GED: GAUSSIAN EXPECTED DEFORMATION FOR MASK-FREE DAMAGED-SECTION REGISTRATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Folds and tears in serial-section electron microscopy create abrupt motion changes between neighboring tissue regions. We present Gaussian Expected Deformation (GED), a pairwise registration method that learns the spatial influence of local displacement fields without damage masks. Anisotropic Gaussian supports can extend along damage boundaries and narrow across them, reducing the mixing of incompatible motions. Sequential RANSAC extracts multiple motion groups from feature correspondences to initialize the fields. Per-pair optimization then refines their parameters and supports through image-similarity and correspondence objectives. On the synthetic benchmark, GED reduces valid endpoint error from 2.09 to 0.215 pixels relative to GPO, a reduction of 89.7%. Near-damage error falls from 6.66 to 1.10 pixels, a reduction of 83.4%. Controlled support ablations and learned-weight visualizations link anisotropy to lower near-damage error and cross-boundary mixing.

1 INTRODUCTION

Connectomic reconstruction requires accurately aligning serial-section electron microscopy (ssEM) images into a coherent volume (Saalfeld et al., 2010; Popovych et al., 2024). Mechanical stresses during sectioning can produce folds, tears, and local tissue displacements that disrupt spatial correspondence between adjacent sections (Scheffer et al., 2013). If left uncorrected, the resulting misalignments accumulate through the image stack and compromise neuron tracing and connectivity analysis (Lee et al., 2019; Popovych et al., 2024).

Registering a damaged section requires estimating its spatial mapping to a reference section. Within intact tissue, deformation varies smoothly, but it can change abruptly across folds and tears (Mitchell et al., 2019; Popovych et al., 2020; Xin et al., 2022). This creates a trade-off between propagating motion within tissue and limiting its influence across damage boundaries. Estimating the local motions is also difficult because compression, overlap, and disrupted structures produce sparse or incorrect correspondences (Scheffer et al., 2013; Xin et al., 2022). Figure 1 illustrates a synthetic fold. Global RANSAC captures the dominant motion, while GED recovers the opposing motions.

General methods (Saalfeld et al., 2012; Khairy et al., 2018) estimate dense deformations with spatially homogeneous smoothness assumptions, which can blend incompatible motions across tears and between independently displaced tissue regions (Mitchell et al., 2019; Xin et al., 2022). Damage-aware approaches use detected folds, tears, or tissue fragments to apply region-specific transformations or adapt regularization near damage boundaries (Mitchell et al., 2019; Popovych et al., 2020; Xin et al., 2022). They can keep incompatible motions apart but require damage masks, auxiliary defect detectors, or predefined tissue partitions.

We propose Gaussian Expected Deformation (GED), a mask-free deformation representation for damaged sections. GED learns anisotropic Gaussian supports to propagate motion within tissue and limit mixing across damage boundaries. Normalized support weights combine local displacement fields into a dense correction field. Sequential RANSAC (Fischler & Bolles, 1981) provides the initialization by retaining multiple motion-consistent groups of correspondences. Per-pair optimization refines the local fields and their supports using image similarity and correspondence constraints.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

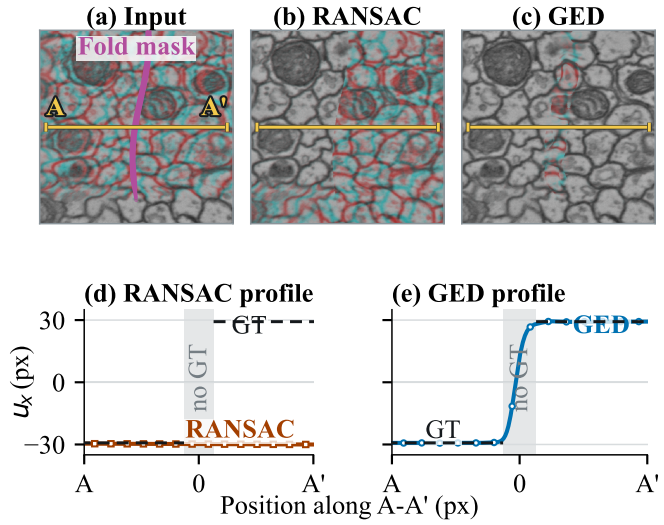

Figure 1: Global RANSAC and GED at a synthetic fold. (a–c) Input, RANSAC, and GED overlays show the reference in red, damaged image in cyan, and fold mask in magenta. (d,e) Horizontal correction u_x along reference-image line A–A', compared with ground truth. Gray regions have no valid ground-truth correspondences.

