

Poisson Image Editing Extended

Daniel Leventhal*

Bernard Gordon*
Brown University

Peter G. Sibley*



Figure 1: (a): Mixed Seamless Cloning image of a sand spiral applied to a grassy field. Notice the abrupt end of the spiral. (b): Alpha Interpolation in RGB space. Notice the brightly colored artifacts. (c): Alpha Interpolation in YIQ space, copying IQ channels directly from the target image. Shimmer effect from (b) has been removed, but there are unnatural bright/dark spots on the spiral. (d): Alpha Interpolation in YIQ space, with Luminance Rescaling ($\gamma = 0.6$). The texture is much more subtle than in (c), and the pots have been removed

We present a method for transferring textures from a source image to a target with fewer artifacts than previous implementations. We based our work on [Pérez et al. 2003], improving their mixed seamless cloning technique. We created a technique called Alpha Interpolation to remove abrupt edges in the resulting image caused by mixed seamless cloning, as well as Luminance Rescaling to allow the user to scale the influence of the source image on the result. Finally, we applied Poisson Image Editing to video to explore the coherence issues that it presents.

Poisson Image Editing creates a new image f from a source image g and a destination image f^* by composing a guiding vector field v based on f^* and g , and then finding f such that the gradient of f is closest in the L2-norm to v over a domain Ω , while holding $f = f^*$ over the boundary of Ω (Dirchelet boundary conditions), so that the edited image matches up with the rest of the image while the edited region can have features from both f^* and g . It is well known that the solution of the minimization problem is also the solution of the Poisson equation, finding f such that the gradient of f is closest (L2-norm) to some vector field v , called the guiding vector field. Using this interpretation, we can write the Poisson equation with Dirchelet boundary conditions as: $\Delta f = \nabla \cdot v$, with $f^*|_{\partial\Omega} = f|_{\partial\Omega}$. This technique is later referred to as Gradient Domain Fusion when used in [Agarwala et al. 2004]

Alpha Interpolation

Perez et al.’s Mixed Seamless Cloning filter solves the Poisson equation using a guiding vector field composed of the maximum (by norm) of the gradient of the source and destination images, in order to capture visually interesting elements from both images. However, at the edge of the region of interest, even though the Dirchelet boundary conditions ensure that the pixel values match, macroscopic effects end very abruptly (Figure 1(a)).

Instead of using a mask with values in $\{0, 1\}$, the user creates a trimap by selecting a region of definite interest ($\alpha = 1$) and a region that is definitely not of interest ($\alpha = 0$). We then create an alpha mask $\alpha \in [0, 1]$ for the entire source image where the unselected region is blended from 0 to 1 using a series of gaussian blurs. Between each blur we reset the area of interest to 1 and the area not of interest to 0. We then scale the source gradient by our alpha before using it in mixed seamless cloning. This makes the mixed seamless cloning algorithm less likely to select features from the source image near the edges of the domain where α is small, and makes the contribution that the source does make around the edges less prominent. With alpha interpolation, our guiding field looks

like this:

$$\forall x \in \Omega, v(x) = \begin{cases} \nabla f^*(x) & \text{if } \|\nabla f^*(x)\| > \alpha \|\nabla g(x)\| \\ \alpha \nabla g(x) & \text{otherwise} \end{cases}$$

We solve the Poisson equation for each channel (RGB) separately. Unfortunately, this can cause coherence problem resulting in an undesired rainbow-like shimmer effect (Figure 1(b)). To address this, we ran it in YIQ space on only the luminance channel, copying chrominance directly from the destination image. This produces good results, and has the desired effect of smoothly fading out macroscopic effects (Figure 1(c)).

Luminance Rescaling

If a source image has large gradients in the luminance channel, those gradients can cause undesired effects in the resultant image (Figure 1(c)). We addressed this by rescaling the luminance by a constant factor towards the mean luminance, μ , over Ω in the source image: $L_p = \gamma L_p + (1 - \gamma)\mu$. This method allows us to smooth out sharp gradients in the luminance channel and remove the undesired artifacts these can cause, but also decreases the gradient everywhere, removing the texture that we wish to copy (Figure 1(d)).

Video Editing

We also ran Poisson Image Editing on videos, using a constant, regular region that we placed at a sequence of points on the source and destination videos as they played. Our method is similar to Panaramic video textures [Agarwala et al. 2005]. This worked as expected, with each frame being acceptable on its own, but coherence problems between frames causing a very noticeable halo effect.

AGARWALA, A., DONTCHEVA, M., AGRAWALA, M., DRUCKER, S., COLBURN, A., CURLISS, B., SALESIN, D., AND COHEN, M. 2004. Interactive digital photomontage. *ACM Transactions on Graphics* 23, 3 (Aug.), 294–302.

AGARWALA, A., ZHENG, K. C., PAL, C., AGRAWALA, M., COHEN, M., CURLISS, B., SALESIN, D. H., AND SZELISKI, R. 2005. Panoramic video textures. *ACM Transactions on Graphics* 24, 3 (Aug.), 821–827.

PÉREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. *ACM Transactions on Graphics* 22, 3 (July), 313–318.

*e-mail: {dlev, bernie, pgs}@cs.brown.edu