

Article

Not peer-reviewed version

A Variational Screened Poisson Reconstruction for Whole-Slide Stain Normalization

[Junlong Xing](#)^{*}, Hengli Ni, [Qiru Wang](#), Yijun Jing

Posted Date: 28 February 2026

doi: 10.20944/preprints202602.1947.v1

Keywords: stain normalization; screened Poisson equation; Helmholtz operator; domain decomposition



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Variational Screened Poisson Reconstruction for Whole-Slide Stain Normalization

Junlong Xing^{1,*} , Hengli Ni², Qiru Wang¹ and Yijun Jing¹

¹ School of Mathematics, Sun Yat-Sen University, Guangzhou, China

² Department of Pathology, Children's Hospital of Soochow University, Soochow University, Suzhou, Jiangsu, China

* Correspondence: xingjlong@mail2.sysu.edu.cn

Abstract

To address staining variability in digital pathology, we formulate normalization as a variational inverse problem balancing structural fidelity and chromatic alignment. We propose *Screened Poisson Normalization* (SPN), a convex model in CIE $L^*a^*b^*$ space governed by the modified Helmholtz equation. By coupling a first-order gradient consistency term with a zeroth-order chromatic anchor, SPN yields a screened reconstruction controlled by a single screening parameter λ . Beyond establishing well-posedness in $H^1(\Omega)$ and Lipschitz stability, we derive a property crucial for gigapixel whole-slide analysis: the screening term induces an intrinsic localization length $\ell \sim \lambda^{-1/2}$, which leads to exponential decay of boundary-induced errors. This theoretical insight provides a principled justification for a scalable, non-overlapping subdomain implementation; when combined with an efficient DCT-based spectral solver, it enables artifact-free processing without elaborate tiling heuristics or boundary post-processing. Extensive experiments on multi-center datasets further demonstrate that SPN achieves robust chromatic consistency while faithfully preserving diagnostically relevant microstructural detail.

Keywords: stain normalization; screened Poisson equation; Helmholtz operator; domain decomposition

1. Introduction

1.1. Background

Histopathological analysis is the current gold standard for disease diagnosis, particularly in oncology. With the rapid evolution of digital pathology [1], the slides containing stained tissue sections are routinely digitized into ultra-high-resolution whole-slide images (WSIs) using advanced microscopy scanners. As a result, computational analysis of digital pathology images is increasingly complementing, and in some cases replacing, manual microscopic evaluation by pathologists. Within this paradigm, the quality, consistency, and standardization of digitized pathological images constitute pivotal determinants of diagnostic reproducibility and the performance of subsequent computational analyses.

The WSIs imaging pipeline begins with conventional tissue preparation, followed by whole-slide scanning into standard RGB image formats. Figure 1 provides an overview of the complete acquisition workflow [28].




Figure 1. Typical workflow of WSIs imaging , illustrating color variations arising from staining and scanner differences. Staining protocol: standard H&E (HE), long time stained Eosin (longE), and long time stained H&E (longHE). Scanners: Aperio ScanScope CS2 (CS2), NanoZoomer 2.0-HT (NZ20), NanoZoomer S210 (NZ210), and Aperio GT 450 (GT450).

Clinically, tissue specimens (obtained via surgical resection or biopsy) are fixed in formalin to preserve structural integrity, dehydrated through graded ethanol solutions, and embedded in paraffin wax to form a block. Thin tissue sections (3–6 μm) are cut from this block, mounted onto slides, and stained, most commonly with hematoxylin and eosin (H&E). Finally, a whole-slide scanner digitizes the slide at subcellular resolution to generate a WSI that is stored in a gigapixel images. However, deviations inevitably accumulate due to heterogeneous tissue preparation (e.g., sectioning, reagents, pH) and diverse digitization conditions (e.g., scanner optics, sensor noise). These artifacts introduce substantial domain shifts that impair diagnostic consistency and algorithmic performance. To address this, robust stain normalization is required to harmonize WSIs from heterogeneous sources, thereby standardizing the input for reliable downstream computational analysis.

1.2. Related Works

Research on stain normalization has advanced significantly since the digitalization of pathology. The seminal work of Reinhard et al. [6] introduced statistical color alignment via mean–variance matching in the decorrelated Lab space. However, this global statistics-based normalization inevitably results in the over-smoothing of fine image details.

To improve illumination robustness and chromatic fidelity, subsequent variations modified the transformation domain. Zarella et al. [14] and Magee et al. [15] adopted color projection in HSV/Lab spaces; Roy et al. [17] and Vijh et al. [18] introduced fuzzy-logic constraints; and Nadeem et al. [20] reformulated stain normalization as an optimal-transport problem via multimarginal Wasserstein barycenters. While these methods enhance perceptual smoothness, they remain fundamentally pixel-domain and global in nature, limiting their ability to preserve tissue microstructures.

Macenko et al. [7] shifted attention to structure preservation through stain vector estimation in optical-density (OD) space via singular value decomposition (SVD). Vahadane et al. [8] further incorporated sparse non-negative matrix factorization (SNMF), which jointly optimizes stain basis and concentration matrices. Adaptive deconvolution [19] and IDA decomposition [16] have since been proposed to address spatial heterogeneity while maintaining textural and morphological integrity. Gupta et al. [9] proposed a geometry-inspired, chemical- and tissue-invariant stain-normalization framework (GCTI). GCTI uses SVD/NMF and explicit geometric alignment to correct color variations caused by illumination, stain color vectors, and stain quantity.

Meanwhile, with the advancement of digital pathology, data-driven deep learning models have demonstrated superior performance to conventional stain normalization techniques in recent years. Existing approaches are mainly characterized by two classes: (1) supervised learning models that require paired or carefully curated training data [3,5,13,22–25], and (2) generative frameworks based

on Generative Adversarial Networks (GANs) [4,11,12,21,26,33]. Despite achieving strong quantitative results, such improvements do not necessarily translate into better clinical diagnostic performance. When applied across institutions, sample-driven deep learning models are vulnerable to distribution shifts, which undermine reliability in real-world deployment. Moreover, their lack of explicit physical modeling of the staining process limits interpretability and hinders clinical acceptance.

1.3. Main Contributions

We propose the **Screened Poisson Normalization (SPN)** model, a variational framework that reformulates stain normalization as a strictly convex inverse problem governed by reaction-diffusion physics. This approach integrates physical interpretability with rigorous mathematical analysis, establishing the well-posedness and stability of the solution to guarantee the preservation of diagnostic morphology against input perturbations. To address the gigapixel scale of clinical pathology, we further derive a domain decomposition algorithm underpinned by the exponential decay of the screened Poisson Green's function; this theoretical property ensures that interface errors remain exponentially localized, enabling the seamless and parallel reconstruction of Whole-Slide Images without the boundary artifacts typical of patch-based processing.

The paper is organized as follows. Section 2 investigates the mathematical nature of the stain normalization problem and establishes its theoretical foundations. Section 3 then derives *Screened Poisson Normalization* (SPN) from a reaction-diffusion interpretation of histological staining and formulates stain standardization as a variational inverse problem in the perceptual CIE *Lab* space. Section 4 presents the rigorous analysis of the model, establishing well-posedness, Lipschitz stability, and the exponential localization property that motivates the non-overlapping subdomain implementation. Section 5 details the numerical realization, including the finite-difference discretization, the DCT-based spectral solver, and practical implementation considerations. Section 6 reports comprehensive experimental validation on multi-center datasets. Finally, Section 7 concludes the paper and discusses limitations and directions for future work.

2. Mathematical Foundations

We define stain normalization as an inverse problem on a bounded Lipschitz domain $\Omega \subset \mathbb{R}^2$. Let $I_S, I_T : \Omega \rightarrow \mathbb{R}^3$ denote the source and reference images, respectively. The task is to recover a normalized field $u : \Omega \rightarrow \mathbb{R}$ (processed channel-wise in the decorrelated CIE *Lab* space [6]) that simultaneously satisfies structural fidelity to I_S and statistical alignment with I_T . This section analyzes the ill-posed nature of this reconstruction and establishes its theoretical connection to reaction-diffusion physics.