The contributions are summarized below.

- A mask-free deformation representation that controls the spatial influence of local displacement fields. Learned anisotropic Gaussian supports limit cross-boundary mixing to approximate abrupt motion transitions near damage.
- A motion initialization that preserves the distinct motions present in noisy correspondences, combining sequential RANSAC decomposition with correspondence-guided Gaussian anchors.
- Geometric evaluation on synthetic damage and real sections using deformation ground truth and manual landmarks. Support ablations examine the effect of anisotropy near damage boundaries.

2 Related Work

Serial-section EM registration. Elastic spring-mesh models construct smooth deformation fields from local correspondences (Saalfeld et al., 2012; Cardona et al., 2012). Learning-based methods improve axial continuity through learned features, correlation volumes, and deformation integration (Yoo et al., 2017; Zhou et al., 2019; Liu et al., 2023; Xin et al., 2023; Zhang et al., 2025). Damage-aware methods adapt registration to folds, tears, or tissue fragments. They register segmented regions separately (Scheffer et al., 2013; Popovych et al., 2020) or relax smoothness near detected tears (Mitchell et al., 2019; Popovych et al., 2024). Expected Affine combines local transformations using damage-mask-aware path distances (Xin et al., 2022). Other approaches jointly correct deformation and restore missing content (Lv et al., 2024; Huang et al., 2020).

Gaussian primitives. Gaussian primitives are central to 3D Gaussian splatting for scene reconstruction and novel-view synthesis (Kerbl et al., 2023; Lu et al., 2024). Recent work has adapted them to image registration. GPO (Tian et al., 2025) uses Gaussian control nodes to propagate displacements from salient retinal vessels into texture-poor regions. GaussianDIR (Li et al., 2025) employs multiscale Gaussian primitives to model smooth anatomical deformations in volumetric CT and MRI. GED uses learned anisotropic Gaussian supports to approximate abrupt motion transitions in damaged ssEM sections.

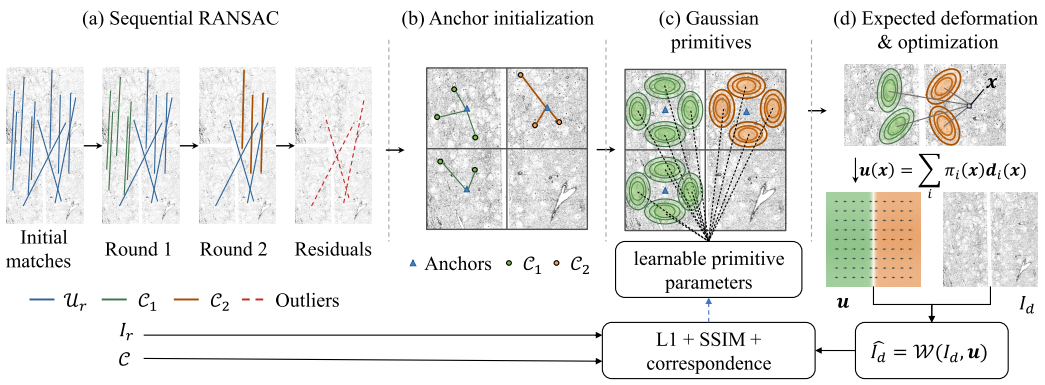

Figure 2: Overview of GED. (a) Sequential RANSAC groups correspondences by motion. (b) Accepted correspondences initialize grid anchors. (c) Each anchor generates anisotropic Gaussian primitives. (d) Weighted local displacements form the correction field $\mathbf{u}(\mathbf{x})$ used to warp I_d.

