

Error-Tolerant Image Compositing

Michael W. Tao¹, Micah K. Johnson², and Sylvain Paris³

¹ University of California, Berkeley

² Massachusetts Institute of Technology

³ Adobe Systems, Inc.

Abstract. Gradient-domain compositing is an essential tool in computer vision and its applications, e.g., seamless cloning, panorama stitching, shadow removal, scene completion and reshuffling. While easy to implement, these gradient-domain techniques often generate bleeding artifacts where the composited image regions do not match. One option is to modify the region boundary to minimize such mismatches. However, this option may not always be sufficient or applicable, e.g., the user or algorithm may not allow the selection to be altered. We propose a new approach to gradient-domain compositing that is robust to inaccuracies and prevents color bleeding without changing the boundary location. Our approach improves standard gradient-domain compositing in two ways. First, we define the boundary gradients such that the produced gradient field is nearly integrable. Second, we control the integration process to concentrate residuals where they are less conspicuous. We show that our approach can be formulated as a standard least-squares problem that can be solved with a sparse linear system akin to the classical Poisson equation. We demonstrate results on a variety of scenes. The visual quality and run-time complexity compares favorably to other approaches.

Keywords: gradient-domain compositing, visual masking.

1 Introduction

Gradient-domain compositing is an essential technique at the core of many computer vision applications such as seamless cloning [1–4], panorama stitching [5–7], inpainting [8], shadow removal [9], scene completion [10], and reshuffling [11]. These methods first delineate the composited regions, then compute a target gradient field and boundary conditions from these regions, and finally solve the Poisson equation to reconstruct an image. A major issue with gradient-domain compositing is that the combined gradient field may not be integrable; that is, an image with gradients that match the target field as well as the specified boundary conditions may not exist. Existing work mitigates this issue by moving the boundary to more carefully combine the merged regions. However, when the combined images are widely different, this strategy may not be sufficient. Or, if the user has specified the boundary by hand, he or she may not want it to be altered. For instance in Figure 1, the selection cannot be modified




Fig. 1. We present an image-compositing technique tolerant to selection inaccuracies. In this example, a user wishes to add trees to an image of the Egyptian pyramids, but it is not possible to select the trees without cutting through the foliage (a). Moreover, to ensure a good insertion behind the pyramids, it is not possible to modify the selection boundary. A direct copy of the pixels yields a undesirable visible seam (b). Standard gradient-domain compositing minimizes the seam, but leads to bleeding artifacts where the foliage is cut (c). Our method characterizes where color leakage should be avoided, producing a seamless composite without bleeding artifacts (d).

because the tree trunks have to abut the pyramids. Even with boundary refinement, the target gradient fields may be far from integrable, yielding color leaks and halos typical of Poisson-based methods.

In this paper, we present an algorithm for minimizing artifacts in gradient-domain image compositing. We characterize the origin of typical bleeding artifacts and analyze the image to locate the areas where they would be most and least conspicuous. Based on this analysis, we propose a two-step algorithm. First, we process the gradient values on the boundary to minimize artifacts in regions where bleeding would be visible. Second, we describe a weighted integration scheme that reconstructs the image from its gradient field so that residuals are located in textured regions where they are less visible. Our results show that the combination of these two steps yields significantly better composites. Moreover, our method is formulated as a least-squares optimization that can be solved using a sparse linear system, which makes our approach computationally efficient. We demonstrate our approach on scenarios in which boundary mismatches are likely to occur: user-driven seamless cloning [1], heterogeneous panorama stitching [7], and scene reshuffling [11].

1.1 Related Work

Gradient-domain techniques are useful to a variety of problems in computer vision, including image stitching, intrinsic images, shadow removal, and shape-from-shading [5, 12–15]. In most of these problems, the gradient field contains non-integrable regions and many authors have noted that reconstruction artifacts are often due to boundary conditions. As a result, a variety of methods have been introduced to minimize artifacts by refining the boundary location [2, 4, 5, 16]. Rather than moving the boundary, which may not always be possible, we focus on reconstructing the final image from the target gradient field once the boundary is specified. Our approach is complementary and orthogonal to boundary-refinement methods. We show that our image analysis combined with a careful study of the numerical scheme reduces visible

artifacts. Our approach could benefit many computer vision algorithms that rely on gradient-domain reconstruction as a subroutine.

