

Independent Color-Channel Adjustment for Seamless Cloning based on Laplacian-Membrane Modulation

Bernardo Henz¹, Frederico A. Limberger², Manuel M. Oliveira³

Instituto de Informática – UFRGS – Porto Alegre, Brazil

Abstract

Image compositing combines visual elements from multiple images aiming to produce natural-looking results. Poisson-image-editing methods can generate seamless compositions, but introduces color changes in the inserted elements, which often generate unrealistic and displeasing results. Recently, Henz et al. have presented an approach for controlling the amount of color changes in the inserted image patches [1]. We improve their work by using new sets of color basis vectors tailored to maximize color separation on transition zones, and by independently controlling the amount of adjustment on individual color channels. With our technique, artists have greater control over the final composites color, besides significantly enhancing their freedom to achieve effects that were unreachable by previous methods. We performed two user studies to evaluate the perceptual effectiveness of the proposed technique. The studies show that our method provides the flexibility to produce even more natural and pleasing composites than previous seamless-cloning techniques.

Keywords: Color adjustment, Color preservation, Poisson image editing, Seamless cloning, Image composition.

1. Introduction

Image compositing tries to combine multiple images in a natural-looking way. Poisson image editing (PIE) [2] can produce seamless compositions and has been adopted by several recent techniques [3, 4, 5, 6]. However, despite its elegant formulation, PIE may introduce significant color changes in the inserted elements [7], leading to some displeasing results (Fig. 1 (center)).

Recently, Henz et al. have proposed a color-adjustment technique for controlling the amount of color change in PIE [1]. Such a technique works by modulating a Laplacian membrane based on an alpha mask and on some interpolation parameter w . While capable of achieving more natural-looking compositions than previous seamless-cloning approaches, the range of possible results is still limited by the use of a single w value applied to all color channels.

In this paper, we improve the color-adjustment method of Henz et al. [1] by using new sets of basis vectors to encode the color space, and by independently adjusting the influence of the individual color channels in the inserted elements. This new strategy significantly enhances the freedom of the users to produce composites that were unreachable by the original technique. While our approach can be used with any color space and any set of basis vectors that span such a space, we show that the use of particular bases offer some advantages. For instance, an orthonormal base obtained by aligning one of its basis vectors with the (approximate) background color (in the transition

zone) of the target image maximizes color separation, allowing for great control over how much the background color actually affects the inserted objects. Fig. 1 illustrates the effectiveness of our technique. Note how Poisson cloning (*i.e.*, PIE) severely shifts the colors of the inserted characters (Fig. 1 (center)). The result obtained with our technique, on the other hand, naturally preserves the character's original colors (Fig. 1 (right)).

To validate our technique, we have performed two user studies to compare our results with ones obtained with previous seamless-cloning approaches, including [1]. The results of these studies indicate that, on average, our method generates more pleasing composites, which are preferred by the users. The flexibility of our technique allows it to benefit from fast solutions for computing Laplacian membranes [8, 9], guaranteeing high-quality and quick results, even when compositing large patches containing over a million pixels.

The **contributions** of this work include:

- An efficient technique to perform seamless cloning that provides user control over the degree of color preservation in the inserted patches for different color channels, and in arbitrary color spaces (Section 3);
- A technique for automatically obtaining an orthonormal basis for the selected color space that maximizes color separation for the transition zone in the target image. Such a basis significantly improves the user control over the amount of target background color that diffuses into the inserted patches (Section 3.1). Combined with the previous contribution, it produces more pleasing and natural composites.

¹bhenz@inf.ufrgs.br

²falimberger@inf.ufrgs.br

³oliveira@inf.ufrgs.br



Figure 1: Compositing elements from a source image into a target one. (Left) Source and target images. (Center) Compositing generated by Poisson cloning [2]. (Right) Compositing created with our method by modulating a Laplacian membrane using an automatically computed basis for the RGB color space and interactively-chosen adjustment parameters for the new color channels $((w_1, w_2, w_3) = (0.74, 0.91, 0.90))$. Note how our method preserves the original colors while ensuring smooth transitions between the foreground and the background elements.

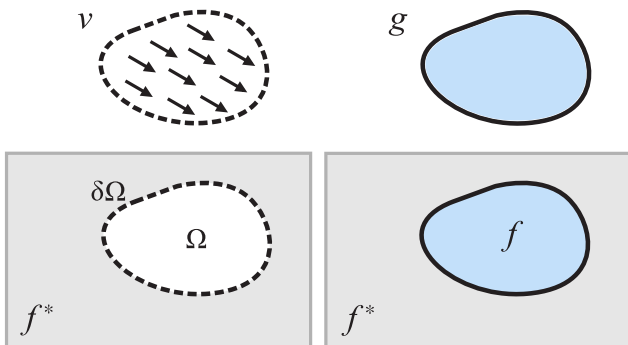


Figure 2: Poisson cloning [2] places a source patch g , with an associated vector field v and support Ω , over a target background f^* . To produce a seamless result, f^* is composited with f (a modified version of g computed by the algorithm).