2.1. Variational Formulation of the Inverse Problem

Mathematically, recovering u is fundamentally under-determined. The statistical constraint operator $\mathcal{C}(u) \approx \mathcal{C}(I_T)$ is non-injective, as infinite spatial configurations can share identical global moments. Conversely, relying solely on gradient fidelity $\nabla u \approx \nabla I_S$ determines the solution only up to an additive constant (the null space of the Neumann operator). Consequently, the unconstrained optimization is *ill-posed* in the sense of Hadamard [38].

To restore well-posedness, the problem must be cast within a Tikhonov regularization framework [39]. We seek a solution $u \in H^1(\Omega)$ that minimizes a strictly convex energy functional of the form:

$$\mathcal{J}(u) := \frac{1}{2} \|\nabla u - \mathbf{g}\|_{L^2(\Omega)}^2 + \frac{\lambda}{2} \|u - T\|_{L^2(\Omega)}^2, \quad \lambda > 0. \quad (1)$$

Here, the first term enforces structural consistency with a guidance field $\mathbf{g}$ derived from the source, while the second term anchors the solution to a chromatic prior T . The regularization parameter λ balances these competing objectives and ensures the coercivity of the bilinear form, guaranteeing the existence of a unique, stable weak solution.

2.2. Physical Isomorphism: Reaction–Diffusion Kinetics

The variational structure of (1) is not merely a mathematical regularization; it is isomorphic to the physical transport laws [30] governing histological staining. The accumulation of dye within porous tissue follows reaction–diffusion kinetics [27,29,31]. For a dye species with concentration $C(\mathbf{x}, t)$, the governing equation is:

$$\partial_t C = \nabla \cdot (D \nabla C) + R(C, \mathbf{x}), \quad (2)$$

where D is the diffusion coefficient and R describes the binding reaction. Near chemical equilibrium, the binding process approximates linear kinetics $R(C) \approx -k(C - C_{\text{eq}})$, where k is the reaction rate and C_{eq} is the local equilibrium concentration determined by tissue affinity.

At the thermodynamic steady state ($\partial_t C = 0$), assuming isotropic diffusion, (2) reduces to the screened Poisson equation:

$$-\Delta C + \kappa C = \kappa C_{\text{eq}}, \quad \text{with } \kappa := k/D. \quad (3)$$

This physical equilibrium (3) shares the exact Euler-Lagrange form of the variational problem (1). The regularization parameter λ thus corresponds to the physical ratio of reaction to diffusion rates (κ), while the anchor T corresponds to the chemical equilibrium potential. This theoretical equivalence provides the physical basis for modeling stain normalization as a screened Poisson process.

3. The Screened Poisson Normalization (SPN) Model

This section details the Screened Poisson Normalization (SPN) model. Operating in the CIE *Lab* color space [6] and using optical density (OD) to represent dye concentration [33,34], we cast stain normalization as a variational reconstruction problem.

For each channel $c \in \{L^*, a^*, b^*\}$, we define the normalized field $u_c \in H^1(\Omega)$ as the unique minimizer of the strictly convex energy functional:

$$\mathcal{J}_c(u_c) := \int_{\Omega} \left(\frac{1}{2} \|\nabla u_c - \mathbf{g}_c\|^2 + \frac{\lambda_c}{2} (u_c - T_c)^2 \right) d\mathbf{x}. \quad (4)$$

The first term ensures structural fidelity by penalizing gradient deviations from a guidance field $\mathbf{g}_c$, while the second acts as a Tikhonov regularizer, anchoring the solution to a chromatic prior T_c . The parameter $\lambda_c > 0$ governs the trade-off between these objectives and ensures the problem is well-posed.

Minimizing (4) yields the associated Euler-Lagrange equation, which takes the form of a steady-state screened Poisson equation:

$$-\Delta u_c + \lambda_c u_c = \lambda_c T_c - \nabla \cdot \mathbf{g}_c \quad \text{in } \Omega, \quad (5)$$

subject to the adiabatic boundary condition $(\nabla u_c - \mathbf{g}_c) \cdot \mathbf{n} = 0$ on $\partial\Omega$. Physically, this equation describes a diffusion process driven by a source term $f_c := \lambda_c T_c - \nabla \cdot \mathbf{g}_c$, which balances the injection of high-frequency morphological details with low-frequency chromatic information.

To close the system, we construct the chromatic anchor T_c and guidance field $\mathbf{g}_c$ from the source I_S and reference I_T . The anchor T_c is defined by matching the global mean μ and standard deviation σ of the reference:

$$T_c(\mathbf{x}) = \frac{\sigma_T^{(c)}}{\sigma_S^{(c)}} ((I_S)_c(\mathbf{x}) - \mu_S^{(c)}) + \mu_T^{(c)}. \quad (6)$$

This linear mapping provides a stable baseline for chromaticity that is robust to local variations. To preserve morphology, we parameterize the guidance field $\mathbf{g}_c$ as a spatially varying blend of source and anchor gradients:

$$\mathbf{g}_c(\mathbf{x}) = (1 - w(\mathbf{x})) \alpha_c \nabla (I_S)_c(\mathbf{x}) + w(\mathbf{x}) \nabla T_c(\mathbf{x}), \quad (7)$$

where α_c is a robust contrast factor defined by the ratio of median gradient magnitudes, i.e., $\alpha_c = \text{median} \|\nabla(I_T)_c\| / \text{median} \|\nabla(I_S)_c\|$. The structural weight $w(\mathbf{x}) \in (0, 1]$ is derived from the source luminance channel I_{S,L^*} via:

$$w(\mathbf{x}) = \exp\left(-\frac{\|\nabla I_{S,L^*}(\mathbf{x})\|^2}{\tau^2}\right), \quad \tau > 0. \quad (8)$$

In structurally salient regions ($w \rightarrow 0$), the guidance field strictly adheres to the source morphology $\nabla(I_S)_c$, whereas in homogeneous regions, the model relaxes toward the smooth chromatic anchor. The resulting solution corresponds to the steady state of a system with screening length $\ell_c \sim \lambda_c^{-1/2}$, preventing local artifacts from propagating globally.

4. Theoretical Analysis

This section analyzes the mathematical properties of the Screened Poisson Normalization (SPN) model. We treat the variational formulation as a linear inverse problem in the Sobolev space $H^1(\Omega)$, establishing the well-posedness of the solution, its stability with respect to input data, and the convergence of the domain decomposition scheme used for gigapixel processing.

The variational problem derived in Section 3 is formally equivalent to finding a weak solution $u \in H^1(\Omega)$ such that

$$B(u, v) = F(v), \quad \forall v \in H^1(\Omega), \quad (9)$$

where the bilinear form $B : H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{R}$ and the linear functional $F : H^1(\Omega) \rightarrow \mathbb{R}$ are defined by:

$$B(u, v) := \int_{\Omega} (\nabla u \cdot \nabla v + \lambda uv) d\mathbf{x}, \quad F(v) := \int_{\Omega} (\mathbf{g} \cdot \nabla v + \lambda T v) d\mathbf{x}.$$

We assume $\Omega \subset \mathbb{R}^2$ is a bounded Lipschitz domain, $\lambda > 0$, $\mathbf{g} \in [L^2(\Omega)]^2$, and $T \in L^2(\Omega)$.

Theorem 1 (Existence and Uniqueness). *For any screening parameter $\lambda > 0$, guidance field $\mathbf{g} \in [L^2(\Omega)]^2$, and chromatic anchor $T \in L^2(\Omega)$, the variational problem (9) admits a unique solution $u^* \in H^1(\Omega)$.*

Proof. The result follows from the Lax–Milgram theorem. First, the bilinear form B is continuous on $H^1(\Omega) \times H^1(\Omega)$ by the Cauchy–Schwarz inequality:

$$|B(u, v)| \leq \|\nabla u\|_{L^2} \|\nabla v\|_{L^2} + \lambda \|u\|_{L^2} \|v\|_{L^2} \leq \max(1, \lambda) \|u\|_{H^1} \|v\|_{H^1}.$$

Second, B is strictly coercive due to the strictly positive screening parameter λ . Testing with $v = u$:

$$B(u, u) = \|\nabla u\|_{L^2}^2 + \lambda \|u\|_{L^2}^2 \geq \min(1, \lambda) \|u\|_{H^1}^2.$$

Since F defines a bounded linear functional on $H^1(\Omega)$, the existence of a unique minimizer is guaranteed. $\square$

With the solution uniquely defined, we quantify its stability with respect to perturbations in the input data $(\mathbf{g}, T)$.

