Poisson Image Editing using Perona-Malik Diffusion

Vishnu Nittoor

Abstract—Poisson Image Editing is a well-established technique for seamlessly compositing objects from one image into another by working in the gradient domain. However, when users draw coarse masks that do not conform precisely to object boundaries, the mask often intersects strong edges in the target image. Poisson Image Editing has no mechanism to preserve these edges, resulting in blurred background structures near the compositing boundary. Prior approaches address this by modifying the guidance field or adding edge-aware regularization terms, but these either introduce transparency effects or require additional constraints. We propose a different approach inspired by Perona-Malik anisotropic diffusion: rather than modifying the guidance or adding constraints, we replace the standard Laplacian with an edge-aware diffusion operator. The conductivity is derived from gradient magnitudes in both source and target images, inhibiting diffusion at strong edges in either. A sensitivity parameter controls edge-stopping strength, with standard Poisson editing recovered in the limit. We demonstrate improved Structural Similarity Index (SSIM) scores on multiple test scenes, showing better preservation of background edges while maintaining or improving source object structure.

Index Terms—Poisson Image Editing, Perona-Malik Diffusion, Computational Photography, Image Blending

1 Introduction

Poisson Image Editing [1] is a well-established technique for seamlessly compositing objects from one image into another. Rather than directly copying pixel values—which creates visible seams due to lighting and color mismatches—the method works in the gradient domain. A *guidance field*, typically the gradient of the source image, specifies the desired local contrast within the compositing region, while boundary conditions anchor the solution to the target image. The result is a blended image that preserves the internal structure of the source object while adapting its overall appearance to match the surrounding background (Figure 1).

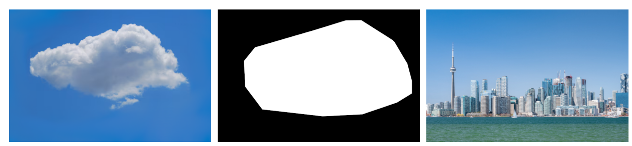

Fig. 1. The seamless cloning problem. Given a source image (left), a coarse mask defining the region to import (center), and a target image (right), the goal is to produce a composite where the source object appears naturally integrated into the target scene.

In practice, users often draw *coarse masks* that do not conform precisely to object boundaries. When these masks intersect background structures—fences, window blinds, building edges—standard Poisson Image Editing blurs across them. Although the framework provides mechanisms to incorporate target edge information by modifying the guidance field, these approaches render the source object partially transparent rather than preserving it opaquely.

Prior work has made Poisson Image Editing edge-aware through two main strategies. One approach modifies the guidance field itself, transitioning smoothly from source gradients in the interior to target-influenced gradients near the boundary [2]. Another adds explicit edge-aware regularization terms to the optimization, penalizing edits that cross strong color boundaries [3]. These approaches modify what is being diffused or how the solution is constrained, but not the underlying mechanism responsible for blurring, that is the diffusion process itself.

We explore a different direction inspired by Perona-Malik anisotropic diffusion [4]. Rather than modifying the guidance field or adding constraints, we replace the standard Laplacian with an edge-aware diffusion operator whose conductivity depends on gradient magnitudes in both the source and target images. Where either image contains a strong edge, diffusion is inhibited, preserving that structure in the composite. A sensitivity parameter κ controls the strength of edge-stopping; in the limit $\kappa \to \infty$, the method recovers standard Poisson Image Editing.

Our contributions are as follows:

- A Poisson image editing formulation with an edge-aware diffusion operator that respects edges in both source and target images.
- A sensitivity parameter that controls edge-stopping strength, with standard Poisson editing as a limiting case.
- Qualitative and quantitative evaluation demonstrating improved preservation of background edges while maintaining or improving source object structure.

Our method does not address color bleeding artifacts and can sometimes exacerbate them; additionally, the sensitivity parameter may require qualitative tuning, as the chosen metric does not always correspond to visual quality.