2. Background and Related Work

Let g be a source image patch to be composited on some target image f^* (Fig. 2). The PIE framework [2] computes a new image patch f that seamlessly integrates itself into f^* by solving the Poisson equation

$$\Delta f = \nabla \cdot v, \quad (1)$$

with Dirichlet boundary conditions $f|_{\delta\Omega} = f^*$. Here, $\delta\Omega$ is the boundary of the region Ω (the support of both g and f). In general, $v = \nabla g$, enforcing that the gradient of f be as close as possible to the gradient of g . We refer to the process of com-

puting and pasting f into the background image f^* as *Poisson cloning* (Fig. 2).

2.1. Optimization Methods

The solution of the Poisson equation is obtained by solving an $N \times N$ sparse linear system, where N is the number of pixels in the region Ω . For large values of N , finding the solution of such a system becomes computationally and memory intensive. Trying to address this problem, some techniques have been proposed to efficiently solve large sparse linear systems [10, 11], while others try to reduce the size of the region Ω [12]. Although they alleviate the problem, they are still not sufficiently fast for use in interactive applications.

Farbman et al. [8] describe an efficient method to compute a Laplacian membrane \tilde{f} without solving a linear system. This is obtained by approximating the solution of the following Laplacian equation:

$$\Delta \tilde{f} = 0, \quad \tilde{f}|_{\delta\Omega} = f^* - g \quad (2)$$

with Dirichlet boundary conditions. f is then computed as

$$f = \tilde{f} + g. \quad (3)$$

Eq. 2 calculates a smooth membrane that interpolates the difference $f^* - g$, computed at $\delta\Omega$, over the entire region Ω (Fig. 3). In a subsequent work, Farbman et al. [9] use a multiscale convolution scheme to approximate the membrane \tilde{f} . The main limitation of these methods is that Eqs. 1 and 2 are only equivalent if the vector field v is conservative (*i.e.*, the line integral between any two points in the vector field is independent of the

chosen path). Thus, v must be the gradient of some scalar function and, therefore, these methods cannot be used to perform Poisson cloning using mixed gradients [2], for instance. Nevertheless, these optimizations are very handy when one wants to perform seamless cloning of large patches, on several images, or in video sequences.

2.2. Color-Correction Techniques

Since PIE works in the gradient domain, it has little control over changes in color space. Thus, if the tones of the inserted patch and the background image differ significantly, the colors of the inserted elements may be severely affected (Fig. 1 (center)). Many methods have been developed specifically to control this color problem. Dizdaroglu and Ikibas use information from each color channel to build a *mixed guided vector field* that tries to preserve the foreground color [13]. Li et al. [14] also build a new guided vector field as a weighted average between the gradient of the source image and a matting composition. Yang et al. propose a multi-resolution framework that uses a variational model in a minimization equation to generate more natural and realistic compositions [15]. Their method also modifies the guidance vector field. Guo and Sim modify the Poisson equation by adding different weights for groups of image pixels [16]. The user specifies regions that should be preserved (using markups), and the technique adjusts their weights accordingly. Wu and Xu use geodesic-distance maps for computing a color-belief measure for each pixel in the source image [17]. This measure is used as a weight in a closed-form energy function that is minimized by solving a linear system.

While all these methods achieve color preservation by solving new and modified minimization equations, or by building new guidance vector fields, the resulting systems cannot be written as Laplace equations because there is no assurance that the new vector fields will be conservative. Therefore, they cannot take advantage of fast methods that approximate the Laplace membrane [8, 9]. Our technique performs color preservation while taking advantage of these fast methods.

Some color-transfer techniques focus on matching the mean and variance values of color histograms of the source and target images [18, 19]. Lalonde and Efros proposed a method to re-color cloned image regions [20]. Recently, Xue et al. performed statistical and visual perception experiments to determine factors that influence realism on image composites, proposing an automatic method to try to improve them [21]. Their method uses machine learning for choosing zones in the histograms of the source and target images for matching luminance, CCT (Correlated Color Temperature), and saturation over the images. While these methods focus on achieving realistic composites, there is no assurance of color preservation, i.e., object colors may significantly change. In addition, they do not offer control over the color adjustments.

3. The Color-Adjustment Method

According to Eq. 3, the Laplace membrane \tilde{f} is the difference between the computed patch f and the original one g .

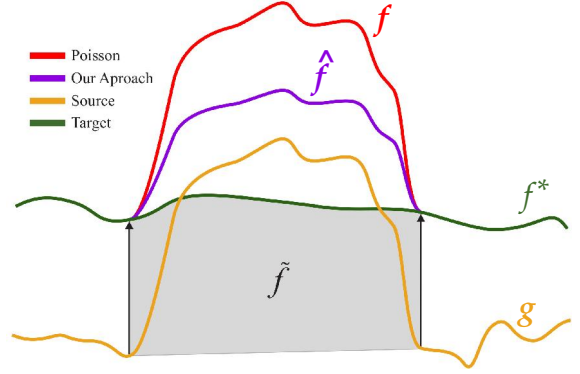


Figure 3: Poisson cloning computes $f = \tilde{f} + g$. Thus g (orange curve) is shifted by \tilde{f} to f (red curve). Our technique modulates the contribution of \tilde{f} , allowing the user to obtain results that range from alpha blending (orange curve) to Poisson cloning (red curve). \hat{f} (purple curve) shows a typical intermediate result produced by our technique. Note how its borders nicely blend with the target image f^* .