3 Method

3.1 Problem Formulation

Let $I_r, I_d : \Omega \to [0,1]$ denote the reference and damaged images over the image domain $\Omega \subset \mathbb{R}^2$. GED estimates a dense correction field $\mathbf{u} : \Omega \to \mathbb{R}^2$ on the reference domain. For each reference coordinate $\mathbf{x}$, the corresponding damaged coordinate is $\Psi_{\mathbf{u}}(\mathbf{x}) = \mathbf{x} - \mathbf{u}(\mathbf{x})$, so $\mathbf{u}(\mathbf{x})$ points from the damaged location to the reference location, and the registered image is $\hat{I}_d(\mathbf{x}) = I_d(\Psi_{\mathbf{u}}(\mathbf{x}))$. Figure 2 summarizes the pipeline.

3.2 Matching-Driven Motion Decomposition

We extract keypoints and descriptors with ALIKED (Zhao et al., 2023) and establish correspondences with LightGlue (Lindenberger et al., 2023), which rejects matches below a confidence threshold. A correspondence is $(\mathbf{p}_n, \mathbf{q}_n)$, where $\mathbf{p}_n$ and $\mathbf{q}_n$ are coordinates in the reference and damaged images. A damaged pair still contains both incorrect matches and correct matches from different tissue motions, so a single transformation retains only the dominant motion and treats the others as outliers.

We apply RANSAC (Fischler & Bolles, 1981) repeatedly to separate motion-consistent groups. Let $\mathcal{U}_r$ be the unassigned correspondences before round r. We fit a damaged-to-reference similarity transformation T_r and collect its consensus set $\mathcal{C}_r = \{n \in \mathcal{U}_r : \|\mathbf{p}_n - T_r(\mathbf{q}_n)\|_2 \le \tau_m\}$. We remove $\mathcal{C}_r$ from $\mathcal{U}_r$ before the next round. Rounds stop when no sufficiently supported consensus remains or the hypothesis limit is reached. Unassigned matches are discarded. Appendix B gives the thresholds and limits.

3.3 Correspondence-Guided Gaussian Anchors

We lay a regular grid over the reference image at the optimization resolution and assign each accepted correspondence to the cell containing $\mathbf{p}_n$. An anchor is placed at the center of each cell containing at least one correspondence. The accepted correspondences initialize the local field parameters of K primitives at distinct offsets around each anchor. Appendix B gives the initialization settings.

3.4 Adaptive Gaussian Deformation Field

Each active anchor generates K primitives. Primitive i has a center $\boldsymbol{\mu}_i$, two positive scales $s_{i,x}, s_{i,y}$, a support orientation ψ_i, and an opacity α_i. Its covariance is $\boldsymbol{\Sigma}_i = R(\psi_i)\,\mathrm{diag}(s_{i,x}^2, s_{i,y}^2)R(\psi_i)^\mathsf{T}$, where R denotes the two-dimensional rotation matrix. Its unnormalized weight is

$$w_i(\mathbf{x}) = \alpha_i \exp\big[-\tfrac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_i)^\mathsf{T}\boldsymbol{\Sigma}_i^{-1}(\mathbf{x}-\boldsymbol{\mu}_i)\big]. \tag{1}$$

The two scales and support orientation allow influence along and across a boundary to vary independently. Each primitive carries a local displacement field $\mathbf{d}_i(\mathbf{x})$ with a center displacement $\mathbf{t}_i$ and a linear residual controlled by a scalar ϕ_i. The field parameters are optimized jointly with the Gaussian support. Shared MLPs decode the support scales, orientation, and opacity from anchor features. Appendix B specifies the decoder and primitive settings.

3.5 Expected Deformation Aggregation and Optimization

Weights are evaluated per 16×16-pixel tile. Let $\mathcal{N}(\mathbf{x})$ be the candidate primitives of the tile containing $\mathbf{x}$. The correction field is

$$\pi_i(\mathbf{x}) = \frac{w_i(\mathbf{x})}{\max\big(\sum_{l\in\mathcal{N}(\mathbf{x})} w_l(\mathbf{x}),\, \epsilon\big)}, \qquad \mathbf{u}(\mathbf{x}) = \sum_{i\in\mathcal{N}(\mathbf{x})} \pi_i(\mathbf{x})\,\mathbf{d}_i(\mathbf{x}), \tag{2}$$

with $\epsilon = 10^{-8}$ guarding against division by zero. The weights are nonnegative and sum to one wherever the total weight exceeds ϵ, forming a spatially adaptive partition of unity. This aggregation follows the weighted expectation of Expected Affine (Xin et al., 2022). Its damage-mask-aware path weights are replaced by learned Gaussian weights.