The general formulation of the gradient-domain reconstruction problem is to seek an image I that approximates the target field \mathbf{v} in a least-squares sense (with ∇ , the gradient operator):

$$\operatorname{argmin}_I \int \|\nabla I - \mathbf{v}\|^2 \quad (1)$$

which can be minimized by solving the Poisson equation:

$$\Delta I - \operatorname{div}(\mathbf{v}) = 0 \quad (2)$$

where Δ is the Laplacian operator $\partial^2/\partial x^2 + \partial^2/\partial y^2$ and div is the divergence operator $\partial/\partial x + \partial/\partial y$. To solve this equation, one also needs boundary conditions that depend on the application. We illustrate how to compute the target gradient \mathbf{v} in the context of seamless compositing using three inputs: the background image, B ; the foreground image, F ; and a selection, \mathcal{S} with a boundary β [5].

$$\mathbf{v}(x, y) = \begin{cases} \nabla F & \text{if } (x, y) \in \mathcal{S}, (x, y) \notin \beta \\ \nabla B & \text{if } (x, y) \notin \mathcal{S} \\ \frac{1}{2}(\nabla F + \nabla B) & \text{if } (x, y) \in \beta \end{cases} \quad (3)$$

Other cases such as panorama stitching are similar except that the images are not named “foreground” and “background.” For the sake of simplicity, we will name the images foreground and background.

The gradients from the foreground image F and background image B are integrable since they are computed directly from images. But the gradients along the boundary between the two images may not be integrable, creating a source of errors that the integration routine must manage. Farbman et al. [17] address this issue by relying on users to identify the leaks. The gradients of marked regions are ignored, which removes the leaks. In comparison, our method analyzes the image to automatically adapt the integration process. Our approach shares similarities with the method of Lalonde et al. [16] who propose to take the image gradient magnitude into account during the reconstruction process. However, color leaks may still appear with this technique when boundaries are not accurate.

Besides image compositing, gradient-domain methods have also been used in computer vision for surface reconstruction problems, such as shape-from-shading and photometric stereo. In these problems, an algorithm estimates the gradient of a surface at every pixel and then a robust Poisson solver is used to find the surface that best fits the estimated gradients. We refer to the recent work of Agrawal et al. [14], Reddy et al. [18], and the references therein for detail. Although image compositing and robust integration techniques both reconstruct a 2D signal from its gradients, the two problems are fundamentally different. The gradients from surface-reconstruction methods are noisy everywhere, whereas image-compositing gradients are problematic only at the boundary between foreground and background. In this paper, we exploit this specificity to improve the quality of the results. We also rely on visual masking to locate integration residuals where they are less conspicuous.

1.2 Contributions

In this paper, we introduce several contributions.

▷ *Low-curl boundaries.* We describe a method that limits the artifacts by minimizing the curl of the target gradients on the foreground-background boundary.

▷ *Weighted Poisson equation.* We show how to add weights to the Poisson equation so that integration residuals lie in textured regions where they are less visible due to visual masking.

▷ *Efficient non-bleeding compositing.* We combine the two previous contributions to obtain a compositing algorithm that prevents bleeding artifacts while remaining linear akin to the original Poisson equation.

1.3 Overview

Our algorithm consists of two steps. First, we focus on the boundary between the foreground and background regions. We characterize the origin of the bleeding artifacts and we show how to modify the gradient field \mathbf{v} to minimize them. The second step focuses on the location of the integration residuals. We show that artifacts are less visible in textured regions due to visual masking. We describe an algorithm that controls the integration residuals such that they are located in textured areas. In the results section, we show that the combination of these two steps yields visually superior results.

2 Low-Curl Boundary

A necessary condition for a gradient field \mathbf{u} to be integrable is to have a zero curl¹. That is, if there exists an image I such that $\nabla I = \mathbf{u}$, then $\text{curl}(\mathbf{u}) = \partial \mathbf{u}_y / \partial x - \partial \mathbf{u}_x / \partial y = 0$. For example, consider the configuration illustrated in Figure 2(a). When all pixels come from one image, in this case the foreground image, the derivatives are consistent and the curl is zero. Therefore, this region is integrable, i.e., the image can be reconstructed from its gradients. The same observation holds for regions from the background image.