Therefore, the bigger the values over the membrane, the bigger the difference between f and g . The central idea of Henz et al.'s method is to *control this difference by adjusting the membrane accordingly*. Figs. 3 and 4 illustrate the concept.

Similar to other approaches for color preservation in seamless cloning [7, 14], Henz et al.'s technique [1] uses two masks: a pasting mask M_p (same as the mask used for PIE, and used for computing the membrane \tilde{f} in Eq. (2)), and an alpha mask M_α for the object(s) one would like to preserve colors. M_p includes the object plus some *transition zone*, and can be defined manually, using some interactive tools [22, 23], or obtained as a dilated version of M_α .

The method then modulates the Laplacian membrane using

$$\tilde{f}_m = (1 - wM_\alpha)\tilde{f}, \quad (4)$$

where \tilde{f} is the Laplace membrane computed using [9], \tilde{f}_m is the modulated membrane, and $w \in [0, 1]$ controls the modulation. Fig. 4 shows the membrane before and after the modulation (Eq. 4). The compositing is obtained as

$$\hat{f} = \tilde{f}_m + g. \quad (5)$$

Note the influence of the parameter w in Eq. 4. When $w = 0$, the equation reduces to Poisson image editing, with no color adjustment (i.e., $\tilde{f}_m = \tilde{f}$). When $w = 1$, Eqs. 4 and 5 produce an alpha blending of the patch g and the Poisson result \tilde{f} . $w \in [0, 1]$ controls how much of the color will be changed. Bigger values will make the colors on the inserted object(s) to take values similar to the original ones; lower values will produce results closer to PIE [2].

Henz et al. [1] provide an automatic way to compute w . Note that, at the transition zones, if the colors in the target image differ too much from the colors in the source image, Poisson image editing will diffuse such difference into f . Their idea is to use a measure of this difference to compute a suitable w value. As compositing results are strictly related to human perception, they use Pele and Werman's perceptual color-difference measure [24] for this purpose.

A transition zone Γ (Fig. 7 bottom left), whose colors are computed by PIE, is defined as $\Gamma = M_p - M'_\alpha$, where M'_α is a binary mask obtained as

$$M'_\alpha(x, y) = \begin{cases} 1, & \text{if } M_\alpha(x, y) > 0 \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Henz et al. evaluate the perceptual difference between the target and source images. Let c_T be the average color over Γ in the target image; likewise, let c_S be the average color over Γ in the source image. Henz et al. compute $\text{diff}_{pw} = |c_T - c_S|_{pw}$, where $|\cdot|_{pw}$ is the Pele and Werman's measure. Since diff_{pw} falls in the $[0, 1]$ interval, with lower values indicating similar colors, they set $w = \text{diff}_{pw}$.

While w , computed as described, produces good results in practice, the use of a single parameter value to modulate all color channels limits the range of compositing results that can be achieved. Next, we show how one can enhance the control over color preservation/adjustment, allowing users to explore a much wider range of compositing possibilities.

3.1. Improved Color Bases for Seamless Cloning

This Section introduces two extensions to improve the quality of the results obtained with the technique presented by Henz et al. [1]. The first one is an automatic technique to compute an orthonormal basis (for the color space) more suitable for color preservation/adjustment in seamless cloning applications. The second extension consists of independent modulation of the individual color channels. Together, these two extensions allow users to achieve results that were not possible with the original technique.

Our goal is to improve the control of the amount of color diffusion from the transition zone Γ (of the target image) into the inserted patch. Thus, let the *dominant color* (dc) be the average color over Γ in the target image. For the purpose of this description, we will assume the use of the RGB color space, but any other color space would be equally fine. Thus, let R , G , and B be three orthonormal vectors defining the canonical basis of the RGB color space. After normalizing the dc vector, we find a 3×3 transformation matrix T that maps R into dc . T rotates R around the vector $U = R \times dc$ by an angle $\alpha = \arccos(R \cdot dc)$ (Fig. 5 (left)), where the symbols \times and \cdot are the cross and dot product, respectively. The vectors of the new basis \mathbf{B} can be computed as:

$$[\mathbf{B}_1 \ \mathbf{B}_2 \ \mathbf{B}_3] = T [R \ G \ B],$$

where \mathbf{B}_i , R , G , and B are column vectors. Since $[RGB] = I_{3 \times 3}$ (i.e., the canonical basis for the RGB color space equals the identity matrix):

$$[\mathbf{B}_1 \ \mathbf{B}_2 \ \mathbf{B}_3] = T. \quad (7)$$

The actual entries of matrix T are provided in Appendix A (Eq. (A.1)). Basis \mathbf{B} is illustrated in Fig. 5 (right). Recall that \mathbf{B}_1 corresponds to the dominant color in the transition zone Γ . Therefore, the fundamental idea behind our method can be summarized as follows: *The basis \mathbf{B} allows us to control the amount*




Figure 4: The Laplacian membranes. (top) \tilde{f} membrane computed using [9]. (bottom) \tilde{f}_m membrane generated by our color-adjustment approach. The lower values of the membrane \tilde{f}_m reduce the difference between \tilde{f} and g .