Transition width. Consider two Gaussian primitives centered at $s = \pm a$ along a boundary normal. They have equal opacity and normal scale σ_n. The width over which the normalized mixture weight rises from 10% to 90% is

$$W_{10\text{–}90} = \frac{\sigma_n^2 \ln 9}{a}. \tag{3}$$

For these two primitives, reducing the normal scale narrows the transition at fixed center separation. The tangential scale can remain large to preserve coverage along the boundary. Appendix D.2 gives the derivation. Section 4.3 examines the learned supports through ablations and cross-boundary weight measurements.

Objective. All trainable variables are jointly optimized for each pair. Let $\mathcal{W}(I_d, \mathbf{u}) = I_d \circ \Psi_\mathbf{u}$ denote differentiable warping and $\mathcal{C} = \bigcup_r \mathcal{C}_r$ the accepted correspondences. We minimize

$$\mathcal{L} = \frac{\lambda_1}{|\Omega|}\,\|\mathcal{W}(I_d,\mathbf{u}) - I_r\|_1 + \lambda_s \mathcal{L}_{\mathrm{SSIM}}\left(\mathcal{W}(I_d,\mathbf{u}), I_r\right) + \frac{\lambda_m}{|\mathcal{C}|}\sum_{n\in\mathcal{C}} \|\Psi_\mathbf{u}(\mathbf{p}_n) - \mathbf{q}_n\|_2^2, \tag{4}$$

where the L1 term is averaged over pixels and $\mathcal{L}_{\mathrm{SSIM}}$ is the structural similarity loss (Wang et al., 2004). The correspondence term aligns the sampled coordinates with the accepted matches.

4 Experiments

4.1 Experimental Setup

Datasets. We evaluate on synthetic pairs with dense deformation ground truth and on five real pairs from the Damaged Section benchmark (Xin et al., 2022). Synthetic pairs are generated from 125 CREMI sections[1] and 100 AC4 sections (Kasthuri et al., 2015) with instance segmentations at 1024×1024 pixels. We synthesize folds and tears through piecewise tissue displacements across damage boundaries. Thin-plate-spline deformations provide smooth controls (Bookstein, 1989). Each dataset contributes 50 pairs for each of three fold and tear severities, plus 50 smooth controls, yielding 700 pairs in total. Appendix A summarizes the synthesis procedure. The real pairs contain folds, a tear, and smooth nonlinear deformation, with image sizes from 1000×1000 to 6400×3776 pixels. Appendix E shows all five cases.