Theorem 2 (Lipschitz Stability). *Let $u_1, u_2 \in H^1(\Omega)$ be the unique weak solutions corresponding to data $(\mathbf{g}_1, T_1)$ and $(\mathbf{g}_2, T_2)$, respectively. The solution depends continuously on the data:*

$$\|u_1 - u_2\|_{H^1(\Omega)} \leq \frac{1}{\min(1, \lambda)} \left(\|\mathbf{g}_1 - \mathbf{g}_2\|_{L^2(\Omega)} + \lambda \|T_1 - T_2\|_{L^2(\Omega)} \right). \quad (10)$$

Proof. Let $w = u_1 - u_2$. Subtracting the respective weak formulations yields

$$B(w, v) = \int_{\Omega} ((\mathbf{g}_1 - \mathbf{g}_2) \cdot \nabla v + \lambda(T_1 - T_2)v) d\mathbf{x}, \quad \forall v \in H^1(\Omega).$$

Choosing the test function $v = w$ and applying the coercivity of B alongside the Cauchy–Schwarz inequality gives

$$\min(1, \lambda) \|w\|_{H^1}^2 \leq B(w, w) \leq (\|\mathbf{g}_1 - \mathbf{g}_2\|_{L^2} + \lambda \|T_1 - T_2\|_{L^2}) \|w\|_{H^1}.$$

Dividing by $\|w\|_{H^1}$ completes the proof. $\square$

We next quantify the trade-off between morphological preservation and color normalization. The following estimate ensures that structural modifications are strictly bounded by the magnitude of the chromatic transfer.

Lemma 1 (Energy Comparison Principle). *Let u^* be the global minimizer of $\mathcal{J}(u)$. For any comparison function $v \in H^1(\Omega)$,*

$$\|\nabla u^* - \mathbf{g}\|_{L^2(\Omega)}^2 + \lambda \|u^* - T\|_{L^2(\Omega)}^2 \leq \|\nabla v - \mathbf{g}\|_{L^2(\Omega)}^2 + \lambda \|v - T\|_{L^2(\Omega)}^2. \quad (11)$$

Proof. This is an immediate consequence of the optimality condition $\mathcal{J}(u^*) \leq \mathcal{J}(v)$. $\square$

Theorem 3 (Structural Fidelity Bound). *Assume $\mathbf{g} \approx \nabla I_S$. The L^2 -deviation of the normalized gradients ∇u^* from the guidance field satisfies:*

$$\|\nabla u^* - \mathbf{g}\|_{L^2(\Omega)} \leq \sqrt{\lambda} \|I_S - T\|_{L^2(\Omega)} + \|\nabla I_S - \mathbf{g}\|_{L^2(\Omega)}. \quad (12)$$

Proof. Set $v = I_S$ in Lemma 1. Neglecting the non-negative term $\lambda \|u^* - T\|_{L^2}^2$ on the left-hand side yields

$$\|\nabla u^* - \mathbf{g}\|_{L^2}^2 \leq \|\nabla I_S - \mathbf{g}\|_{L^2}^2 + \lambda \|I_S - T\|_{L^2}^2.$$

The result follows from the subadditivity of the square root function, $\sqrt{a^2 + b^2} \leq a + b$. $\square$

Finally, we address the convergence of the tiled reconstruction required for gigapixel Whole Slide Images. Let Ω be partitioned into non-overlapping subdomains $\{\Omega_k\}$. The local solution u_{tile} on Ω_k satisfies the natural boundary condition $(\nabla u_{\text{tile}} - \mathbf{g}) \cdot \mathbf{n} = 0$ on $\partial\Omega_k$. The discrepancy $e := u_{\text{global}} - u_{\text{tile}}$ arises solely from the flux mismatch $\psi := (\nabla u_{\text{global}} - \mathbf{g}) \cdot \mathbf{n}$ at the interface $\Gamma := \partial\Omega_k \setminus \partial\Omega$.

Lemma 2 (Green's Function Asymptotics). *The fundamental solution $G_\lambda(\mathbf{x})$ of the operator $-\Delta + \lambda\mathcal{I}$ in $\mathbb{R}^2$ satisfies the decay estimate for $|\mathbf{x}| > 0$:*

$$|G_\lambda(\mathbf{x})| \leq C \frac{e^{-\sqrt{\lambda}|\mathbf{x}|}}{\sqrt{|\mathbf{x}|}}. \quad (13)$$

Theorem 4 (Exponential Localization of Interface Error). *For any $\mathbf{x} \in \Omega_k$, the error $e(\mathbf{x})$ decays exponentially with the distance to the interface Γ :*

$$|e(\mathbf{x})| \leq K \exp\left(-\frac{d(\mathbf{x}, \Gamma)}{\ell_c}\right), \quad \ell_c = \lambda^{-1/2}. \quad (14)$$

Proof. The error function satisfies $-\Delta e + \lambda e = 0$ in Ω_k with Neumann data $\partial_n e = \psi$ on Γ . Using the boundary layer potential representation

$$e(\mathbf{x}) = \int_\Gamma G_\lambda(\mathbf{x} - \mathbf{y}) \psi(\mathbf{y}) d\sigma(\mathbf{y}),$$

and noting that $|\mathbf{x} - \mathbf{y}| \geq d(\mathbf{x}, \Gamma)$ for all $\mathbf{y} \in \Gamma$, we invoke the monotonicity of the Green's kernel:

$$|e(\mathbf{x})| \leq \left(\sup_{r \geq d(\mathbf{x}, \Gamma)} |G_\lambda(r)| \right) \|\psi\|_{L^1(\Gamma)}.$$

Combining this uniform bound with the exponential decay of G_λ yields the stated estimate. $\square$

Remark 1. Theorem 4 implies that boundary artifacts are confined to a layer of characteristic width $\mathcal{O}(\ell_c)$. For sufficiently large tiles, the interior solution converges exponentially to the global reconstruction, mathematically justifying the use of non-overlapping domain decomposition.

5. Optimization and Implementation

For each channel $c \in \{L^*, a^*, b^*\}$, the SPN model seeks a scalar field $U_c : \Omega \rightarrow \mathbb{R}$ that solves the screened Poisson equation (5), where U_c denotes the discrete (grid-based) version of u_c .

5.1. Discrete and Non-Overlapping Domain Decomposition

The gigapixel resolution of WSI images renders global algorithms computationally prohibitive. Leveraging the localization properties of the SPN operator established in Section 4, we adopt a patch-wise implementation strategy.

We discretize the continuous domain Ω on a regular $H \times W$ lattice:

$$\Omega_h = \{(x, y) \in \mathbb{Z}^2 \mid 0 \leq x < H, 0 \leq y < W\},$$

with unit grid spacing. The scalar field U_c is represented as a matrix $\mathbf{U}_c \in \mathbb{R}^{H \times W}$, and the vector field $\mathbf{g}_c$ as a tensor in $\mathbb{R}^{H \times W \times 2}$. To avoid directional bias, we employ centered finite differences for the interior gradients:

$$(\nabla_h U_c)_{x,y} = \begin{bmatrix} D_x U_c \\ D_y U_c \end{bmatrix}_{x,y} = \begin{bmatrix} (U_{c,x+1,y} - U_{c,x-1,y})/2 \\ (U_{c,x,y+1} - U_{c,x,y-1})/2 \end{bmatrix}, \quad 1 \leq x \leq H-2, 1 \leq y \leq W-2. \quad (15)$$

At the image boundaries, we utilize forward/backward differences consistent with the homogeneous Neumann condition $\partial_n U_c = 0$.

The discrete divergence $\nabla_h \cdot \mathbf{g}_c$ is defined as the negative adjoint of ∇_h , ensuring that the discrete operator retains the symmetric positive-definite structure of the continuous formulation. Combining ∇_h and $\nabla_h \cdot$ yields the standard five-point Laplacian stencil:

$$(\Delta_h U_c)_{x,y} = U_{c,x+1,y} + U_{c,x-1,y} + U_{c,x,y+1} + U_{c,x,y-1} - 4U_{c,x,y}, \quad (16)$$

with appropriate modifications at the boundaries. On the vectorized domain, this defines a sparse block-Toeplitz with Toeplitz-blocks (BTTB) matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$ (where $N = HW$). Solving the linear system

$$(\mathbf{L} - \lambda_c \mathbf{I}) \mathbf{u}_c = \mathbf{r}_c$$

directly via sparse factorization is infeasible at the WSI scale. Instead, we exploit the spectral properties of $\mathbf{L}$ for efficient inversion.