- *V. Nittoor is with the University of Toronto.*

2 Related Work

2.1 Gradient-Domain Image Editing

The foundational framework for our research is that of gradient-domain image editing, primarily established by Pérez et al. [1] in their seminal work on Poisson Image Editing. Motivated by the observation that the human visual system is significantly more sensitive to local contrast changes (gradients) than to absolute luminance values, this paradigm shifts the editing focus to incorporate the gradient information of a source image S into a target image T.

Mathematically, a scalar function on a bounded domain is uniquely defined by its values on the boundary and its Laplacian in the interior. This framework solves for a blended image f (or B) over the mask region Ω that matches the target image T on the boundary $\partial\Omega$, while its internal diffusion behavior matches that of a *guidance vector field* $\mathbf{v}$:

$$\Delta f = \nabla \cdot \mathbf{v} \quad \text{in } \Omega, \tag{1}$$

$$f = T \quad \text{on } \partial\Omega \tag{2}$$

This is equivalent to solving the variational problem that minimizes the difference between the gradient of the reconstructed image and the guidance field:

$$\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \quad \text{with} \quad f|_{\partial\Omega} = T|_{\partial\Omega} \tag{3}$$

To implement this numerically, the continuous domain is discretized onto the image pixel grid. The variational problem above translates to minimizing a discrete energy functional involving the finite differences between nearest neighbors. Specifically, for each pixel p and its neighbors $q \in N_p$, we seek the pixel values f_p that minimize the sum of squared differences between the gradient of the reconstructed image and the guidance field:

$$\min_{f|_\Omega} \sum_{\langle p,q\rangle \cap \Omega \neq \emptyset} (f_p - f_q - v_{pq})^2 \quad \text{with} \quad f_p = T_p \text{ for } p \in \partial\Omega \tag{4}$$

where v_{pq} is the projection of the guidance vector $\mathbf{v}$ onto the directed edge (p, q). Minimizing this quadratic energy is equivalent to solving a sparse system of linear equations, typically involving a 5-point Laplacian stencil.

2.1.1 Guidance Field Strategies

Several options exist for constructing the guidance field $\mathbf{v}$. In this work, we focus on *seamless cloning* (also called *importing gradients*), where $\mathbf{v} = \nabla S$. This produces an opaque transfer of the source object into the target scene, adjusting its color and lighting to match the surrounding background while preserving its internal structure.

Alternative strategies exist that incorporate gradient information from the target image. An α-blended combination $\mathbf{v} = \alpha\nabla S + (1 - \alpha)\nabla T$ mixes both gradient fields, while *mixed gradients* selects the larger-magnitude gradient at each pixel. However, both approaches produce a transparency effect: the target's texture bleeds through the source object in regions where the source is relatively smooth. While useful for certain applications, these strategies are unsuitable when opaque insertion of the source object is required, which is precisely the setting we address.

Zhang et al. [2] proposed a spatially-varying approach that creates a *gradient transition map* near the mask boundary. Rather than applying uniform blending, they reconstruct the guidance field to smoothly transition from source gradients in the interior to target-influenced gradients near $\partial\Omega$, producing more natural compositions. Their method also avoids solving the full linear system by using an interpolation-based approximation. While this gradient transition approach modifies the guidance field $\mathbf{v}$, our method instead modifies the diffusion operator's conductivity—the two strategies are complementary and could potentially be combined.