Figure 5: Color basis transformation. (left) The RGB color basis, the *dominant color* vector, and the α angle between R channel and dc . (right) The new orthonormal basis \mathbf{B} , whose vector $\mathbf{B}_1 = dc$.

of background color diffusion into the inserted patch simply by controlling the influence of the coefficient associated to \mathbf{B}_1 in the resulting Laplacian membrane (Eq. (9)).

For any given color $c = (c_R, c_G, c_B)$ expressed in the R, G, B basis, its corresponding coefficients (b_1, b_2, b_3) expressed in the new color basis \mathbf{B} are immediately obtained through the change-of-basis matrix:

$$\begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_{1.1} & \mathbf{B}_{1.2} & \mathbf{B}_{1.3} \\ \mathbf{B}_{2.1} & \mathbf{B}_{2.2} & \mathbf{B}_{2.3} \\ \mathbf{B}_{3.1} & \mathbf{B}_{3.2} & \mathbf{B}_{3.3} \end{bmatrix} \begin{bmatrix} c_R \\ c_G \\ c_B \end{bmatrix}, \quad (8)$$

where $\mathbf{B}_{i,j}$ is the j -th component of the new basis vector \mathbf{B}_i .

The Laplacian membrane \tilde{f} can be directly computed using the new color basis \mathbf{B} . We refer to the Laplacian membrane expressed in the new color basis as $\tilde{f}_{\mathbf{B}}$. One can then modulate this Laplacian membrane using different parameter values $w_m \in [0, 1]$ for each channel \mathbf{B}_i , obtaining $\tilde{f}_{\mathbf{B}m}$ (the multi-channel analogous to Eq. (4) in the new color basis):

$$\tilde{f}_{\mathbf{B}m} = (1 - w_m M_\alpha) \tilde{f}_{\mathbf{B}}. \quad (9)$$

Eq. (9) provides great flexibility for controlling the amount of background color diffusion into the inserted patch, allowing the creation of results that are unreachable by previous seamless-cloning techniques.

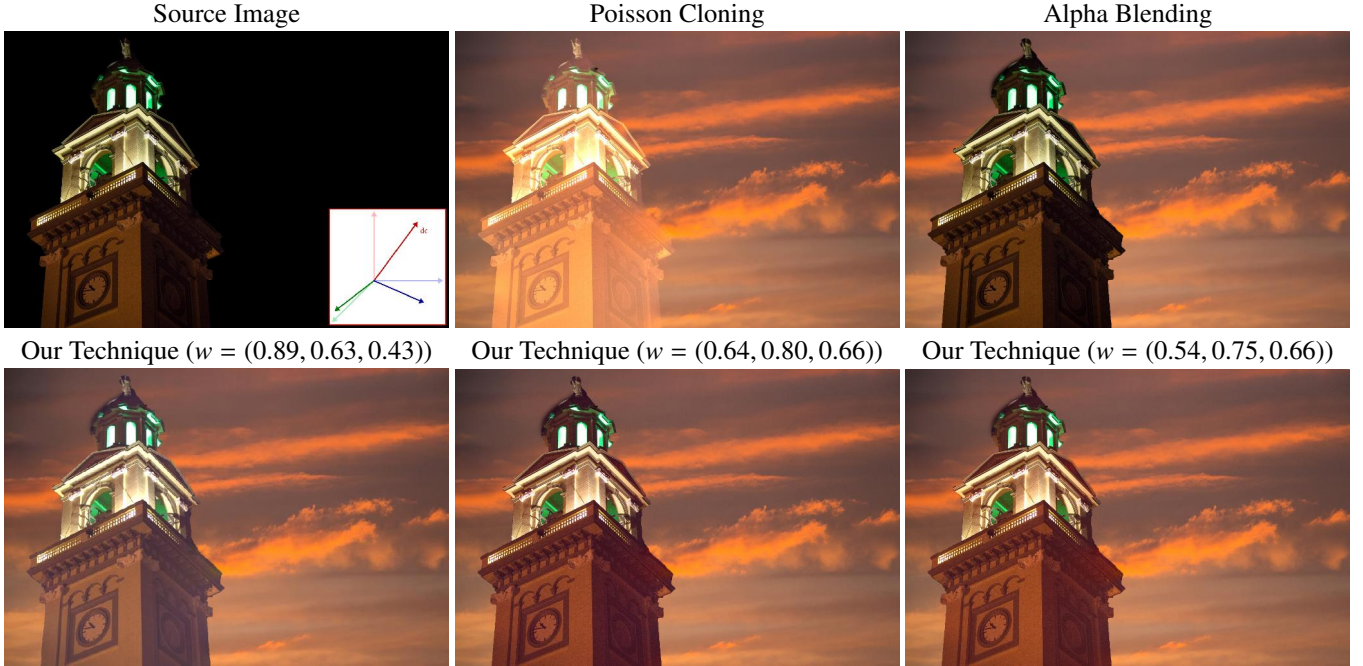


Figure 6: Compositing a bell tower from a photograph taken at night onto a sunset background. Comparison of the results produced by conventional Poisson cloning, alpha blending, and our technique using distinct sets of values for the parameter w . Poisson cloning introduces significant color shift in the bell tower, producing some unnatural result. Although alpha blend preserves the original colors of the tower, the result also does not look natural, as there is no interaction of the ambient light with the tower. By using independent w_m values for each color channel in the new basis \mathbf{B} (see inset in the source image), our technique provides additional degrees of freedom for achieving different results.