In the image compositing problem, a non-integrable gradient field only occurs on the boundary, as illustrated in Figure 2(b). On the boundary, the gradient field is a mixture of two fields and may have non-zero curl since gradients come from mixed sources. When the gradient field has a non-zero curl, we cannot minimize the Poisson equation (2) exactly and residuals remain. These residuals are often visible in composited images as halos, bleeding, or other artifacts.

2.1 Reducing the Curl on the Boundary

Since the non-integrability of regions along boundary is the source of artifacts, we seek to alter the desired gradient field to minimize the bleeding artifacts. Let \mathbf{v} represent the desired gradient field of the composited image. To preserve

¹ Note that for a 2D vector field $\mathbf{u} = (\mathbf{u}_x, \mathbf{u}_y)$, the curl is a scalar value that corresponds to the z component of the 3D curl applied to the 3D vector field $(\mathbf{u}_x, \mathbf{u}_y, 0)$.




Fig. 2. *Estimating the curl on a discrete grid.* Circles denote pixels and arrows denote differences between pixels. If the curl is computed within the foreground region (a), all the derivatives come from F and the curl is null. The background case is equivalent (not shown). On the boundary (b), derivatives from diverse sources are used and in general the curl is not zero.

the visual information of the two images, we do not modify the foreground or background gradients in \mathbf{v} . We only modify \mathbf{v} values on the boundary such that the curl is as small as possible.

A naive solution would be to seek $\text{curl}(\mathbf{v}) = 0$ everywhere. But the following counterexample shows that this approach would not achieve our goal. Consider the standard copy-and-paste operation that directly combines the pixel values and produces an image I_{seam} with visible seams. The curl of the gradient field of I_{seam} is null since it is computed from an actual image. And, inside the selection, gradients are equal to the foreground values since pixels have been copied. The same holds outside the selection with the background values. However, on the boundary, gradients are different from either the foreground or the background, which generates the seams. We address the shortcomings of this naive solution by seeking gradient values that minimize the curl and are close to the gradients of the input images.

A Least-squares Approach. We formulate our goal using a least-squares energy where the desired gradients \mathbf{v} on the boundary are the unknowns. The first term minimizes the curl: $\int_{\beta} [\text{curl}(\mathbf{v})]^2$ and the second term keeps the values close to the input gradients $\int_{\beta} (\mathbf{v} - \nabla F)^2 + \int_{\beta} (\mathbf{v} - \nabla B)^2$. This last term has the same effect as keeping \mathbf{v} close to the average gradient. We combine the two terms to obtain:

$$\operatorname{argmin}_{\mathbf{v}} \int_{\beta} \left([\text{curl}(\mathbf{v})]^2 + W_{\beta} \left[\mathbf{v} - \frac{1}{2}(\nabla B + \nabla F) \right]^2 \right) \quad (4)$$

where W_{β} controls the importance of the second term.

Adaptive Weights. Figure 3 shows results for several values of W_{β} . For large W_{β} , we only minimize the proximity to the input gradients, which is the standard gradient compositing with seamless boundaries but leaking artifacts. For a small W_{β} , we have the naive solution described above where we only minimize the curl. There are no bleeding artifacts but the boundary is visible. We combine these two behaviors by varying the weights according to the local image structure.




Fig. 3. *Influence of the curl term.* With high weights W_β , the composite is seamless but suffers from bleeding (a). With low W_β , bleeding disappears but seams become visible (b). Our adaptive approach locally adjusts the weights to achieve seamless results with no leaks (c).

Intuitively, a seamless boundary is desirable when both sides of the boundary are smooth. This is the case for instance when we stitch a sky region with another sky region. A seamless boundary is also acceptable when both sides are textured because leaking is a low-frequency phenomenon that will be hidden by visual masking. Figure 4 illustrates this effect that has also been used in the rendering literature [19–22]. In these two cases, we seek high values for W_β . But when a textured region is composited adjacent to a smooth region, we want to prevent bleeding because such regions would generate unpleasing artifacts on the smooth side, e.g. in the sky. In this case, we want low values of W_β . The following paragraph explains how we compute W_β based on the local amount of texture.