[1] https://cremi.org

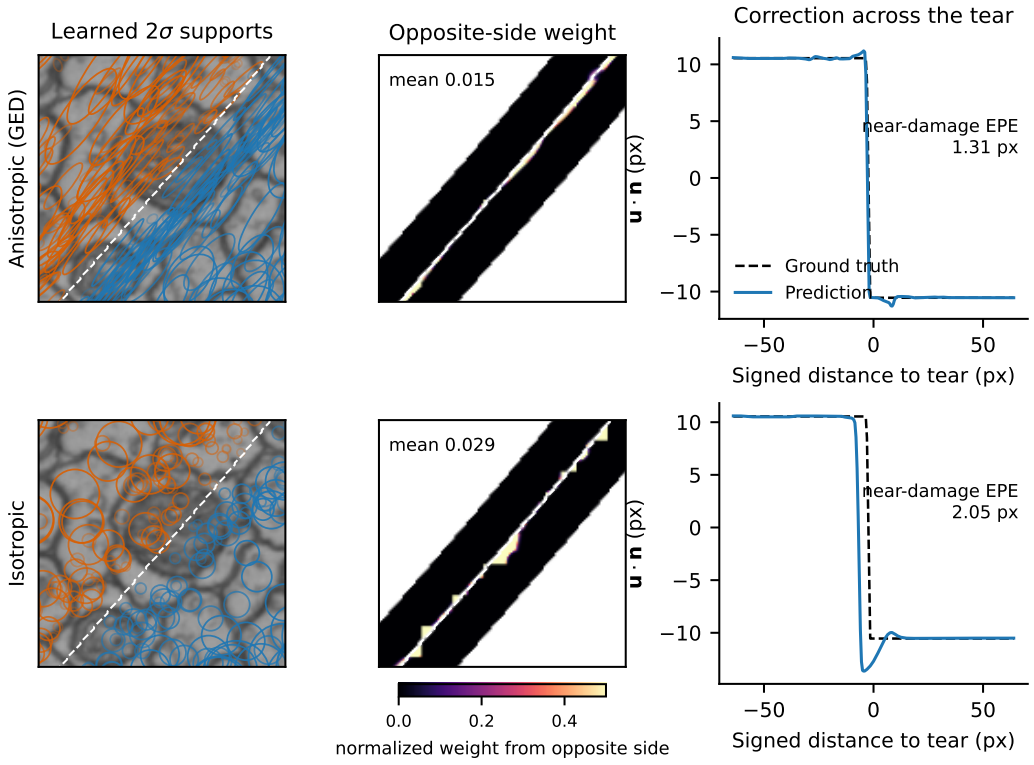

Figure 3: Learned supports for a medium-severity synthetic tear from CREMI. Anisotropic GED is above the isotropic variant. Columns show support ellipses, opposite-side weights, and displacement profiles.

Design and evaluation split. Ablations and optimization-step selection use 40 medium-severity synthetic pairs, with ten pairs per dataset and damage type. The main comparison uses the remaining 660 pairs. The 50-step setting is selected using design-set EPE.

Baselines. We compare against global affine RANSAC, Expected Affine (EA) (Xin et al., 2022), Elastic (Saalfeld et al., 2012; Cardona et al., 2012), SEAMLeSS (Popovych et al., 2024), GPO (Tian et al., 2025), and a 2D adaptation of GaussianDIR* (Li et al., 2025). We use the official implementations with the settings and adaptations described in Appendix B. EA receives the generated damage masks for synthetic pairs and the benchmark masks for real pairs.

Implementation. All methods register images at a 512-pixel short edge. Their backward maps are upsampled to the original resolution for evaluation. GED uses a 24×24 anchor grid, $K = 4$, and 50 Adam steps (Kingma & Ba, 2015) with $(\lambda_1, \lambda_s, \lambda_m) = (0.5, 0.2, 0.005)$.

Metrics. In method comparison tables, bold marks the best value and underlining the second best. We report global normalized cross-correlation (NCC) (Lewis, 1995). For synthetic pairs, valid endpoint error (EPE) is the Euclidean distance between predicted and ground-truth sampling coordinates in original-resolution pixels (Baker et al., 2011). It is computed densely over pixels with a valid ground-truth correspondence. Fold overlaps, tear gaps, and out-of-image pixels are excluded. Near-damage EPE uses valid pixels within a 32-pixel dilation of the damage mask and is averaged over the 495 pairs containing such pixels. We compute mIoU over foreground instances after nearest-neighbor label warping (Everingham et al., 2010; Long et al., 2015).

For support analysis, we define cross-boundary leakage using the ground-truth damage boundary. At each valid near-damage pixel, we sum the normalized weights of primitives centered on the opposite side. We report the mean of these sums. Lower values indicate less cross-boundary mixing.

Table 1: Synthetic results on the primary evaluation set. Values are means; ± denotes s.d. EPE is in original-resolution pixels. Near-damage EPE uses pairs with valid near-damage ground truth.

Method	Valid EPE ↓	Near EPE ↓	NCC ↑	mIoU ↑
RANSAC	13.47±12.61	22.59	.634±.161	.591±.163
EA	3.91±3.19	6.98	.794±.086	.661±.165
Elastic	13.37±10.82	23.08	.535±.157	.490±.158
SEAMLeSS	17.93±19.48	17.94	.583±.340	.498±.313
GPO	2.09±4.66	6.66	.908±.108	.909±.097
GaussianDIR*	7.22±12.50	18.43	.954±.051	.831±.220
GED	**.215±.129**	**1.10**	**.969±.028**	**.966±.048**

*Our 2D adaptation of the official implementation, originally designed for 3D volumes.