Figure 7: Seamless cloning of a clownfish. (left) From top to bottom: source image, alpha mask, and transition zone $\Gamma = M_p - M'_\alpha$, respectively. (right) Results produced by PIE (top) and by our technique using $w = (0.65, 0.90, 0.90)$ (bottom). The inset shows the new color basis \mathbf{B} .

4. Results and Discussion

As a direct extension of the method of Henz et al. [1], our technique can be easily combined with the techniques proposed by Farbman et al. [8, 9] to compute an approximation to \tilde{f} . For all the results shown in this paper, we used the Convolution Pyramids technique [9] to compute the Laplacian membranes. According to Farbman et al., the computation cost for obtaining the new membrane \tilde{f}_m is (almost) linear in the number of pixels. In addition to the computation of \tilde{f}_m , the cost of our approach is linear in the number of pixels in the region Ω . Thus, the total

cost of the compositing process is bounded by the computation of \tilde{f} . We now present results obtained using our enhanced color basis and independent values for w_1 , w_2 , and w_3 . In these examples, the alpha masks were computed by the available MATLAB implementation of Shahrian et al.'s technique [25]. The corresponding pasting masks M_p were obtained from the alpha masks M_α , after applying binarization and dilation filters.

The example in Fig. 6 shows the compositing of a bell tower (top left) from a picture taken at night, into a sunset background. Note how Poisson cloning yields a significant color shift to the tower (top center). The use of alpha blending (top right) preserves its original colors, but produces some unnatural compositing. Our results (bottom row) range between these two extremes, achieving good color-preservation and smooth transitions. For these images, the values of the corresponding w_m triplets were manually chosen to exhibit variations among the results.

Fig. 7 compares the results of Poisson cloning (top) and our method (bottom). Note how our result gracefully preserves the colors of the original fish in the new background.

Fig. 8 shows the compositing of a portrait into a background image with unusual ambient illumination. Fully preserving the original skin tones (top right) makes the person to stand out from the rest of the image, as if the subject were not originally part of the scene. Strongly changing the skin color (bottom left) tends to produce unpleasant results. Our method allows for smooth diffusion of the green color as ambient light over the subject. Fig. 9 provides another example comparing our technique to Poisson cloning and to alpha blending. It shows




Figure 8: Compositing a portrait on a background with unusual ambient illumination. Pure alpha blending or Poisson cloning produces less realistic results. Our technique allows for smooth diffusion of the green color as ambient light over the subject. The inset shows the new basis \mathbf{B} .



Figure 9: Image composition of two planes into a sunset background. (first column) Input: source image, alpha mask, and pasting mask, respectively. The remaining columns show compositing results obtained using Poisson cloning, our method ($w = (0.55, 0.63, 0.64)$), and alpha blending. Our method diffuses the background color over the inserted patches, simulating the effect of ambient illumination. Note the colors of the planes and of the steam. The inset shows the new basis \mathbf{B} .

two planes inserted on a sunset sky. By diffusing sunset’s color over the steam while controlling the colors of the planes, our technique produces more natural-looking results.

Our technique requires the specification of an alpha mask. This, however, is not a big constraint, as there are many high-quality alpha-matting techniques available [25, 26], including one that works in real time [27]. Although refined alpha masks are preferred, our technique still achieves good results with rough M_α masks. For instance, the masks in Fig. 6 were produced manually using a quick-selection tool. In fact, we show that even when using very precise alpha masks, our technique tends to produce more naturally-looking results than pure alpha blending. This is illustrated in Figs.8 and 9, where the shading of the objects inserted using pure alpha blending is not "influenced" by the lighting of the target environment.

Our color-adjustment method works over a Laplacian membrane and can handle both conservative and non-conservative

vector fields. In the first case, the membrane is computed using fast seamless-cloning approaches [8, 9]. If the vector field is non-conservative, we compute the membrane as $\tilde{f} = f - g$, where f is obtained solving a linear system in the conventional way. Fig. 10 shows our color-preservation method applied with a variation of the mixed-gradient approach [2]. In this example, we have used mixed gradients only in the transition zone (region in the pasting mask not covered by the alpha mask). This solution avoids blurring the transition zone, while also preserving the gradient of the bear fur. Such result cannot make use of the fast membrane-approximation approaches [8, 9].

Fig. 11 shows an example where a patch from the source image has been inserted twice on a target background whose transition zones contain distinct textures with colors significantly different from the source ones. Note how the resulting composite looks pleasing. Likewise, our technique is not limited to target image regions containing only variations of a single