5.2. Spectral Diagonalization and DCT Solver

Under homogeneous Neumann boundary conditions, the eigenfunctions of the discrete Laplacian (16) coincide with the basis functions of the 2D Type-II Discrete Cosine Transform (DCT-II).

Let $\mathcal{F}$ denote the orthonormal forward 2D DCT-II operator and $\mathcal{F}^{-1}$ its inverse (DCT-III) [42]. For an input image $\mathbf{F} \in \mathbb{R}^{H \times W}$, the spectral coefficients $\hat{\mathbf{F}} = \mathcal{F}\{\mathbf{F}\}$ are given by:

$$\hat{F}_{u,v} = \alpha_u \alpha_v \sum_{x=0}^{H-1} \sum_{y=0}^{W-1} F_{x,y} \cos\left[\frac{\pi u(2x+1)}{2H}\right] \cos\left[\frac{\pi v(2y+1)}{2W}\right], \quad (17)$$

where $0 \leq u < H, 0 \leq v < W$, and α_u, α_v are normalization constants. In this basis, the Laplacian is diagonal:

$$\mathcal{F}\{\Delta_h U_c\}_{u,v} = \Lambda_{u,v} \hat{U}_{c,u,v},$$

with eigenvalues given analytically by:

$$\Lambda_{u,v} = 2\left(\cos\frac{\pi u}{H} + \cos\frac{\pi v}{W}\right) - 4, \quad \Lambda_{u,v} \in [-8, 0]. \quad (18)$$

The mode $(u, v) = (0, 0)$ corresponds to the DC component with $\Lambda_{0,0} = 0$.

Let $\mathbf{R}_c = \nabla_h \cdot \mathbf{g}_c - \lambda_c \mathbf{T}_c$ denote the discrete right-hand side. Applying $\mathcal{F}$ to the discrete screened Poisson equation decouples the system:

$$(\Lambda_{u,v} - \lambda_c) \hat{U}_{c,u,v} = \hat{R}_{c,u,v}, \quad \forall (u, v). \quad (19)$$

The solution is obtained via a simple spectral filter:

$$\hat{U}_{c,u,v} = \frac{\hat{R}_{c,u,v}}{\Lambda_{u,v} - \lambda_c}, \quad \mathbf{U}_c = \mathcal{F}^{-1}\{\hat{\mathbf{U}}_c\}. \quad (20)$$

Stability and Regularization.

Since $\lambda_c > 0$ and $\Lambda_{u,v} \leq 0$, the denominator $\mathcal{D}_{u,v} := \Lambda_{u,v} - \lambda_c$ satisfies $\sup_{u,v} \mathcal{D}_{u,v} = -\lambda_c < 0$. In particular, the DC mode $(0, 0)$ is well-defined, eliminating the global offset ambiguity typical of pure Poisson reconstruction. From an optimization perspective, the screening term $\lambda_c \|\mathbf{U}_c - \mathbf{T}_c\|_2^2$ acts as a Tikhonov regularizer, stabilizing the inversion of low-frequency modes and anchoring the solution to $\mathbf{T}_c$.

5.3. Tiled SPN Implementation for Gigapixel WSIs

Direct application of the spectral solver to a full-resolution WSI is impractical due to memory constraints. We therefore adopt a non-overlapping domain decomposition [43] where the slide domain Ω is partitioned into tiles $\{\Omega_k\}_{k=1}^K$. The key requirement is that each local problem must approximate the restriction of the *single* global SPN functional.




Figure 2. Tiled SPN pipeline. In the first phase, we compute slide-level global prior statistics. In the second phase, we restrict these priors to each tile and solve the screened Poisson equation using a DCT-based spectral solver, assembling the normalized tiles into a coherent WSI.

Our implementation proceeds in two phases:

1. Global Prior Estimation. We first compute the global statistics of the source (I_S) and reference (I_T) slides. For each channel c , we calculate the global moments $(\mu_S^{\text{glob}}, \sigma_S^{\text{glob}})$ and (μ_T, σ_T) , define the chromatic anchor T_c via the affine map (Eq. 6), and estimate the global contrast factor α_c (Eq. 7). These quantities are computed once at the slide level and provide coherent priors for all local reconstructions.

2. Local Spectral Reconstruction. We solve the screened Poisson equation independently on each tile. For a tile Ω_k , we restrict the global priors ($T_{c,k} = T_c|_{\Omega_k}$ and $g_{c,k} = g_c^{\text{glob}}|_{\Omega_k}$), assemble the local right-hand side $\rho_{c,k} = \nabla_h \cdot g_{c,k} - \lambda_c T_{c,k}$, and solve:

$$(\Delta_h - \lambda_c) U_{c,k} = \rho_{c,k} \quad \text{on } \Omega_k. \quad (21)$$

The homogeneous Neumann boundary conditions are implicitly enforced by the DCT-II/III basis. The normalized WSI is obtained by concatenating the tile solutions, $(\tilde{I}_S)_c|_{\Omega_k} = U_{c,k}$. By Theorem 4, interface discrepancies are exponentially localized near the boundaries with characteristic width $\mathcal{O}(\lambda_c^{-1/2})$.

Algorithm 1 Tiled Screened Poisson Normalization (SPN)

Require: Source I_S , Reference I_T , parameters τ, λ_c , tiling $\{\Omega_k\}_{k=1}^K$
Ensure: Normalized WSI $\tilde{I}_S$

- 1: **Phase 1: Global Statistics**
- 2: **for** each channel $c \in \{L^*, a^*, b^*\}$ **do**
- 3: Compute global source moments on I_S and target moments on I_T
- 4: Define global chromatic anchor T_c via Eq. (6)
- 5: Compute global contrast factor α_c via Eq. (7)
- 6: **end for**
- 7: **Phase 2: Tile-wise Spectral Reconstruction**
- 8: **for** each tile Ω_k in parallel **do**
- 9: Extract source patch $P_k \leftarrow I_S|_{\Omega_k}$
- 10: Compute structure weight $w_k(\mathbf{x}) \leftarrow \exp(-\|\nabla(P_k)_{L^*}(\mathbf{x})\|^2/\tau^2)$
- 11: **for** each channel $c \in \{L^*, a^*, b^*\}$ **do**
- 12: Restrict anchor: $T_{c,k} \leftarrow T_c|_{\Omega_k}$
- 13: Form guidance field:
- 14: $\mathbf{g}_{c,k} \leftarrow (1 - w_k) \alpha_c \nabla(P_k)_c + w_k \nabla T_{c,k}$
- 15: Assemble RHS: $\rho_{c,k} \leftarrow \nabla_h \cdot \mathbf{g}_{c,k} - \lambda_c T_{c,k}$
- 16: Solve $(\Delta_h - \lambda_c) U_{c,k} = \rho_{c,k}$ via DCT spectral filter (20)
- 17: Update $\tilde{I}_S|_{\Omega_k} \leftarrow U_{c,k}$
- 18: **end for**
- 19: **end for**
- 20: **return** $\tilde{I}_S$

The boundary behavior of this tiled scheme is governed by the screened Poisson operator analyzed in Section 4. The screening term $\lambda_c > 0$ introduces a characteristic length scale $\ell_c \approx \lambda_c^{-1/2}$, ensuring exponential decay of flux mismatches away from tile interfaces. Consequently, the assembled reconstruction is structurally and chromatically consistent with the global solution, with seams reduced to negligible boundary effects.

5.4. Computational Complexity

The computational cost is dominated by the 2D DCT spectral solves. For a single tile with $N_k = H_k W_k$ pixels:

- **Assembly:** Computing gradients, weights, and the right-hand side $\rho_{c,k}$ scales linearly, i.e., $\mathcal{O}(N_k)$.
- **Spectral Solve:** The 2D DCT is implemented via separable 1D fast cosine transforms, with a cost of $\mathcal{O}(N_k \log N_k)$.

The total cost for a WSI with $N = \sum_k N_k$ pixels is $\mathcal{O}(N \log N)$, with memory complexity $\mathcal{O}(N_k)$ per tile. This approach allows for trivial parallelization over tiles. In comparison, sparse direct solvers typically scale as $\mathcal{O}(N^{1.5})$ for 2D problems, while multigrid methods, though $\mathcal{O}(N)$, require complex data structures. The tiled DCT implementation thus provides an effective balance between asymptotic efficiency and implementation simplicity for gigapixel images.

