry cat), Vadim Kurland (rock arch), Bill Damon (northern cardinal), Christopher Michel (whale), Dun.can (swans), Joshua Tree National Park (skull rocks), Evan Schaaf (eye), Lemuel Cantos (old lady), Loren Kerns (waterfall), Pug Girl (Venice), CanyonlandsNPS (stones), Alastair Campbell (cow), Anja Disseldorp (Chinese man), Cristian Bortes (Etretat), and Henry Burrows (rock formation). Benchmark images were assembled by Mould and Rosin [Mould and Rosin 2016]; the photographs were taken by MrClean1982 (barn), trevorklatko (Oparara), Theen Moy (cat), Charles Roffey (desert), Richard Messenger (Yemeni). We also thank to Paul Debevec for giving us the right to use the original

Figure 12: Brightness comparison. Left to right: original image of Memorial Church , contrast enhancement of Lee et al. (LDR), gradient domain HDR compression of Fattal et al., our result. Original *Stanford Memorial Church* image ©Paul Debevec.

Figure 13: Comparison of detail enhancement. Top row, left image: original image; middle: local Laplacian filter of Paris et al.; right: our result. The bottom row shows a close-up of the wall. Original *Belgium House* image ©Raanan Fattal.

Stanford Memorial Church image and Raanan Fattal for the original Belgium House.

REFERENCES

Soonmin Bae, Sylvain Paris, and Frédo Durand. 2006. Two-scale Tone Management for Photographic Look. *ACM Transactions on Graphics* 25, 3 (2006), 637.

Xue Bai and Guillermo Sapiro. 2007. A Geodesic Framework for Fast Interactive Image and Video Segmentation and Matting. In *Proceedings of the IEEE International Conference on Computer Vision.*

Eric P. Bennett and Leonard McMillan. 2005. Video Enhancement Using Per-Pixel Virtual Exposures. *ACM Transactions on Graphics* 24, 3 (2005), 845.

John Canny. 1986. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1986).

John Collomosse and Peter M. Hall. 2002. Painterly rendering using image salience. In *20th Eurographics UK Conference, EGUK 2002.* Institute of Electrical and Electronics

Figure 14: Failure cases. Strong color shifts of sky in *Venice*; presence of many textures in *Etretat* hiding the abstract background; noise effect in *Chinese* man next to the edges.

Engineers Inc., 122–128.

John Collomosse and Peter M. Hall. 2005. Genetic Paint: A Search for Salient Paintings. In *Proc. EvoMUSART*. 437–447.

Antonio Criminisi, Toby Sharp, Carsten Rother, and Patrick Perez. 2010. Geodesic Image and Video Editing. *ACM Transactions on Graphics* 29, 5 (2010).

Richard N Czerwinski, Douglas L Jones, and William D O 'Brien. 1999. Detection of Lines and Boundaries in Speckle Images - Application to Medical Ultrasound. *IEEE Transactions On Medical Imaging* 18, 2 (1999), 126–136.

Guang Deng. 2011. Generalized Unsharp Masking Algorithm. *IEEE Transactions on Image Processing* 20, 5 (2011), 1249–1261.

James H Elder and Richard M Goldberg. 2001. Image Editing in the Contour Domain. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 3 (2001), 291–296.

Raanan Fattal, Dani Lischinski, and Michael Werman. 2002. Gradient domain high dynamic range compression. *ACM Transactions on Graphics* 21, July 2002 (2002), 249–256.

Arthur L. Guptill. 2014. *Rendering in Pen and Ink: The Classic Book On Pen and Ink Techniques for Artists, Illustrators, Architects , and Designers.* Potter/TenSpeed/Harmony.

Paul Haeberli. 1990. Paint by numbers: abstract image representations. *ACM SIGGRAPH Computer Graphics* 24, 4 (1990), 225–232.

Aaron Hertzmann. 1998. Painterly Rendering with Curved Brush Strokes of Multiple Sizes. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM, New York, NY, USA, 453–460.

Stefan Jeschke. 2016. Generalized Diffusion Curves: An Improved Vector Representation for Smooth-Shaded Images. *Computer Graphics Forum* 35, 2 (May 2016), 71–79.