2.2 Edge-Aware Optimization

Hua et al. [3] incorporates edge-awareness by reinterpreting the gradient domain framework as a variational problem with explicit edge-aware constraints. They proposed augmenting the standard energy functional with a term $E_e(p)$ for each pixel p in its 4-connected neighborhood $N_4(p)$ based on *edit propagation*:

$$E_e(p) = \sum_{q \in N_4(p)} w^e_{pq} \left[(f(p) - u(p)) - (f(q) - u(q))\right]^2 \tag{5}$$

Here, the term $(f - u)$ represents the "edit" or change applied to a pixel. The weight w^e_{pq} is defined using a Gaussian of the color difference in the Lab color space:

$$w^e_{pq} = \exp\left(-\frac{||u(p) - u(q)||^2}{\sigma^2}\right) \tag{6}$$

By minimizing this term, the framework penalizes differences in the *change* applied to neighboring pixels that are similar color. This effectively creates a barrier to diffusion at strong edges: if $u(p)$ and $u(q)$ are very different (an edge), w^e_{pq} approaches zero, and the constraint is relaxed, allowing the edits at p and q to diverge independently.

Hua et al.'s method represents the closest prior work to ours, as both aim to make Poisson editing edge-aware while preserving opacity. However, our approach operates differently: rather than adding constraint terms to the energy functional, we modify the diffusion operator itself. By replacing the standard Laplacian with the nonlinear Perona-Malik operator, we embed edge awareness directly into the diffusion process. The two approaches are complementary—a modified diffusion operator could be integrated into a constraint-based variational framework—making the investigation of Perona-Malik diffusion in Poisson Image Editing a distinct and valuable avenue of study.

2.3 Nonlinear Diffusion

Our approach draws on the theory of anisotropic diffusion introduced by Perona and Malik [4] as a method for scale-space analysis and edge detection. Their work was motivated by the limitations of standard linear scale-space filtering (Gaussian blurring), which suffers from two key defects: it displaces edge locations at coarse scales and blurs region boundaries, destroying semantic edge junctions.

To resolve this, Perona and Malik introduced a nonlinear diffusion process that satisfies three criteria: causality (no new features are generated at coarse scales), immediate localization (edges remain sharp and in place), and piecewise smoothing (intra-region smoothing is preferred over inter-region smoothing).

The anisotropic diffusion equation is defined as:

$$\frac{\partial I}{\partial t} = \nabla \cdot (c(x, y, t)\nabla I) \tag{7}$$

where $c(x, y, t)$ is the diffusion coefficient. Unlike the heat equation where c is constant, Perona and Malik define c as a monotonically decreasing function of the gradient magnitude $|\nabla I|$. This allows the diffusion process to adapt locally:

$$c(|\nabla I|) = g(|\nabla I|) \tag{8}$$

In homogenous regions where $|\nabla I|$ is low, $g(|\nabla I|) \to 1$, acting as standard Gaussian smoothing. At strong edges where $|\nabla I|$ is high, $g(|\nabla I|) \to 0$, effectively stopping the flux and preserving the edge structure.

One of the two functions proposed for the conduction coefficient is

$$g_1(\nabla I) = \exp\left(-\left(\frac{|\nabla I|}{K}\right)^2\right) \tag{9}$$

where K is a constant parameter regulating the sensitivity to edges, known for privileging high-contrast edges.

Our work imports this anisotropic diffusion framework into the Poisson image editing problem, using the Perona-Malik conductivity function to inhibit diffusion across strong edges near the mask boundary. A key departure from standard anisotropic diffusion is that we compute the conductivity map *a priori* from the known source and target images, rather than from the evolving solution itself. In the original Perona-Malik formulation, conductivities depend on $|\nabla I|$ where I is the image being diffused, leading to an iterative nonlinear process. In contrast, our conductivities are fixed before solving, yielding a single linear system. This design choice reflects the structure of the editing problem: we know exactly where edges exist in both S and T, and we wish to preserve them in B.