6. Experiments and Results

6.1. Experimental Setup

Datasets

To evaluate the effectiveness and robustness of the proposed algorithm, we conducted comparative experiments on multiple publicly available and clinical datasets exhibiting diverse staining conditions and imaging variations. In total, the evaluation encompassed 88 patients and more than 300 whole-slide histopathological images (WSIs).

- **Dataset 1:** MITOS-ATYPIA-14 [35], released as part of the MITOS-ATYPIA Grand Challenge, was annotated by the Pathology Department of Pitié-Salpêtrière Hospital in Paris. It consists of 14 pairs of breast carcinoma slides, each scanned using two distinct digital pathology scanners—Aperio

Table 1. Summary of datasets used for experimental evaluation.

Dataset	#Samples	#WSIs	Magnification	Color variation source	Body part
MITOS-ATYPIA-14	14	28	40×/0.25μm	Two imaging scanners	Breast
SCC dataset	44	220	40×/0.25μm	Five imaging scanners	Skin
HEV dataset	1	9	40×/0.25μm	Nine staining protocols	Thyroid
SUSY-BF-10	25	50	40×/0.25μm	Color fading in 10 years	Breast

ScanScope XT and Hamamatsu NanoZoomer—allowing controlled analysis of inter-scanner color variation.

- **Dataset 2:** Multi-Scanner Squamous Cell Carcinoma (SCC) dataset [36] contains 44 samples of canine cutaneous squamous cell carcinoma digitized using 5 different scanning systems, producing a total of 220 WSIs.
- **Dataset 3:** HE-Staining Variation (HEV) dataset [37] provided by Heidelberg University. The HEV dataset comprises follicular thyroid carcinoma slides prepared under nine distinct staining protocols: standard H&E, over-stained H&E (longHE), under-stained H&E (shortHE), hematoxylin-only (onlyH), eosin-only (onlyE), over-stained hematoxylin (longH), over-stained eosin (longE), under-stained hematoxylin (shortH), and under-stained eosin (shortE). These variations facilitate the evaluation of color normalization under chemically induced staining inconsistencies.
- **Dataset 4:** SUSY-BF-10 clinical dataset (local collection). To address the temporal color degradation that public datasets do not capture, we constructed a clinical dataset from **Sun Yat-sen Memorial Hospital, Sun Yat-sen University**. The SUSY-BF-10 dataset includes 25 long-term archived breast carcinoma slides, each preserved for over ten years and rescanned to analyze fading and chromatic shifts over time.

A summary of dataset characteristics and representative visualizations is provided in Table 1 and Figure 3.




Figure 3. Visualizations for the four datasets. The HEV dataset (row a) demonstrates the five different staining protocols on serial sections: (a1) longE, (a2) onlyE, (a3) H&E, (a4) shortE, and (a5) shortH. The SYSU-BF-10 dataset (row b), from (b1) to (b5), displays stained sections from five different patients, illustrating staining and fading conditions over a 10-year period. The SCC dataset (row c) presents imaging results acquired from five different scanning devices. Finally, in the MITOS-ATYPIA-14 dataset (row d), images (d1), (d2), and (d3) represent one set at three different magnification levels (10×, 20×, and 40×), while images (d4), (d5), and (d6) form a color control group scanned with a different device at the same three magnification levels.

Implementation Details

In surveying the literature on stain normalization, we observe that most existing methods are evaluated primarily on pre-extracted WSI tiles and are often positioned as a preprocessing step for downstream tasks such as classification or segmentation. In contrast, the ultimate goal of digital pathology in the clinical setting is to provide pathologists with visually consistent slides, especially in cross-site, multi-center collaborations and remote diagnostic workflows. Under these scenarios, a robust and interpretable color normalization scheme at the *whole-slide* level is of central importance.

Motivated by this observation, we design our experiments to assess performance from two complementary perspectives:

1. **Full-resolution WSI normalization.** We evaluate the behavior of each method when applied directly to entire, full-size histopathology slides.
2. **Tile-based reconstruction.** We further consider the practically common pipeline in which a large WSI is first partitioned into tiles, normalized tile-wise, and then reassembled into a full-resolution slide. This setting is particularly relevant for memory-constrained or streaming-based deployments.

In addition to the four real-world datasets used in our study, we also construct a synthetic test image (`synthetic_image`) with controlled color and structural patterns. This synthetic benchmark provides an intuitive and stress-test-style assessment of robustness and generalization across a wide range of stain variations.

All experiments are conducted on a server equipped with an Intel(R) Xeon(R) Platinum 8352V CPU @ 2.10 GHz, an NVIDIA RTX 4090 GPU with 24 GB of VRAM, and 512 GB of system memory. The software stack is based on Python 3.8 and PyTorch 2.5.

Baselines

We compare the proposed method against three representative baselines:

- **Reinhard.** The classical global Reinhard color transfer algorithm [6] serves as our primary reference for global-statistics-based normalization. It matches low-order statistics between source and target images in a decorrelated color space and remains widely used in digital pathology.
- **GCTI.** We include a geometry-aware method with local alignment in the stain vector space (GCTI) [9], which explicitly exploits the underlying geometric structure of color distributions to compensate for scanner- and protocol-induced variations in microscopy images within a unified framework.
- **CycleGAN.** As a deep learning-based baseline, we employ a CycleGAN-style stain transfer model [10]. This method learns a non-linear mapping between source and target stain domains from unpaired data and represents a strong data-driven alternative to hand-crafted normalization schemes.

These baselines jointly cover global-statistics approaches, geometry-aware model-based methods, and modern generative adversarial networks, providing a comprehensive context for evaluating the proposed gradient-domain diffusion normalization.

6.2. Evaluation Metrics

In practice, there is still no widely accepted, unified standard for quantitative evaluation in the stain normalization task. Existing work typically adopts a mixture of structural similarity measures, color-space or gamut-based criteria, and performance gains on downstream tasks (e.g., classification or segmentation). From a clinical perspective, however, the ideal outcome is clear: the normalized image should (i) *preserve* the structural and morphological content of the source image and (ii) *match* the color space and distribution of the target image as closely as possible. Achieving this would allow pathology images acquired at different hospitals, regions, and from heterogeneous devices to be mapped to a common color standard, thereby reducing color-induced variability, improving diagnostic efficiency for pathologists, and ultimately promoting more standardized diagnostic criteria.

To better assess the performance of different algorithms under this perspective, we evaluate: (1) the structural similarity between the source image and the normalized result, and (2) the chromatic alignment between the normalized result and the target image. Concretely, we use the Structural Similarity Index (SSIM) [44] as a morphology-oriented metric between the source I_S and the normalized image $\tilde{I}_S$, and the 1st-Wasserstein distance [45] between the color distributions of $\tilde{I}_S$ and the target I_T . Considering the entanglement of structure and color in the RGB space, all metrics are computed in the CIELab color space. The specific definitions are as follows.

Structural Similarity Index (SSIM). To quantify the preservation of tissue morphology and local texture, we employ the SSIM metric. Let $\mathbf{x}$ and $\mathbf{y}$ denote local window patches extracted from the luminance channels of the source image I_S and the normalized image $\tilde{I}_S$, respectively. The SSIM is defined as:

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \quad (22)$$

where μ_x, μ_y represent the local means, σ_x^2, σ_y^2 denote the local variances, and σ_{xy} is the cross-covariance between the two patches. The constants $C_1 = (k_1L)^2$ and $C_2 = (k_2L)^2$ are included to ensure numerical stability, where L is the dynamic range of pixel values (e.g., $L = 255$ for 8-bit images). The final SSIM score is computed as the mean over all sliding windows in the image domain.