Estimating the Local Amount of Texture. Our strategy relies on the presence or absence of texture in a given neighborhood. In this paragraph, we describe a simple and computationally efficient texture estimator although one could use other models [23, 24]. Formally, our scheme is:

$$T_{\sigma_1, \sigma_2}(\mathbf{g}) = \frac{G_{\sigma_1} \otimes \|\mathbf{g}\|}{G_{\sigma_2} \otimes \|\mathbf{g}\|} n(\|\mathbf{g}\|) \quad (5)$$

where \mathbf{g} is a gradient field, G_σ is a Gaussian of width σ , σ_1 and σ_2 are two parameters such that $\sigma_1 < \sigma_2$, \otimes is the convolution operator, and $n(\cdot)$ a noise-controlling function. Our scheme relies on image gradients, for instance $T(\nabla I)$ is the texture map of the image I . We compare the average amplitude of the




Fig. 4. We show the same dots on a uniform background (a) and on a photograph (b) but the two dots on the tree are not visible because of the texture of the foliage

gradients in two neighborhoods defined by σ_1 and σ_2 . If the image is locally textured, then the average in the small neighborhood will be higher than in the large neighborhood, corresponding to $T > 1$. Conversely, $T < 1$ corresponds to regions with locally less texture than in the larger neighborhood. This scheme would be sensitive to noise in smooth regions where gradients are mostly due to noise. We address this issue with the function n that is equal to 0 for very small gradient and 1 otherwise. In practice, we use a smooth step equal to 0 for the bottom 2% of the intensity scale and 1 for 4% and above. In our context, the goal is to differentiate textureless areas from textured regions; the relative amount of texture in textured regions does not matter. Consequently, we use the estimator $\bar{T} = \min(1, T)$ that considers all textured regions to be equal.

Computing the Boundary Weights. Recall that we want W_β to be large when both foreground and background regions have the same degree of texture, either both smooth or both textured. If one of them is smooth and the other textured, we want W_β to be small. We consider the difference $D = |\bar{T}(\nabla F) - \bar{T}(\nabla B)|$ and define W_β using a function that assigns a small value w when D is large, 1 when D is small, and linearly transitions between both values. Formally, we use:

$$W_\beta = \begin{cases} w & \text{if } D > \tau \\ \min(1, w + \lambda(1 - D/\tau)) & \text{otherwise} \end{cases} \quad (6)$$

where λ and τ control the linear transition. We found that $\lambda = 4$, $w = 0.05$, $\tau = 0.005$, $\sigma_1 = 0.5$, and $\sigma_2 = 2$ work well in practice. All results are computed with these values unless otherwise specified.

Discussion. Figure 5 illustrates the effect of our approach that reduces the curl on the compositing boundary. Bleeding artifacts are significantly reduced. In next section, we describe how to remove the remaining leaks. For color images, we use RGB gradients in Equation (5) so that we account for luminance and chrominance textures. From an efficiency standpoint, an important characteristic of our approach is that it can be solved with a sparse linear system since our least-squares energy (Eq. 4) involves only sparse linear operators and W_β depends only on the input data.

3 Controlling the Location of the Residuals

Although our boundary treatment reduces the curl of the gradient field \mathbf{v} , in general \mathbf{v} is not integrable. As with other gradient-domain methods, our goal is to produce an image with a gradient field ∇I as close as possible to \mathbf{v} . Our strategy is to modify the Poisson equation (Eq. 2) in order to locate the residuals as much as possible in regions where they will be the least objectionable. Intuitively, we want to avoid errors in smooth regions such as the sky where they produce color leaks and halos, and put them in textured areas where visual masking will conceal the artifacts (Fig. 4).