Michael Kass and Justin Solomon. 2010. Smoothed Local Histogram Filters. *ACM Trans. Graph* 29, 4 (2010), 10 pages.

Jan Eric Kyprianidis, Henry Kang, and Jürgen Döllner. 2009. Image and Video Abstraction by Anisotropic Kuwahara Filtering. 28, 7 (2009).

Chul Lee, Chang Su Kim, and Chulwoo Lee. 2013. Contrast Enhancement Based on Layered Difference Representation of 2D Histograms. *IEEE transactions on image processing* 22, 12 (2013), 5372–84.

David Mould. 2013. Image and Video Abstraction Using Cumulative Range Geodesic Filtering. *Computers and Graphics (Pergamon)* 37, 5 (2013), 413–430.

David Mould and Paul L Rosin. 2016. A Benchmark Image Set for Evaluating Stylization. In *Expresive '16 Proceedings of the Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*. 11–20.

Alexandrina Orzan, Adrien Bousseau, Pascal Barla, and Joëlle Thollot. 2007. Structure-preserving manipulation of photographs. *Proc. NPAR* 1, d (2007), 103–110.

Alexandrina Orzan, Adrien Bousseau, Holger Winnemöller, Pascal Barla, Joëlle Thollot, and David Salesin. 2009. Diffusion Curves: A Vector Representation for Smooth-Shaded Images. *ACM Transactions on Graphics (TOG)* 28, 5 (2009).

Giuseppe Papari, Nicolai Petkov, and Patrizio Campisi. 2007. Artistic Edge and Corner Enhancing Smoothing. *IEEE Transactions On Image Processing* 16, 10 (2007).

Sylvain Paris, Samuel W. Hasinoff, and Jan Kautz. 2011. Local Laplacian Filters: Edge-aware Image Processing with a Laplacian Pyramid. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 68, 12 pages.

Sylvain Paris, P. Kornprobst, J. Tumblin, and F. Durand. 2008. Bilateral Filtering: Theory and Applications. *Foundations and Trends in Computer Graphics and Vision* 4, 1 (2008), 1–75.

Patrick Pérez, Michel Gangnet, and Andrew Blake. 2003. Poisson Image Editing. *ACM Transactions on Graphics* 22, 3 (2003), 313.

E D Pisano, S Zong, B M Hemminger, M DeLuca, R E Johnston, K Muller, M P Braeuning, and S M Pizer. 1998. Contrast Limited Adaptive Histogram Equalization Image Processing to Improve the Detection of Simulated Spiculations in Dense Mammograms. *Journal of digital imaging* 11, 4 (1998), 193–200.

Pizer S Puff D, Pisano E, Muller K, Johnston R, Hemminger B, Burbeck C, McLelland R. 1994. A method for determination of optimal image enhancement for the detection of mammographic abnormalities. *Journal of Digital Imaging* 7, 4 (1994), 161–171.

Amir Semmo, Daniel Limberger, Jan Eric Kyprianidis, and Jürgen Döllner. 2016. Image stylization by interactive oil paint filtering. *Computers and Graphics (Pergamon)* 55 (April 2016), 157–171.

Jean Serra. 1998. Connectivity on Complete Lattices. *Journal of Mathematical Imaging and Vision* 9, 3 (1998), 231– 251.

Carlo Tomasi and R. Manduchi. 1998. Bilateral Filtering for Gray and Color Images. *International Conference on Computer Vision* (1998), 839–846.

Shuhang Wang, Jin Zheng, Hai Miao Hu, and Bo Li. 2013. Naturalness Preserved Enhancement Algorithm for Non-Uniform Illumination Images. *IEEE Transactions on Image Processing* 22, 9 (2013), 3538–3548.

Holger Winnemöller, Sven C. Olsen, and Bruce Gooch. 2006. Real-time video abstraction. *ACM Transactions on Graphics* 25, 3 (2006), 12–21.

Hongteng Xu, Guangtao Zhai, Xiaolin Wu, and Xiaokang Yang. 2014. Generalized Equalization Model for Image Enhancement. *IEEE Transactions on Multimedia* 16, 1 (2014), 68–82.

Lu Yuan and Jian Sun. 2012. *Automatic Exposure Correction of Consumer Photographs.* Springer Berlin Heidelberg, Berlin, Heidelberg, 771–785.