Wasserstein Distance (W-D). To evaluate chromatic alignment, we calculate the 1st-Wasserstein distance (also known as Earth Mover's Distance) between the color distributions of the target image I_T and the result $\tilde{I}_S$. For a specific color channel $c \in \{a^*, b^*\}$, let P_c and Q_c represent the probability

distributions of pixel intensities in I_T and $\tilde{I}_S$, respectively. The Wasserstein distance $W_1(P_c, Q_c)$ is defined as the minimal cost to transform one distribution into the other:

$$W_1(P_c, Q_c) = \inf_{\gamma \in \Pi(P_c, Q_c)} \mathbb{E}_{(u,v) \sim \gamma} [|u - v|], \quad (23)$$

where $\Pi(P_c, Q_c)$ denotes the set of all joint distributions $\gamma(u, v)$ whose marginals are P_c and Q_c . Since we operate on 1D marginal distributions of color channels, Eq. (23) has a closed-form solution expressed via the cumulative distribution functions (CDFs), $F_T(z)$ and $F_R(z)$:

$$W_1(P_c, Q_c) = \int_{-\infty}^{+\infty} |F_T(z) - F_R(z)| dz. \quad (24)$$

A lower W_1 score indicates superior alignment of the statistical color profile, without assuming Gaussianity of the underlying distributions.

6.3. Results and Visual Analysis

We complement the quantitative evaluation with a detailed visual comparison across four representative datasets and a synthetic benchmark. This analysis focuses on both full-size whole-slide views and tiled reconstructions.

Full size baseline test

As illustrated in the first two rows of Figure 4, both Reinhard and SPN achieve faithful chromatic adaptation. This performance is consistent with their explicit statistical matching in the CIELab space, where the a^* and b^* channels isolate chromaticity from luminance. In contrast, the third and fourth rows reveal the limited generalization of CycleGAN on out-of-distribution data: when applied to non-MITOS samples, the model introduces a characteristic “deep red” bias inherent to its training set rather than correctly adapting to the target domain. Furthermore, compared to the global Reinhard method, SPN and GCTI preserve high-frequency morphological details, exhibiting significantly sharper cellular boundaries at high magnification.




Figure 4. Full-size results of different methods on the MITOS-ATYPIA-14 and SYSU-BF-10 datasets. Normalization setup for the MITOS-ATYPIA-14 (cross-scanner: Aperio to Hamamatsu, 10× / 40×) and SYSU-BF-10 (temporal drift: 10-year rescan) datasets. Note that the CycleGAN results are obtained using a model trained only on mitosis dataset.

SCC dataset: full-size vs. tiled normalization.

The SCC dataset contains slides acquired from five different scanners with codes cs2, gt450, nz20, nz210, and p100. We use cs2 as the source device and treat the remaining four devices as distinct targets. Figure 5 compares color normalization on both full-size images and tiled images. The full-size resolution is 2000×1600 , and the tiled setting uniformly partitions each image into a 5×4 grid of non-overlapping patches.







































Source(cs2)	Target1(gt450)		Target2(nz20)		Target3(nz210)		Target4(p1000)	
	Full-size[col_1] (2000*1600)	Tiled-imgs[col_2] (5*4 grid)	Fullsize[rcol_3] (2000*1600)	Tiled-imgs[col_4] (5*4 grid)	Full-size[col_5] (2000*1600)	Tiled-imgs[col_6] (5*4 grid)	Full-size[col_7] (2000*1600)	Tiled-imgs[col_8] (5*4 grid)
								
CycleGan (trained by mitos- dataset)								
Reinahrđ								
GCTI								
SPN								

Figure 5. Full-size and tiled normalization results on the multi-scanner SCC dataset (cs2, gt450, nz20, nz210, p100). Due to limited generalization outside the MITOS-14 training set, CycleGAN produces identical outputs across different target domains. We therefore display the base CycleGAN reconstruction in Column 2, with corresponding high-magnification details shown in col 4, 6, and 8.

As shown in Figure 5, tile-wise normalization via Reinahrđ and GCTI introduces distinct boundary discontinuities arising from inconsistent local statistics. While SPN effectively eliminates these seams through gradient preservation, CycleGAN displays a diminished checkerboard effect relative to the statistical methods. This partial suppression is attributed to the shared generator, which enforces a consistent color transform across tiles. Yet, this same mechanism limits adaptability, causing CycleGAN to hallucinate the MITOS style onto the SCC samples rather than respecting the target domain’s specific appearance.

HEV dataset: Multi-stained pathology images.

The HEV dataset encompasses nine distinct staining configurations, offering the richest color diversity among the datasets analyzed. We utilize standard HE as the source domain, treating the remaining eight protocols as targets. For brevity, and because full-slide comparisons show limited variation, Figure 6 focuses exclusively on the tiled results.




Figure 6. Tiled-image results of different methods on the HEV dataset. The CycleGAN model is still trained on MITOS-ATYPIA-14, leading to nearly identical outputs across the eight HEV targets; we therefore show a single CycleGAN row with selected regions magnified.

The CycleGAN results in Figure 6 mirror those in Figure 5: blocking artifacts are diminished relative to Reinhard and GCTI due to the global constraints of the generator. More significantly, this experiment highlights the robustness limitations of the GCTI framework. GCTI employs a local wedge-finding algorithm that projects pixel intensities onto a subspace defined by SVD, inferring stain vectors via angular percentiles (e.g., 0.1th and 99.9th). This mechanism is susceptible to failure in extreme scenarios: (1) Illumination correction singularities, where near-zero background estimates lead to division-by-zero errors and intensity saturation; and (2) Thresholding instability in uniform regions, where the lack of color variance renders percentile-based thresholds arbitrary. Figure 6 clearly demonstrates the former, where instability in the background estimation step produces extensive artifacts characterized by signal saturation and loss of texture.

Evaluation on Synthetic Image

To rigorously assess algorithmic stability under limiting conditions, we introduce a synthetic image characterized by a smooth background gradient and a central geometric object (Figure 7).




Figure 7. Robustness evaluation on the synthetic image. **Row 1:** Full-resolution normalization. SPN preserves the background gradient faithfully. GCTI exhibits a division-by-zero singularity due to background misclassification, causing saturation. **Rows 2–3:** Tiled reconstruction using a non-overlapping grid. CycleGAN trained by mitosis data Reinhard fails on homogeneous tiles where statistical moments are undefined. GCTI suffers from rank deficiency in zero-gradient regions (pure background); tiles where the algorithm failed to converge are masked in gray.

The first row of Figure 7 presents the full-resolution normalization results. SPN and Reinhard achieve the most faithful reconstruction of the background gradient, demonstrating robust structural consistency across both low-frequency trends and high-frequency edges. In contrast, GCTI exhibits a catastrophic failure mode driven by algorithmic instability: the central large circle is misclassified as background, causing the denominator in the illumination correction step to approach zero. This singularity leads to numerical explosion, resulting in a saturated, near-uniform output in the background region.

Additionally, the magnified regions in the second row reveal that CycleGAN produces spurious ringing artifacts around the circular structures, a phenomenon inconsistent with the behavior of the variational and statistical methods. This artifact stems from the generative network’s reliance on learned upsampling filters (transposed convolutions) which, unlike rigorous interpolation schemes, can hallucinate high-frequency noise to satisfy the discriminator. Such artifacts underscore the critical deficit of deep generative models regarding interpretability: unlike the proposed framework, which is constrained by a physical diffusion process, the GAN introduces opaque, architectural artifacts that compromise diagnostic confidence. The tiled results (Row 3) further expose the stability limits of the baseline methods, particularly regarding well-posedness on disjoint subdomains. First, statistical methods break down on homogeneous patches. Reinhard normalization relies on matching first- and second-order moments; when a target tile consists of a pure background (zero variance), the scaling factors become undefined, rendering the transfer function ill-posed. Similarly, GCTI fails in these

homogeneous regions because the lack of spectral diversity renders the SVD-based stain estimation rank-deficient, leading to numerical breakdown.

Conversely, SPN maintains both global gradient continuity and local geometric fidelity across all tiles. This robustness confirms that the proposed variational framework effectively regularizes the normalization problem, preventing the degenerate behavior observed in purely statistical approaches when local data is sparse.

6.4. Quantitative Results and Discussion

Table 2 presents a comprehensive quantitative comparison of our proposed SPN against three state-of-the-art methods: CycleGAN, Reinhard, and GCTI. We utilize the Structural Similarity Index Measure (SSIM) to evaluate structure preservation ($\uparrow$) and the Wasserstein Distance (W-D) to assess stain distribution matching ($\downarrow$).

Table 2. Quantitative comparison of different stain normalization methods. The best results (max SSIM and min W-D) for each row are highlighted in **bold**.