3.1 Adapting the Poisson Equation

Let's assume that we have a scalar map W_P with high values in regions where errors would be visible and low values otherwise. We discuss later how to compute




Fig. 5. To reduce the color bleeding artifacts visible in a Poisson composite (a), we proceed in two steps. We adjust the gradient values at the boundaries to minimize the curl and weight the reconstruction process so that residual is mostly concentrated in the textured regions. While these two steps improve the results when applied separately (b,c), combining them achieves a visually superior composite (d).

such a function using our texture estimator \bar{T} . Given W_P , we modulate the least-squares strength so that we penalize less the regions where we prefer the residuals to be, that is, regions with low W_P values:

$$\operatorname{argmin}_I \int W_P \|\nabla I - \mathbf{v}\|^2 \quad (7)$$

Since we want to reduce the difference between ∇I and \mathbf{v} , W_P has to be strictly positive everywhere. Moreover, to keep our approach computationally efficient, we will design W_P such that it does not depend on the unknown image I . In this case, Equation 7 is a classical least-squares functional that can be minimized by solving a linear system. To obtain a formula similar to the Poisson equation (2), we apply the Euler-Lagrange formula [25]. Recall that W_P does not depend on I . Thus, we obtain the following linear system:

$$\operatorname{div}(W_P(\nabla I - \mathbf{v})) = 0 \quad (8)$$

In Section 3.2, we show that although this equation is simple, it has favorable properties.

Computing the Weights. To keep our scheme linear, we do not use any quantity related to the unknown I . We use the desired gradient field \mathbf{v} to estimate the texture location in the image. Although \mathbf{v} is not equal to the gradient of final output, it is a good approximation that is sufficient to compute the weights W_P . Since we want high weights in smooth regions and low weights in textured areas, we use the following formula: $W_P = 1 - p \bar{T}(\mathbf{v})$ where p is a global parameter that indicates how much we control the residual location. For instance, $p = 0$ corresponds to no control, that is, to the standard Poisson equation, whereas larger values impose more control. p has to be strictly smaller than 1 to keep $W_P > 0$. We found that values close to 1 performs better in practice. We use $p = 0.999$ in all our results. We also found that it is useful to have a more local estimate of the texture, which we achieve using $\sigma_1 = 0$ to compute \bar{T} while keeping the other parameters unchanged.

3.2 Analysis of the Residual Structure

Independent of the actual definition of W_P , we can show that the residuals produced by our approach have structure that is aligned with the image content. Wang et al. [26] have demonstrated that such structural similarity produces more acceptable results. To better understand the role of W_P , we distribute the divergence in Equation 8: $W_P \operatorname{div}(\nabla I - \mathbf{v}) + \nabla W_P \cdot (\nabla I - \mathbf{v}) = 0$. With $W_P \neq 0$, the relation $\operatorname{div}(\nabla I) = \Delta I$, and the logarithmic gradient $\nabla W_P / W_P = \nabla \log W_P$, we obtain:

$$\underbrace{\Delta I - \operatorname{div}(\mathbf{v})}_{\text{Poisson term}} + \underbrace{\nabla \log W_P \cdot (\nabla I - \mathbf{v})}_{\text{new term}} = 0 \quad (9)$$

The left term is the same as the standard Poisson equation (2) while the right term is new. In regions where W_P is constant, the new term is null and our scheme behaves as the Poisson equation, that is, it spreads the residuals uniformly. In other regions where W_P varies, our scheme differs from the Poisson equation and allows for discontinuities in the residual. Since W_P measures the amount of texture, it means that residual variations are aligned with texture edges, which ensures the structural similarity that has been shown desirable by Wang et al. [26]. We provide illustrations of this property in supplemental material.

3.3 Relationship with Existing Methods

For this section, we make explicit the variable W_P , that is, Equation 7 becomes $\int W_P(\mathbf{v}) \|\nabla I - \mathbf{v}\|^2$, and Equation 8, $\operatorname{div}(W_P(\mathbf{v}) (\nabla I - \mathbf{v})) = 0$. We discuss the relationships among our work and related methods independently of the actual definition of W_P .

The Poisson Equation and its Variants. Rewriting the Poisson equation (2) as $\operatorname{div}(\nabla I - \mathbf{v}) = 0$, we see that our linear system has the same complexity since we do not introduce new unknowns nor new coefficients in the system; we only reweight the coefficients. Agrawal et al. [14] also describe an anisotropic variant that is linear. However, while this method performs well in shape-from-shading, it does not prevent bleeding when applied to image compositing (Fig. 6). The L_1 reconstruction method that Reddy et al. [18] propose in the context of shape-from-shading has the same difficulty with image compositing (Fig. 6).