3 METHOD

3.1 Continuous Formulation

We propose to replace the standard Laplacian operator in the Poisson Image Editing framework with the edge-aware Perona-Malik diffusion operator. The classical problem minimizes the Dirichlet energy $\iint |\nabla B - \mathbf{v}|^2$. We instead consider a *weighted* Dirichlet energy:

$$E(B) = \iint_\Omega c(\mathbf{x})|\nabla B - \mathbf{v}|^2 dx \tag{10}$$

where $c(\mathbf{x})$ is a spatially varying conductivity map. Minimizing this energy leads to the following Euler-Lagrange equation (the weighted Poisson equation):

$$\nabla \cdot (c(\mathbf{x})\nabla B) = \nabla \cdot (c(\mathbf{x})\mathbf{v}) \quad \text{in } \Omega \tag{11}$$

$$B = T \quad \text{on } \partial\Omega \tag{12}$$

Here, $c(\mathbf{x}) \in [0, 1]$ acts as a ”permeability” or ”diffusivity” map. In regions where $c(\mathbf{x})$ is high, the equation behaves like the standard Poisson equation, effectively smoothing out differences between the reconstructed gradient ∇B and the guidance field $\mathbf{v}$. However, where $c(\mathbf{x})$ is low, the diffusion flux is throttled. If we align low conductivity regions with the strong edges of the target image, we effectively ”block” the seamless cloning process from blurring across those edges, thereby preserving the structural integrity of the background.

3.2 Proposed Method: Edge-Aware Poisson

Our proposed method defines the conductivity map $c(\mathbf{x})$ based on the fixed structures of the source and target images. Following Perona and Malik [4], we define a diffusivity function g that decays with increasing gradient magnitude:

$$g(|\nabla I|) = \frac{1}{1 + \left(\frac{|\nabla I|}{\kappa}\right)^2} \tag{13}$$

where κ is a user-defined sensitivity parameter. This function evaluates to near 1 when the local gradient is small (flat regions) and approaches 0 as the gradient magnitude significantly exceeds κ (edges). The effect of varying κ on the composite is illustrated in Figure 2.

Since our goal is to seamlessly blend a source object S into a target background T, we must respect edges in *both* images. We wish to preserve the internal details of the source object while preventing the diffusion from overrunning strong boundaries in the target. Therefore, we define our combined conductivity map as the pointwise minimum of the diffusivities of both the source and the target:

$$c(\mathbf{x}) = \min\left(g(|\nabla S(\mathbf{x})|),\ g(|\nabla T(\mathbf{x})|)\right) \tag{14}$$

By taking the minimum, we ensure that diffusion is inhibited if *either* the source or the target possesses a strong edge at location $\mathbf{x}$. This conservative approach is crucial for preventing artifacts: it halts smoothing at the object's internal boundaries (preserving source detail) and at the background's external boundaries (preventing bleeding). This method is computationally efficient as the weights are pre-computed once.

3.3 Discrete Formulation

The continuous problem is discretized on a 2D pixel grid. We employ a standard 5-point finite difference stencil. Since the conductivity $c(\mathbf{x})$ is defined at pixel centers but acts on the flux between pixels, we must estimate its value at the inter-pixel edges. We define the edge weights w_{pq} between a pixel $p = (x, y)$ and its neighbor q (e.g., $q = (x + 1, y)$) as the arithmetic mean of their conductivities:

$$w_{pq} \approx \frac{c(p) + c(q)}{2} \tag{15}$$

The discrete weighted Poisson equation at pixel p then becomes:

$$\sum_{q \in N_p} w_{pq}(B_p - B_q) = \sum_{q \in N_p} w_{pq} \cdot v_{pq} \tag{16}$$

where N_p is the set of 4-connected neighbors of p. The term v_{pq} represents the projection of the guidance field $\mathbf{v}$ onto the directed edge from p to q. Specifically, if q is the right neighbor of p (i.e., $q = (x + 1, y)$), then $v_{pq} \approx \mathbf{v}_x(p)$.