Dataset	Pair	CycleGAN [†]		Reinhard		GCTI		SPN (ours)	
		SSIM $\uparrow$	W-D $\downarrow$	SSIM $\uparrow$	W-D $\downarrow$	SSIM $\uparrow$	W-D $\downarrow$	SSIM $\uparrow$	W-D $\downarrow$
SYSU-BF	18-n $\rightarrow$ o	0.742	18.572	0.981	0.751	0.970	1.104	0.875	0.689
mitos-14	A06_01 $\rightarrow$ H	0.790	7.274	0.911	2.558	0.813	7.042	0.822	3.027
HEV data	HE $\rightarrow$ longE	0.957	15.392	0.920	4.280	0.736	6.338	0.983	3.524
	HE $\rightarrow$ onlyE	–	12.134	0.729	2.464	0.832	6.485	0.718	1.940
	HE $\rightarrow$ shortE	–	11.902	0.932	5.577	0.832	13.999	0.992	3.705
	HE $\rightarrow$ longHE	–	5.892	0.765	3.370	0.656	10.915	0.814	4.255
	HE $\rightarrow$ onlyH	–	18.802	0.920	2.540	0.634	4.427	0.979	3.672
	HE $\rightarrow$ shortHE	–	15.113	0.873	3.488	0.667	4.096	0.896	3.118
	HE $\rightarrow$ longH	–	4.792	0.875	4.745	0.755	22.539	0.921	2.146
	HE $\rightarrow$ shortH	–	13.199	0.825	2.918	0.694	13.107	0.857	3.218
SCC data	cs2 $\rightarrow$ gt450	0.921	25.624	0.679	0.495	0.757	1.873	0.641	0.360
	cs2 $\rightarrow$ nz20	–	20.885	0.838	0.801	0.898	1.950	0.882	1.715
	cs2 $\rightarrow$ nz210	–	16.042	0.880	0.524	0.917	2.074	0.827	0.581
	cs2 $\rightarrow$ p1000	–	10.979	0.886	1.022	0.798	4.310	0.912	1.579
synthetic_image		0.816	17.433	0.906	4.352	0.831	17.117	0.950	3.706

[†]CycleGAN model trained on the MITOS-ATYPIA-14 dataset. Since SSIM is calculated between result and source image, for the same source test, we omit meaningless comparisons using “–”.

It is worth noting that the classic Reinhard and GCTI algorithms remain strong baselines, particularly on relatively simple and highly structured datasets such as **SYSU-BF-10** and **MITOS-14**. In these cases, Reinhard achieves competitive or even state-of-the-art results (e.g., SSIM of 0.981 on SYSU-BF-10). This is likely because these datasets exhibit unimodal color distributions where simple global statistical alignment (mean and standard deviation matching) is sufficient. However, its performance degrades significantly on complex datasets like HEV (W-D often > 4.0), where stain variations are non-linear and multi-modal.

As previously discussed, CycleGAN utilizes fixed weights trained on MITOS. While it yields consistent structural outputs (omitted entries in Table 2), its inability to adapt to the specific target reference leads to large domain discrepancies, evidenced by high W-D scores (e.g., 25.624 on SCC).

Our method demonstrates robust performance across diverse scenarios, with its advantages becoming most pronounced on the *synthetic_image*. This dataset serves as a rigorous benchmark for “general stain normalization” as it provides paired ground truth. On this set, SPN achieves the best results in both metrics (SSIM: 0.950, W-D: 3.706). This confirms that our gradient-guided diffusion mechanism effectively learns the intrinsic mapping between domains without overfitting to specific dataset artifacts, demonstrating superior generalization compared to baselines.

7. Conclusions

We presented **Screened Poisson Normalization (SPN)**, a physics-informed framework that reformulates stain normalization as a screened Poisson reconstruction problem. By modeling the staining process as a reaction–diffusion system, SPN explicitly bridges statistical moment matching with physically interpretable gradient-domain regularization.

Our theoretical and experimental analysis yields three key conclusions:

1. **Structural Fidelity and Scalability:** Operating in the gradient domain allows SPN to preserve diagnostic morphological details (e.g., nuclear boundaries) often degraded by pure statistical alignment. Furthermore, the exponential decay of the screened Green’s function enables rigorous non-overlapping domain decomposition, permitting seamless gigapixel processing without tiling artifacts.
2. **Determinism vs. Data-Dependence:** In contrast to deep learning approaches (e.g., CycleGAN) which are prone to domain bias and generative hallucinations on out-of-distribution data, SPN is training-free and deterministic. This ensures that all output structures are causally linked to the input, providing the explainability and reproducibility essential for clinical diagnostics.
3. **Robustness to Heterogeneity:** SPN demonstrates consistent performance across diverse scanners, staining protocols, and archived slides. The global anchoring (screening) term effectively stabilizes the solution in tissue-sparse or background-dominated regions, overcoming the numerical brittleness observed in local geometry-based methods.

In summary, SPN offers a transparent, mathematically well-posed alternative to "black-box" style transfer, establishing a stable foundation for computational pathology in heterogeneous clinical settings. Future work will focus on parallel optimization for high-throughput deployment and extending the gradient-domain formulation to multi-stain separation and volumetric reconstruction.

Author Contributions: Conceptualization, J.X.; Methodology, J.X.; Software, J.X.; Formal analysis, J.X.; Validation, J.X.; Visualization, J.X.; Writing—original draft, J.X.; Writing—review and editing, J.X., H.N. and Q.W.; Investigation, H.N. and Q.W.; Resources, H.N.; Data curation, H.N. and Y.J.; Supervision, J.X. and Q.W.; Project administration, J.X.; Funding acquisition, J.X. ,Q.W., H.N. All authors have read and agreed to the published version of the manuscript.

Funding: Junlong Xing was partially supported by the National Key R&D Program of China (No.2020YFA0712500, No.2021YFA1001300),the Research and Development Project of Pazhou Lab (Huangpu) (Grant Number 2023K0601).Qiru Wang was partially supported by the National Natural Science Foundation of China (No. 12471176) and the Guangdong Basic and Applied Basic Research Foundation (No.2025A1515012221). Hengli Niwas partially supported by the Young Scientists Fund of the National Natural Science Foundation of China (82303584) and the Guangdong Provincial Key Laboratory of Mathematical and Neural Dynamical Systems (2024B1212010004).

Institutional Review Board Statement: The publicly available datasets used in this study are cited in the References. For the curated clinical WSI collection from Sun Yat-sen University mentioned in the Data Availability Statement, institutional approval and an appropriate data-use agreement were obtained; details can be provided by the corresponding author upon reasonable request.

Informed Consent Statement: For the curated clinical WSI collection, informed-consent requirements were handled according to the approving institutional policies and local regulations (i.e., consent was obtained or waived by the Institutional Review Board as appropriate).

Data Availability Statement: The data supporting the findings of this study comprise three publicly available whole-slide histopathology datasets and one locally collected clinical dataset. The public datasets used in our experiments—MITOS-ATYPIA-14 [35], the Multi-Scanner Squamous Cell Carcinoma (SCC) dataset [36], and the HE-Staining Variation (HEV) dataset [37]—are available from their respective providers under the original terms of use and licensing described in the corresponding references. In addition, we curated a clinical WSI collection from Sun Yat-sen Memorial Hospital, Sun Yat-sen University (SYSU-BF-10; referred to as SYSU-BF-10 in the main text), consisting of 25 long-term archived breast carcinoma slides (archived for over ten years and rescanned to

study temporal color degradation). Due to ethics, privacy, and institutional governance requirements for clinical data, SYSU-BF-10 is not publicly available at the time of publication. We will release the SYSU-BF-10 dataset after the subsequent ethics and privacy procedures have been reviewed. Until completion of these procedures, access may be considered on reasonable request to the corresponding author, subject to institutional approval and an appropriate data-use agreement.