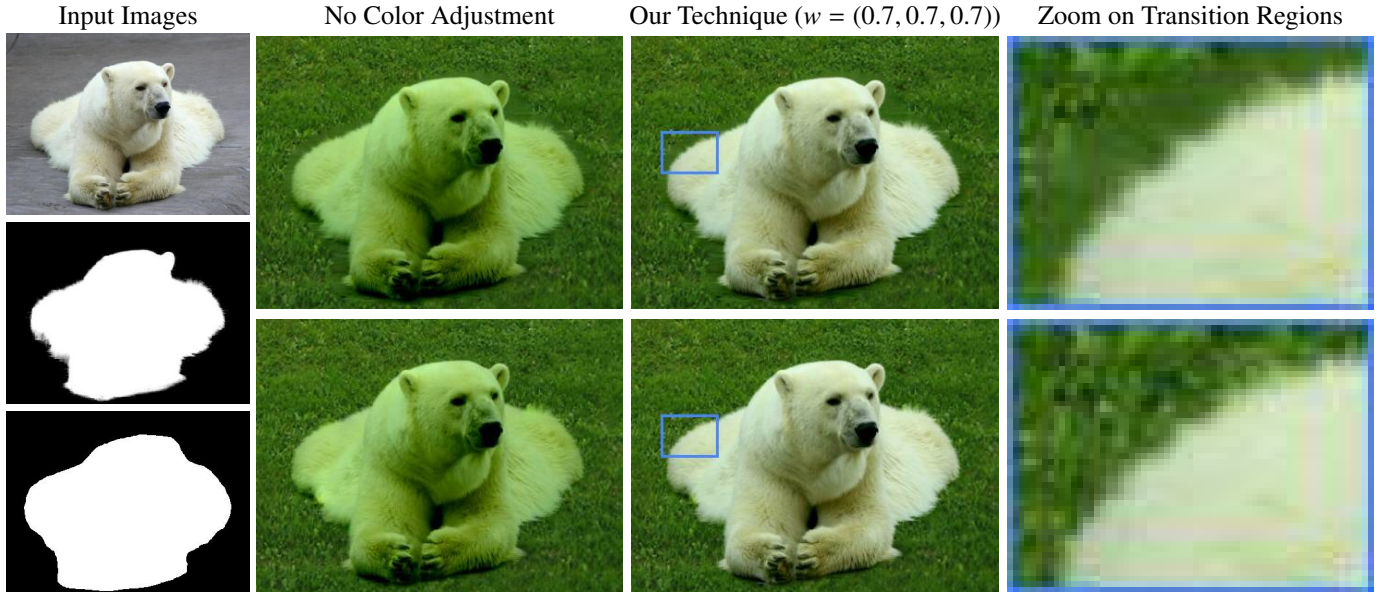


Figure 10: Color adjustments with the use of different gradient strategies at the transition zone, obtained with our technique. Single gradient (top row) and mixed gradients (bottom row). Input images are shown on the left: source image, alpha mask, and pasting mask, from top to bottom, respectively. Note how the mixed gradients are able to match the background texture in the transition regions.




Figure 11: Using our method to composite into a target image with a challenging texture. On the right one sees our result using $w = (0.66, 0.80, 0.81)$. Note that the transition zones contain distinct textures with colors significantly different from each other.

color. This is illustrated in Fig. 12, which shows a background containing multiple colors. The result produced by PIE for this example exhibits strong artifacts due to diffusion of the background colors on the pencil color (Fig. 12 (center)). Our result, on the other hand, properly preserves the pencil color (Fig. 12 (right)). This is achieved through a proper selection of w_m values. For both techniques, the transition zone was defined as a one-pixel wide band around the pencil.

The user should be careful to not introduce textures discontinuities at the boundaries of Ω . Although mixed gradients can be used for treating this problem in some cases, there is no assurance of obtaining good results in general. An alternative for these cases is to use our method with techniques that treat texture discontinuities [4, 28, 29, 17].

We tried to compare the quality of our results with other color-adjustment techniques. Unfortunately, we could not find any available implementation for these techniques. The only method with figures of sufficient resolution for comparison was Guo and Sim's [16]. Fig. 13 compares one of their results with

ours. Note how the colors in the transition zone (window internal wall) in their result look desaturated.

4.1. Performance Evaluation

Table 1 compares the execution times (in seconds) for our technique against an implementation of the bi-conjugate gradient method of the PIE. Both implementations are written in MATLAB. The total cost of our technique is given by the approximation of the Laplacian membrane, change of basis, and membrane modulation. These measurements were performed on an i7 3.5GHz CPU with 16 GB of RAM. As the size of the cloning area increases, the use of PIE becomes impractical.

4.2. Qualitative Evaluation

We have performed two experiments to evaluate the quality of our composites. In the first one, we compared our results against alpha blending and Poisson cloning. The second experiment compared our approach (the use of a new basis for the




Figure 12: Comparison of Poisson cloning and our technique when compositing against a background multiple distinct colors. (left) Source and target input images. (center) The result of PIE (Poisson cloning). Note the artifacts due to interference of the background colors on the pencil color. (right) Result produced by our technique using $w = (0.76, 0.84, 0.84)$. Note how the pencil color is properly preserved. For both techniques, the transition zone Γ was defined as a one-pixel wide band around the pencil.




Figure 13: Comparison between Guo and Sim's method [16] and ours. (left) Source and target images, (center) their composition, (right) our result using $w = (0.78, 0.96, 0.69)$. The inset shows the new color basis \mathbf{B} .

color space and the use of independent w_m parameters) against results produced by the original technique by Henz et al. [1].