$\kappa = 0.01$ (Low) $\kappa = 0.02$ (Optimal) $\kappa = 0.03$ $\kappa = 0.06$ (High)

Fig. 2. Effect of the sensitivity parameter κ on blending quality for the "Cloud" scene. At low κ (left), diffusion is strongly inhibited, causing visible artifacts where edge barriers prevent proper blending. At the optimal $\kappa \approx 0.02$ (center-left), the method achieves the best balance: the cloud is opaquely blended while preserving the sharp building edges in the background. As κ increases (right), the edge-stopping weakens and color from the source begins to bleed across background structures, approaching standard Poisson Image Editing.

Equation 16 forms a system of linear equations $\mathbf{Ax} = \mathbf{b}$, where $\mathbf{x}$ is the vector of unknown pixel colors inside the mask Ω. The matrix $\mathbf{A}$ is a sparse, symmetric, positive-definite matrix (specifically, a weighted Laplacian matrix) with diagonal elements $A_{pp} = \sum_{q \in N_p} w_{pq}$ and off-diagonal elements $A_{pq} = -w_{pq}$.

For pixels p on the boundary of the mask ($\partial\Omega$), the value B_q for a neighbor $q \notin \Omega$ is known (it is the target image value T_q). These known values are moved to the right-hand side of the equation, acting as Dirichlet boundary conditions:

$$b_p = \sum_{q \in N_p} w_{pq} v_{pq} + \sum_{q \in N_p \setminus \Omega} w_{pq} T_q \tag{17}$$

3.4 Implementation Details

We implemented the proposed method in Python. The linear system $\mathbf{Ax} = \mathbf{b}$ is solved using `scipy.sparse.linalg.spsolve`, which wraps the efficient SuperLU direct solver.

The guidance field $\mathbf{v}$ is typically derived from the source image gradients (∇S), but our implementation also supports the "Mixing Gradients" strategy defined in Equation (4). The discrete gradients are computed using forward finite differences. The computational complexity is dominated by the linear solver. Since the matrix $\mathbf{A}$ is sparse (with at most 5 non-zero entries per row) and symmetric, the system can be solved efficiently.

4 Experimental Results

To evaluate compositing fidelity, we utilize the Structural Similarity Index (SSIM) [5] rather than Mean Squared Error (MSE), as it allows us to quantify how well the structural details (edges and textures) of the source object are preserved, even when its color temperature is significantly altered to match the target background. We evaluate our method by SSIM inside the masked region Ω against both the target (SSIM$_T$) and source (SSIM$_S$). Higher SSIM values indicate better preservation of edges and textures. All experiments use 600×400 RGB images.

4.1 Comparison with Standard Poisson Image Editing

We compare our edge-aware method against standard Poisson Image Editing (PIE), which corresponds to the limit $\kappa \to \infty$ where diffusion becomes isotropic. For each test scene, we select the optimal κ by sampling 20 values linearly in $[0, 0.1]$ and choosing the parameter that maximizes the combined fidelity score SSIM$_T$ + SSIM$_S$. An exception is the "Jail" scene, where this heuristic encounters a failure mode discussed in Section 4.3.

TABLE 1
Quantitative Comparison. SSIM measured inside the mask region Ω. For most scenes, we use the κ which maximizes the combined fidelity score SSIM$_T$ + SSIM$_S$. For Jail, we also report the source-maximizing result, which produces qualitatively superior output (see Section 4.3).

Scene	κ	Target (SSIM$_T$)		Source (SSIM$_S$)	
		Std. PIE	Ours	Std. PIE	Ours
Cloud	0.016	0.38	**0.46**	0.50	**0.59**
Spike	0.021	0.53	**0.69**	0.43	**0.61**
Fencedog	0.016	0.29	**0.41**	0.40	**0.87**
Plane	0.021	0.49	**0.53**	0.47	0.46
Jail (sum-max)†	0.032	0.03	**0.05**	0.54	**0.58**
Jail (src-max)†	0.058	0.03	0.04	0.54	**0.58**

† See Sec. 4.3 for discussion of the Jail failure mode.