Acknowledgments: The authors thank the dataset providers and the open-source community for making resources available.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Xu, C., Sun, Y., Zhang, Y., Liu, T., Wang, X., Hu, D., Huang, S., Li, J., Zhang, F., & Li, G. (2025). Stain Normalization of Histopathological Images Based on Deep Learning: A Review. *Diagnostics*, **15**(8), 1032.
2. Hetz, M.J., Bucher, T.C., & Brinker, T.J. (2024). Multi-domain stain normalization for digital pathology: A cycle-consistent adversarial network for whole slide images. *Medical Image Analysis*, **94**, 103149.
3. Wang, C., Li, S., Ke, J., Zhang, C., & Shen, Y. (2024). RandStainNA++: Enhance random stain augmentation and normalization through foreground and background differentiation. *IEEE Journal of Biomedical and Health Informatics*, **28**, 3660–3671.
4. Baykal Kablan, E. (2023). Regional realness-aware generative adversarial networks for stain normalization. *Neural Computation and Applications*, **35**, 17915–17927.
5. Li, Y., Zhou, H., Liu, N., & Shen, Y. (2023). Stain normalization and augmentation in frequency space for histology analysis. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2031–2035.
6. Reinhard, E., Ashikhmin, M., Gooch, B., & Shirley, P. (2001). Color transfer between images. *IEEE Computer Graphics and Applications*, **21**(5), 34–41.
7. Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., & Thomas, N.E. (2009). A method for normalizing histology slides for quantitative analysis. In *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 1107–1110.
8. Vahadane, A., et al. (2016). Structure-preserving color normalization and sparse stain separation. *IEEE Transactions on Medical Imaging*, **35**(8), 1962–1971.
9. Gupta, A., Duggal, R., Gehlot, S., Gupta, R., Mangal, A., Kumar, L., Thakkar, N., & Satpathy, D. (2020). GCTI-SN: Geometry-inspired chemical and tissue invariant stain normalization of microscopic medical images. *Medical Image Analysis*, **65**, 101788.
10. Runz, M., Rusche, D., Schmidt, S., et al. (2021). Normalization of H&E-stained histological images using cycle consistent generative adversarial networks. *Diagnostic Pathology*, **16**, 71.
11. Kang, H., Luo, D., Feng, W., et al. (2021). StainNet: A fast and robust stain normalization network. *Frontiers in Medicine*, **8**, 746307.
12. Shaban, M.T., Baur, C., Navab, N., & Albarqouni, S. (2019). Staingan: Stain style transfer for digital histological images. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 953–956.
13. BenTaieb, A., & Hamarneh, G. (2017). Adversarial stain transfer for histopathology image analysis. *IEEE Transactions on Medical Imaging*, **37**(3), 792–802.
14. Zarella, M.D., Yeoh, C., Breen, D.E., & Garcia, F.U. (2017). An alternative reference space for H&E color normalization. *PLoS One*, **12**(3), e0174489.
15. Magee, D., Treanor, D., Crellin, D., Shires, M., Smith, K., Mohee, K., & Quirke, P. (2009). Colour normalisation in digital histopathology images. *Proceedings of Optical Tissue Image Analysis in Microscopy, Histopathology and Endoscopy (MICCAI Workshop)*, 100–111.
16. Alsubaie, N., Trahearn, N., Raza, S.E.A., Snead, D., & Rajpoot, N.M. (2017). Stain deconvolution using statistical analysis of multi-resolution stain colour representation. *PLoS One*, **12**(1), e0169875.
17. Roy, S., Lal, S., & Kini, J.R. (2019). Novel color normalization method for Hematoxylin & Eosin stained histopathology images. *IEEE Access*, **7**, 28982–28998.
18. Vijh, S., Saraswat, M., & Kumar, S. (2021). A new complete color normalization method for H&E stained histopathological images. *Applied Intelligence*, **51**, 1–14.
19. Zheng, Y., Jiang, Z., Zhang, H., Xie, F., Shi, J., & Xue, C. (2019). Adaptive color deconvolution for histological WSI normalization. *Computer Methods and Programs in Biomedicine*, **170**, 107–120.

20. Nadeem, S., Hollmann, T., & Tannenbaum, A. (2020). Multimarginal Wasserstein barycenter for stain normalization and augmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 362–371.
21. Shrivastava, A., Adorno, W., Sharma, Y., et al. (2021). Self-attentive adversarial stain normalization. In *Pattern Recognition. ICPR Workshops and Challenges*, 120–140. Springer.
22. Nishar, H., Chavanke, N., & Singhal, N. (2020). Histopathological stain transfer using style transfer network with adversarial loss. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 330–340.
23. Liang, H., Plataniotis, K.N., & Li, X. (2020). Stain style transfer of histopathology images via structure-preserved generative learning. In *Machine Learning for Medical Image Reconstruction*, 153–162. Springer.
24. Kausar, T., Kausar, A., Ashraf, M.A., et al. (2021). SA-GAN: Stain acclimation generative adversarial network for histopathology image analysis. *Applied Sciences*, **12**, 288.
25. Cho, H., Lim, S., Choi, G., & Min, H. (2017). Neural stain-style transfer learning using GAN for histopathological images. *arXiv preprint arXiv:1710.08543*.
26. Lee, C.C., Kuo, P.T.P., & Peng, C.H. (2022). H&E stain normalization using U-net. In *2022 IEEE 22nd International Conference on Bioinformatics and Bioengineering (BIBE)*, 29–32.
27. Menning, J.D.M., Wallmersperger, T., Meinhardt, M., & Ehrenhofer, A. (2022). Modeling and simulation of diffusion and reaction processes during the staining of tissue sections on slides. *Histochemistry and Cell Biology*, **158**, 137–148.
28. Gurina, T.S., & Simms, L. (2023). Histology, Staining. In *StatPearls [Internet]*. StatPearls Publishing.
29. Suvarna, K.S., Layton, C., & Bancroft, J.D. (Eds.). (2019). *Bancroft's Theory and Practice of Histological Techniques*, 8th ed. Elsevier.
30. Mayerhöfer, T.G., Pahlow, S., & Popp, J. (2020). The Bouguer-Beer-Lambert law: Shining light on the obscure. *ChemPhysChem*, **21**(18), 2029–2046.
31. Murray, J.D. (2002). *Mathematical Biology I: An Introduction*. 3rd ed. Springer: New York, NY, USA.
32. Sheehan, D.C., & Hrapchak, B.B. (2012). *Theory and Practice of Histotechnology*. 3rd ed. Mosby.
33. Hetz, M.J., Bucher, T.C., & Brinker, T.J. (2024). Multi-domain stain normalization for digital pathology: A cycle-consistent adversarial network for whole slide images. *Medical Image Analysis*, **94**, 103149.
34. Landini, G. (2021). Colour deconvolution with ImageJ and Fiji: a bibliography. *Bioinformatics*, **37**, 4181–4182.
35. Roux, L., Racoceanu, D., Loménie, N., et al. (2013). Mitosis detection in breast cancer histological images: An ICPR 2012 contest. *Journal of Pathology Informatics*, **4**, 8.
36. Wilm, F., Fragoso, M., Bertram, C.A., Stathonikos, N., Öttl, M., Qiu, J., Klopffleisch, R., Maier, A., Breininger, K., & Aubreville, M. (2023). Multi-scanner canine cutaneous squamous cell carcinoma histopathology dataset. *arXiv preprint arXiv:2301.04423*.
37. Runz, M., & Weis, C.-A. (2021). Normalization of HE-Stained Histological Images using Cycle Consistent Generative Adversarial Networks [Dataset]. *heiDATA*, V1. <https://doi.org/10.11588/DATA/8LKEZF>
38. Pérez, P., Gangnet, M., & Blake, A. (2023). Poisson image editing. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2* (pp. 577–582).
39. Habring, A., & Holler, M. (2024). Neural-network-based regularization methods for inverse problems in imaging. *GAMM-Mitteilungen*, **47**(4), e202470004.
40. Hansen, P.C. (2010). *Discrete Inverse Problems: Insight and Algorithms*. SIAM.
41. Benning, M., & Burger, M. (2018). Modern regularization methods for inverse problems. *Acta Numerica*, **27**, 1–111.
42. Fortunato, D., Townsend, A., & Trefethen, L.N. (2020). A Spectral Method for Elliptic Problems in a Rectangle. *IMA Journal of Numerical Analysis*, **40**, 2250–2280.
43. Klawonn, A., Lanser, M., & Weber, J. (2024). Machine learning and domain decomposition methods - a survey. *Computational Science and Engineering*, **1**(1), 2.
44. Bakurov, I., Buzzelli, M., Schettini, R., Bianco, S., & Raimondo, F. (2022). Structural similarity index (SSIM) revisited: A data-driven approach. *Expert Systems with Applications*, **189**, 116087.
45. Nguyen, K., & Ho, N. (2022). Revisiting sliced Wasserstein on images: From vectorization to convolution. *Advances in Neural Information Processing Systems*, **35**, 17788–17801.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.