Edge-preserving Filtering. Our method is also related to Farbman’s edge-preserving filter [27] that minimizes an attachment term plus $\int W_P(I_0) \|\nabla I\|^2$ where I_0 is the input image. Farbman projects the formula on the x and y axes but we believe that it does not have a major impact on the results. More importantly, Farbman’s method and ours share the idea of using a modulation W_P that depends on fixed quantities and preserves the least-squares nature of the problem; Farbman uses the input image I_0 and we use the target gradient field \mathbf{v} . Finally, our work has common points with Perona and Malik’s nonlinear anisotropic diffusion filter [28]: $\partial I / \partial t = \operatorname{div}(W_P(\nabla I) \nabla I)$. The difference is that our modulation term W_P is not a function of the image I which makes our equation linear, and we have a term $\nabla I - \mathbf{v}$ instead of ∇I , which can be interpreted as Perona and Malik “diffuse gradients” whereas we “diffuse integration residuals.”




Fig. 6. We compare several approaches on an example where we composite a tree on a sky background. To test the robustness against selection inaccuracies, we introduce three errors (a): a small error on the left, a large error on the right, and the trunk is inserted in the ground. A direct copy-and-paste produces an image with visible seams in the sky region (b). Poisson compositing [1] (c), maximum gradient [1] (d), diffusion [14] (e), Photo Clip Art [16] (f), and robust Poisson reconstruction using the L_1 norm [18] (g) generate seamless boundaries but suffer from bleeding artifacts where the selection cuts through the foliage and also at the contact between the trunk and the ground. In comparison, our method (h) produces artifact-free results. We provide more comparisons in supplemental material.

4 Results

We demonstrate our approach on a typical hand-made compositing scenario which may generate inaccurate selections (Fig. 6). We also show that our approach applies to heterogeneous panorama stitching [7] (Fig. 10) and image reshuffling [11] (Fig. 11). More results are in our supplemental material. All the results are computed using the same parameters unless otherwise specified. These settings performed well in all of our experiments. Parameter variations are also shown in the supplemental material.

Quantitative Evaluation. We use direct compositing and Poisson compositing as baselines to estimate how much bleeding occurs. For direct compositing, we directly copy pixel values and the result I_d exhibits visible seams but not bleeding. For Poisson compositing, we copy gradient values and solve the Poisson equation. The result I_P is seamless but colors leak where the selection is inaccurate. Then we consider an image I , pick a pixel \mathbf{p} in the potential leaking area, and compute: $\|I(\mathbf{p}) - I_d(\mathbf{p})\| / \|I_P(\mathbf{p}) - I_d(\mathbf{p})\|$. Expressed in percentages, 0% indicates no bleeding at all and 100% indicates as much bleeding as Poisson compositing. Figure 7 compares the results for several methods.

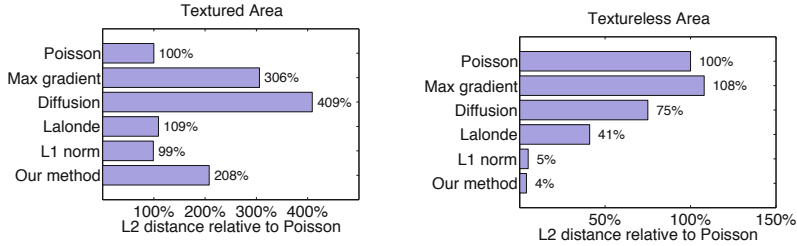


Fig. 7. We numerically evaluate bleeding introduced by different methods. We selected two 11×11 regions in the tree example (Fig. 6), one within the textured area below the trunk and one in the sky on right of the foliage. We compute the L_2 RGB difference between the image before and after compositing, normalized relative to Poisson compositing; that is, 100% indicates as much bleeding as Poisson and 0% indicates no bleeding. In the textured region (left), all methods bleed but the bleeding is masked by the high frequency texture. In the textureless area (right), most methods cause visible bleeding, which is particularly visible in this smooth region. The L_1 -norm and our method achieve similarly low values which confirm minimal bleeding. But in a number of cases, the L_1 -norm method introduces an undesirable color cast shown in the tree example, whereas our method yields a satisfying output.