Quantitative Results. As summarized in Table 1, our method consistently outperforms standard PIE across all test scenes. We achieve improvements in target fidelity (SSIM$_T$) ranging from $+0.04$ on "Plane" to $+0.16$ on "Spike", while often simultaneously improving source fidelity (SSIM$_S$). The largest measured improvement occurs on "Fencedog", where SSIM$_S$ improves from 0.40 to 0.87; however, as discussed below, significant high-frequency noise artifacts in this scene may inflate the SSIM score.

Qualitative Analysis. The quantitative gains correspond to visible improvements in the composited images. On the "Spike" scene (Figure 5), standard PIE blurs the horizontal edges of the window blinds in the background, causing a loss of structural detail near the mask boundary. Our method preserves the sharp edge structure of the blinds while maintaining the original appearance of the source object. On "Fencedog" (Figure 6), our method dramatically improves awareness of the fence structure near the mask boundary, though the nonlinear diffusion introduces some susceptibility to color bleeding artifacts and high frequency noise. The "Plane" scene (Figure 7) shows how standard PIE allows cloud texture to blur into the aircraft structure near the wings, while our method restricts diffusion across these strong edges. Finally, the "Jail" scene (Figure 4) illustrates both the strength and limitation of our approach: while the

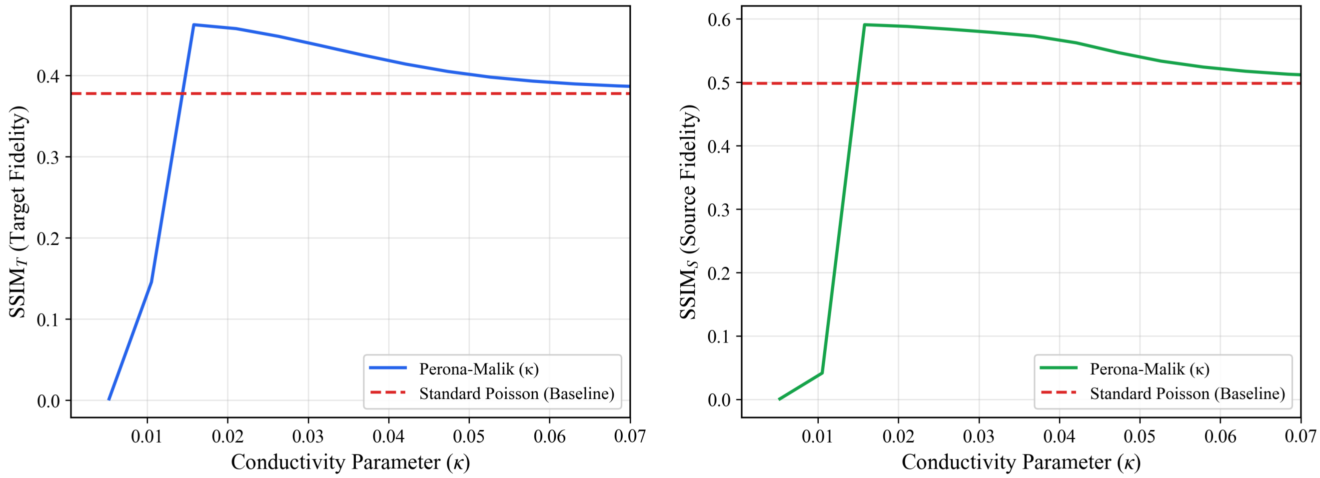

(a) Target Fidelity (SSIM$_T$) (b) Source Fidelity (SSIM$_S$)

Fig. 3. Sensitivity analysis of conductivity parameter κ for the "Cloud" scene. We vary κ linearly from 0 to 0.07. At $\kappa < 0.005$, diffusion is fully inhibited, causing numerical stalling (SSIM ≈ 0). Peak performance occurs around $\kappa \approx 0.012$, significantly outperforming the Standard Poisson baseline (dashed line). Beyond $\kappa > 0.05$, the method asymptotically converges to standard Poisson.