For the first experiment, 119 subjects, ages ranging from 14 to 49 years old, rated a set of 14 groups of composites (shown in random order). For each group, subjects were shown the results created using alpha blending, Poisson cloning, and our technique (the sequence of results on the screen was defined randomly). Each subject gave to each composite a score ranging from 1.0 (not convincing) to 5.0 (very convincing). The average score of the composites generated by alpha blending was 2.48, with standard deviation of 1.3. For the ones created using Poisson cloning the average score was 2.39, with standard deviation of 1.38. The composites created with our technique achieved an average score of 3.15, with standard deviation of 1.17. Additionally, our technique received the best score on 69% of the cases, showing that it achieves better composites than traditional techniques. Fig. 15 shows a histogram of the votes, while Fig. 14 shows examples of images created with our technique and used in these studies.

Table 1: Performance Comparison. Cloning area (in pixels) and time (in seconds) for our method and PIE. The total time for our method includes the computation of the Laplacian membrane (using Convolution Pyramids [9]), changing the color basis and modulating the membrane. The PIE time is from a MATLAB implementation using the bi-conjugate gradient method for solving the linear system.

Source image		Time (in seconds)	
Example	Clone area	Our method	PIE
Fish	93 K	1.11	43.32
Plane	162 K	1.56	122.35
Portrait	359 K	1.48	552.39
Church	765 K	1.75	2,473.40
Dolls	1,983 K	6.29	17,946.55

For the second experiment, we asked subjects to choose the best between two composites: one created using the method proposed by Henz et al. [1], and other obtained using our extension of the original technique. In total, subjects performed 824 comparisons. The results generated by our method were preferred in 55% of the cases, with a p -value of 2.68×10^{-81} , meaning that the result of this test has a high statistical significance. Table 2 summarizes these choices for each image presented in the test.




Figure 15: Score histograms for the first qualitative evaluation. The bins represent the scores (higher values are better) for each image-composing technique. The majority of the scores for our method are greater than or equal to 3, whereas the majority of the scores for other techniques are smaller than or equal to 3.

Our project website [30] provides a MATLAB implemen-



Figure 14: Examples of image compositions created with our technique on the user studies. On the second study, they were compared against the ones achieved by Henz et al.’s method, where they were preferred by the subjects in the following proportions: 65% (bell tower), 45% (house), 54% (dolls), 34% (girl), 77% (portrait), 52% (eye), 57% (planes), and 64% (fish). On average, our results were preferred 55% of the time.

Table 2: Results on the second qualitative experiment performed by 103 subjects. In 6 of 8 examples our technique was preferred over the method of Henz et al. [1]. Figure 14 shows the images used in this evaluation.

Example	Preference
Portrait	77%
Bell Tower	65%
Fish	64%
Planes	57%
Dolls	54%
Eye	52%
House	45%
Girl	34%
Overall	55%

tation of our method, as well the images used in the paper and in the experiments. We encourage the readers to try their own compositions by experimenting with different parameter setups.

5. Conclusion

We presented a generalization of Henz et al.’s technique [1] to perform seamless image compositing. Our generalization provides more creative freedom for controlling the color adjustment in the inserted patches. Our technique consists of using a new basis for the color space. In such a basis, one of its vectors approximates the background colors of the transition zone in the target image. Such colors are the ones that tend to be diffused by Poisson cloning into the inserted patches. Thus, the new basis provides a simple and natural way of avoiding undesired color diffusion. This is achieved by performing composition with independent color adjustments for each component of the new basis. The independent w_m parameter values can be set interactively, or the user can opt to use a single w

value, as in the original technique described by Henz et al. [1]. Our technique takes advantage of fast solutions for estimating Laplacian membranes, allowing it to provide instant feedback on the compositing results.

We have demonstrated the effectiveness of our technique by performing two user studies that compared the results of our technique against the ones obtained with traditional Poisson cloning [2], alpha blending, and Henz et al.’s [1] original formulation. The results of these studies indicate that our technique produces more pleasing composites than the compared ones.

Given the efficiency and freedom provided by our method, we believe it can become a useful tool for photographers and casual users interested in performing challenging composites involving high-resolution images.

As future work, it would be desirable to have an image-statistics-based method to automatically compute per-channel parameter values that can generate pleasing composites.

Acknowledgements

This work was sponsored by CAPES and CNPq-Brazil (grants 306196/2014-0, 134134/2013-3, 482271/2012-4, and 11892-13-7). Fig. 1 (left) is from <http://www.alphamatting.com/>. Figs 13 (left) and (center) are reproduced from [16] (Springer License 3754870262618). The remaining images are from Flickr users bernardohenz, Don McCullough, Bud, slworking, salemstudioz, simens, Wonderlane, didi8, cameliatwu, and Andrew Taylor; and from Pixabay user HebiFot. All the images are under Creative-Commons license.

$$T = \begin{bmatrix} U_1U_1(1 - \cos \alpha) + \cos \alpha & U_1U_2(1 - \cos \alpha) - U_3 \sin \alpha & U_1U_2(1 - \cos \alpha) + U_2 \sin \alpha \\ U_2U_1(1 - \cos \alpha) + U_3 \sin \alpha & U_2U_2(1 - \cos \alpha) + \cos \alpha & U_2U_3(1 - \cos \alpha) - U_1 \sin \alpha \\ U_1U_3(1 - \cos \alpha) - U_2 \sin \alpha & U_3U_2(1 - \cos \alpha) + U_1 \sin \alpha & U_3U_3(1 - \cos \alpha) + \cos \alpha \end{bmatrix} \quad (\text{A.1})$$

Appendix A. The Rotation Matrix T

Eq. (A.1) shows the entries of matrix T used in Eq. (7). The term U_i is the i -th component of vector U , which is the rotation axis.