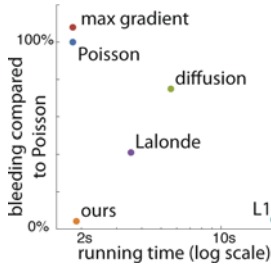


Fig. 8. This plot locates each method according to its speed and how much bleeding it introduces in the sky region on Figure 6 as reported in Figure 7. Our method is as fast as the standard Poisson solver while introducing almost no bleeding. In comparison, the other methods are slower and generate color leaks. Note that the L_1 method does not produce bleeding artifact on this example but it creates a severe color cast (Fig. 6).

Complexity. We compute the final result in two linear steps. This is equivalent to a single linear system because I is a linear function of \mathbf{v} (Eq. 8) and \mathbf{v} is a linear function of B and F (Eq. 4). Further, only sparse operators are involved: divergence, curl, and weights that correspond to diagonal matrices. Compared to the Poisson equation, we solve for the same number of unknowns, that is, the number of pixels in I . The only overhead is the computation of \mathbf{v} , for which the number of unknowns is orders of magnitude smaller, since only pixels on the boundary are concerned. To summarize, our method has the same complexity as the Poisson equation. In comparison, nonlinear methods [18] require more complex iterative solvers. Figure 8 shows that our implementation achieves timings similar to the Poisson reconstruction, resulting in a run-time faster than most other implementations while introducing almost no bleeding.



Fig. 9. In some cases, when compared to the input (a), Poisson compositing (b) and our approach (c) discolor the pasted region. See text for details.



Fig. 10. For heterogeneous panorama, Photoshop Auto Blend [6] produces strong bleeding near the cut. In comparison, our method significantly improves the result. Our approach also performs better than to other methods on this challenging case as shown in supplemental material.




Fig. 11. Compared to the blending approach proposed by Cho et al. [11] (a,b), our approach (c,d) improves the result of image reshuffling. We used the same patch locations and boundaries as Cho et al. but applied our method which yields better results than the Poisson-based blending proposed in the original article [11]. In particular, our result produces more faithful colors but does have local color leaks as can be seen on the close-up (zoom of a region above the girl's hat). This result may be better seen in the supplemental material. Data courtesy of Tim Cho.

Discussion. Although our method produces high quality outputs, a close examination reveals that the boundary can be sometimes overly sharp. This minor issue is difficult to spot at first and less conspicuous than color leaks. Nonetheless, matching the sharpness of other edges in the image would be an interesting extension to this work. As other gradient-domain methods, our method can yield some discoloration (Fig. 9 and supplemental material). This effect is often desirable to achieve seamless blending. If one wishes to preserve the original colors, matting can be solution but it often requires a more careful user input. We also found that our approach is useful in challenging applications such as heterogeneous panorama stitching [7] where mismatches are common place (Fig. 10). In this case, we found that our method performs better with a smoother transition from seamless and leak-free compositing, which is achieved by setting $\tau = 0.01$ in Equation (6).

5 Conclusion

We have described an image-compositing method that is robust to selection inaccuracies. The combination of low-curl boundaries and a weighted reconstruction based on visual masking produces artifact-free results on a broad range of inputs, in particular where other methods have difficulties. In addition, the solution is linear and has similar complexity to the standard Poisson equation. With robust results and speed, our method is a suitable replacement for the standard Poisson equation in many computer vision applications.

Acknowledgments. The authors thank Todor Georgiev for the link with the Poisson equation, Kavita Bala and George Drettakis for their discussion about visual masking, Aseem Agarwala and Bill Freeman for their help with the paper, Tim Cho and Biliانا Kaneva for helping with the validation, Medhat H. Ibrahim for the image of the Egyptian pyramids, Adobe Systems, Inc. for supporting Micah K. Johnson's research, and Ravi Ramamoorthi for supporting Michael Tao's work. This material is based upon work supported by the National Science Foundation under Grant No. 0739255 and No. 0924968.