(a) Standard PIE (b) Sum-max ($\kappa \approx 0.032$) (c) Source-max ($\kappa \approx 0.058$)

Fig. 4. Comparison on the "Jail" scene. (a) Standard Poisson Image Editing causes significant bleeding of the sign colors into the jail bars. (b) The sum-maximizing κ produces artifacts due to the dark background failure mode (see Section 4.3). (c) The source-maximizing κ achieves opaque blending while respecting the jail bar edges.

(a) Standard PIE ($\kappa \to \infty$) (b) Our Method ($\kappa \approx 0.021$)

Fig. 5. Comparison on the "Spike" scene. Standard Poisson Image Editing blurs the horizontal edges of the blinds in the background, causing a loss of structural detail near the mask boundary. In contrast, our method preserves the sharp edge structure of the blinds while maintaining the original appearance of the source object.

(a) Standard PIE ($\kappa \to \infty$) (b) Our Method ($\kappa \approx 0.016$)

Fig. 6. Comparison on the "Fencedog" scene. Our method improves awareness of the fence structure close to the mask boundaries. However, the non-linear diffusion results in increased susceptibility to colour bleeding and high frequency noise artifacts near the grass boundary.

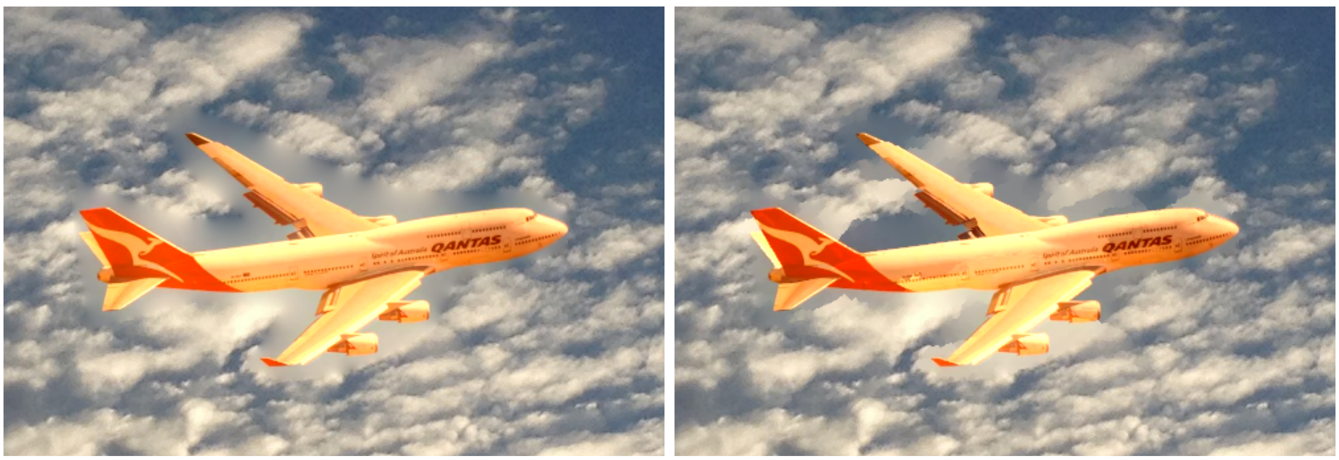

(a) Standard PIE ($\kappa \to \infty$) (b) Our Method ($\kappa \approx 0.021$)

Fig. 7. Comparison on the "Plane" scene. Standard Poisson Image Editing allows the cloud texture to blur into the aircraft structure, particularly noticeable near the wings. In contrast, our method restricts diffusion across these strong edges, resulting in less blurring of the clouds.

source-maximizing κ successfully prevents the sign colors from bleeding into the jail bars, the sum-maximizing heuristic fails due to the dark background (see Section 4.3).