References

- [1] Henz, B., Limberger, F.A., Oliveira, M.M.. Color adjustment for seamless cloning based on Laplacian-membrane modulation. In: Proc of SIB-GRAPI. 2015, p. 196–202.
- [2] Pérez, P., Gangnet, M., Blake, A.. Poisson image editing. ACM TOG 2003;22(3):313–318.
- [3] Jeschke, S., Cline, D., Wonka, P. A gpu laplacian solver for diffusion curves and poisson image editing. ACM TOG 2009;28(5):116:1–116:8.
- [4] Sunkavalli, K., Johnson, M.K., Matusik, W., Pfister, H.. Multi-scale image harmonization. ACM TOG 2010;29:125:1–125:10.
- [5] Tao, M.W., Johnson, M.K., Paris, S.. Error-tolerant image compositing. In: Proc. ECCV. 2010, p. 31–44.
- [6] Farid, M., Mahmood, A.. An image composition algorithm for handling global visual effects. Multimedia Tools and Applications 2014;71(3):1699–1716.
- [7] Li, X., Kim, D.. A color adjustment method for automatic seamless image blending. In: Proc. of HCI 2007, p. 82–91.
- [8] Farbman, Z., Hoffer, G., Lipman, Y., Cohen-Or, D., Lischinski, D.. Coordinates for instant image cloning. ACM TOG 2009;28(3):67:1–67:9.
- [9] Farbman, Z., Fattal, R., Lischinski, D.. Convolution pyramids. ACM TOG 2011;30(6):175:1–175:8.
- [10] Bolz, J., Farmer, I., Grinspun, E., Schröder, P.. Sparse matrix solvers on the gpu: conjugate gradients and multigrid. In: ACM SIGGRAPH 2005 Courses. 2005..
- [11] Wei, F., Yang, H.. Directed transmission method, a fully asynchronous approach to solve sparse linear systems in parallel. In: Proc. SPAA. 2008, p. 365–365.
- [12] Szeliski, R., Uyttendaele, M., Steedly, D.. Fast poisson blending using multi-splines. In: Proc IEEE ICCP 11. 2011..
- [13] Dizdaroğlu, B., İkbaz, C.. A seamless image editing technique using color information. In: IPCV'09. 2009, p. 216–221.
- [14] Li, Y., Xie, Z., Sheng, B., Gui, Y., Ma, L.. Seamless image composition with coherent tone adjustment. In: Proc. VRCAI. 2011, p. 149–154.
- [15] Yang, W., Zheng, J., Cai, J., Rahardja, S., Chen, C.W.. Natural and seamless image composition with color control. Trans Img Proc 2009;18(11):2584–2592.
- [16] Guo, D., Sim, T.. Color me right - Seamless image compositing. In: Proc. CAIP. 2009..
- [17] Wu, H., Xu, D.. Image compositing using dominant patch transformations. Computers and Graphics 2014;38(0):277 – 285.
- [18] Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.. Color transfer between images. IEEE CGA 2001;21(5):34–41.
- [19] Pouli, T., Reinhard, E.. Progressive histogram reshaping for creative color transfer and tone reproduction. In: Proc of NPAR. 2010, p. 81–90.
- [20] Lalonde, J.F., Efros, A.A.. Using color compatibility for assessing image realism. In: IEEE ICCV. 2007..
- [21] Xue, S., Agarwala, A., Dorsey, J., Rushmeier, H.. Understanding and improving the realism of image composites. ACM TOG 2012;31(4):84:1–84:10.
- [22] Mortensen, E.N., Barrett, W.A.. Intelligent scissors for image composition. In: SIGGRAPH '95. ISBN 0-89791-701-4; 1995, p. 191–198.
- [23] Jia, J., Sun, J., Tang, C.K., Shum, H.Y.. Drag-and-drop pasting. ACM TOG 2006;25(3):631–637.
- [24] Pele, O., Werman, M.. Improving perceptual color difference using basic color terms. arXiv 2012;abs/1211.5556.
- [25] Shahrian, E., Rajan, D., Price, B., Cohen, S.. Improving image matting using comprehensive sampling sets. In: Proc. IEEE CVPR. 2013, p. 636–643.
- [26] Varnousfaderani, E., Rajan, D.. Weighted color and texture sample selection for image matting. IEEE TIP 2013;22(11):4260–4270.
- [27] Gastal, E.S.L., Oliveira, M.M.. Shared Sampling for Real-Time Alpha Matting. CGF 2010;29(2):575–584.
- [28] Sun, J., Zhu, H., Xu, Z., Han, C.. Poisson image fusion based on markov random field fusion model. Inf Fusion 2013;14(3):241–254.
- [29] Zhang, Y., Ling, J., Zhang, X., Xie, H.. Image copy-and-paste with optimized gradient. The Visual Computer 2014;30(10):1169–1178.
- [30] Henz, B., Limberger, F.A., Oliveira, M.M.. Independent color-channel adjustment for seamless cloning based on Laplacian-membrane modulation website. 2015. URL: http://inf.ufrgs.br/~bhenz/projects/independent_color_adjustment/.