References

1. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Trans. on Graphics* 22 (2003)
2. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D.H., Cohen, M.F.: Interactive digital photomontage. *ACM Trans. on Graphics* 23 (2004)
3. Georgiev, T.: Covariant derivatives and vision. In: *European Conf. on Computer Vision* (2006)
4. Jia, J., Sun, J., Tang, C.K., Shum, H.Y.: Drag-and-drop pasting. *ACM Trans. on Graphics* 25 (2006)
5. Levin, A., Zomet, A., Peleg, S., Weiss, Y.: Seamless image stitching in the gradient domain. In: *European Conf. on Computer Vision* (2006)
6. Agarwala, A.: Efficient gradient-domain compositing using quadtrees. *ACM Trans. on Graphics* 26 (2007)

7. Sivic, J., Kaneva, B., Torralba, A., Avidan, S., Freeman, W.T.: Creating and exploring a large photorealistic virtual space. In: IEEE Workshop on Internet Vision (2008)
8. Whyte, O., Sivic, J., Zisserman, A.: Get out of my picture! Internet-based inpainting. In: British Machine Vision Conf. (2009)
9. Finlayson, G.D., Drew, M.S., Lu, C.: Entropy minimization for shadow removal. *International Journal of Computer Vision* (2009)
10. Hays, J., Efros, A.A.: Scene completion using millions of photographs. *ACM Trans. on Graphics* 26 (2007)
11. Cho, T.S., Avidan, S., Freeman, W.T.: The patch transform. *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2010)
12. Tappen, M.F., Adelson, E.H., Freeman, W.T.: Recovering intrinsic images from a single image. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 27 (2005)
13. Finlayson, G.D., Hordley, S.D., Lu, C., Drew, M.S.: On the removal of shadows from images. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 28 (2006)
14. Agrawal, A., Raskar, R., Chellappa, R.: What is the range of surface reconstructions from a gradient field? In: European Conf. on Computer Vision (2006)
15. Bhat, P., Zitnick, C.L., Cohen, M., Curless, B.: Gradientshop: A gradient-domain optimization framework for image and video filtering. *ACM Trans. on Graphics* 28 (2009)
16. Lalonde, J.F., Hoiem, D., Efros, A., Rother, C., Winn, J., Criminisi, A.: Photo clip art. *ACM Trans. on Graphics* 26 (2007)
17. Farbman, Z., Hoffer, G., Lipman, Y., Cohen-Or, D., Fattal, R., Lischinski, D.: Coordinates for instant image cloning. *ACM Trans. on Graphics* 28 (2009)
18. Reddy, D., Agrawal, A., Chellappa, R.: Enforcing integrability by error correction using L-1 minimization. In: Computer Vision and Pattern Recognition (2009)
19. Drettakis, G., Bonneel, N., Dachsbacher, C., Lefebvre, S., Schwarz, M., Viaud-Delmon, I.: An interactive perceptual rendering pipeline using contrast and spatial masking. *Rendering Techniques* (2007)
20. Ramnarayanan, G., Ferwerda, J., Walter, B., Bala, K.: Visual equivalence: Towards a new standard for image fidelity. *ACM Trans. on Graphics* 26 (2007)
21. Vangorp, P., Laurijssen, J., Dutré, P.: The influence of shape on the perception of material reflectance. *ACM Trans. on Graphics* 26 (2007)
22. Ramnarayanan, G., Bala, K., Ferwerda, J.: Perception of complex aggregates. *ACM Trans. on Graphics* 27 (2008)
23. Su, S., Durand, F., Agrawala, M.: De-emphasis of distracting image regions using texture power maps. In: ICCV Workshop on Texture Analysis and Synthesis (2005)
24. Bae, S., Paris, S., Durand, F.: Two-scale tone management for photographic look. *ACM Trans. on Graphics* 25 (2006)
25. Aubert, G., Kornprobst, P.: Mathematical problems in image processing: Partial Differential Equations and the Calculus of Variations. *Applied Mathematical Sciences*, vol. 147. Springer, Heidelberg (2002)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Trans. on Image Processing* 13 (2004)
27. Farbman, Z., Fattal, R., Lischinski, D., Szeliski, R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Trans. on Graphics* 27 (2008)
28. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis Machine Intelligence* 12 (1990)