4.2 Ablation Study: Sensitivity to Conductivity

We investigate the sensitivity of our method to the conductivity parameter κ. This analysis effectively functions as an ablation study of the edge-aware component: as $\kappa \to \infty$, the diffusivity function $g(\cdot) \to 1$, recovering the isotropic diffusion of the standard Poisson Image Editing baseline.

Quantitative Analysis. Figure 3 reveals three distinct operational regimes. If $\kappa < 0.005$, the diffusion coefficient vanishes ($g \approx 0$), isolating pixels from the boundary conditions and causing numerical stalling in the linear solver (SSIM ≈ 0). The method enters its *optimal range* at $\kappa \approx 0.01$, where it discriminates between noise and structural edges, significantly outperforming the baseline (SSIM$_T = 0.46$ vs 0.38). Beyond $\kappa > 0.05$, the edge-stopping barrier relaxes, and performance asymptotically converges to the baseline.

One might expect a trade-off where improved target blending degrades source fidelity. This was reported in the poster I submitted, which inadvertently used 1D SSIM. However, with SSIM correctly computed in two dimensions, we observe that optimal κ improves *both* metrics, until they decay with diminishing edge-stopping effect. Our hypothesis is as follows: by inhibiting diffusion upto a point, our method effectively relaxes the extent to which source information must compromise non-gradient information such as the average brightness and colour temperature to satisfy the boundary conditions, resulting in higher $SSIM_S$.

Qualitative Analysis. This progression is visually confirmed in Figure 2. At low κ, the excessive inhibition creates "cut-and-paste" artifacts where the source fails to integrate with the target. The optimal setting ($\kappa \approx 0.02$) achieves a seamless blend while maintaining opacity. As κ increases, the method re-introduces the background bleeding artifacts characteristic of standard Poisson editing.

4.3 Discussion

The "Jail" scene (Figure 4) reveals a limitation of the SSIM metric's relevance to this application. At low κ values, numerical difficulties produce near-black output in the blended region. This achieves *high* SSIM_T because a black output is structurally similar to the predominantly black target background, yet clearly a failed blend is not preferable to a successfully imported source object. Similarly, high-frequency noise artifacts—such as those visible in the "Fencedog" result—can inflate SSIM scores by introducing spurious structure. This exposes a fundamental issue: SSIM measures structural similarity without semantic understanding of what constitutes a "good" composite. For scenes with uniform backgrounds or significant noise, alternative metrics or human evaluation may be necessary.

While our results show clear improvements in edge preservation, we acknowledge that the coarse mask problem could also be addressed by image matting techniques that compute precise alpha mattes along the boundary. Our method does not address the more challenging color bleeding artifacts that occur when gradients within the source region conflict with target gradients—the problem that Hua et al.'s constraint-based approach [3] specifically targets.

Despite these limitations, our work demonstrates that edge-awareness can indeed be incorporated into gradient-domain editing by modifying the underlying PDE operator, rather than solely through additional constraint terms. The Perona-Malik conductivity provides a principled mechanism for spatially varying the diffusion strength based on local image structure. Importantly, this operator-based approach is *orthogonal* to both constraint-based formulations like that of Hua et al. and guidance-field modifications like Zhang et al.'s gradient transition maps [2]—a modified diffusion operator could be integrated with either approach. Future work may explore how these complementary strategies can be combined: using nonlinear diffusion to respect edges at the mask boundary, gradient transitions for smoother blending, and edge-aware constraints to handle color bleeding artifacts.

References

[1] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Transactions on Graphics (TOG)*, vol. 22, no. 3, pp. 313–318, 2003.

[2] Y. Zhang, J. Ling, X. Zhang, and H. Xie, "Image copy-and-paste with optimized gradient," *The Visual Computer*, vol. 30, no. 10, pp. 1169–1178, 2014.

[3] M. Hua, X. Bie, M. Zhang, and W. Wang, "Edge-aware gradient domain optimization framework for image filtering by local propagation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 4305–4312.

[4] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–639, 1990.

[5] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.