Table 2: Mean landmark error on the real section pairs in original-resolution pixels.

Pair	RANSAC	EA	Elastic	SEAMLeSS	GPO	GaussianDIR*	**GED**
DOLW7	101.295	4.053	—	92.218	89.498	86.624	**2.309**
SS2L30	3.768	7.488	11.539	4.523	1.186	**1.063**	1.168
SW13L63	19.683	11.063	18.217	40.515	3.703	64.221	**2.110**
ZW35L414	11.028	6.815	20.484	50.802	1.010	4.668	**0.955**
ZW7L18	10.056	6.264	12.563	4.886	0.869	**0.747**	0.817
Mean	29.166	7.136	—	38.589	19.253	31.464	**1.472**

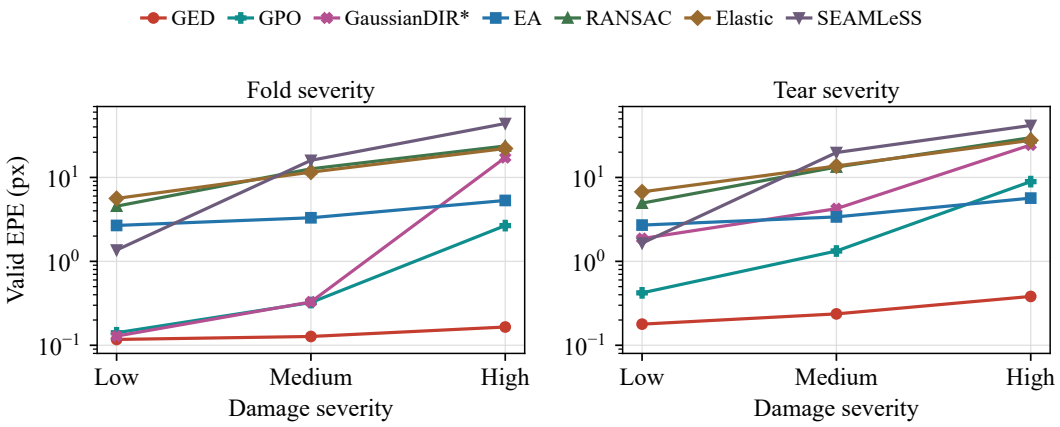

Figure 4: Mean valid EPE across damage severity on CREMI and AC4 after excluding the design pairs. The left panel shows folds and the right panel shows tears. Vertical axes are logarithmic.

For the real pairs, we manually annotate 16 corresponding landmarks per pair. We report the mean Euclidean distance between the predicted and annotated damaged-image coordinates in original-resolution pixels. Appendix E describes the annotation and evaluation procedure.

4.2 Comparison with Existing Methods

Synthetic registration. Table 1 summarizes the primary evaluation results. GED achieves the highest mean NCC and mIoU and the lowest mean valid and near-damage EPE. Relative to GPO, valid EPE drops by 89.7% and near-damage EPE by 83.4%.

Damage type and severity. GED has the lowest mean valid EPE at each fold and tear severity in Figure 4. The gap is wider at high severity. On smooth controls, the Gaussian methods have similarly low mean valid EPE.

Table 3: Near-damage EPE by damage type on primary pairs with valid near-damage ground truth. Errors are in original-resolution pixels.

Damage	RANSAC	EA	Elastic	SEAMLeSS	GPO	GaussianDIR*	**GED**
Fold	17.045	8.780	16.635	10.014	1.880	<u>1.633</u>	**.693**
Tear	26.853	<u>5.606</u>	28.027	24.023	10.338	31.327	**1.420**

Table 4: Results on CREMI and AC4 primary evaluation pairs. EPE is in original-resolution pixels.

	CREMI			AC4		
Method	NCC ↑	Valid EPE ↓	Near EPE ↓	NCC ↑	Valid EPE ↓	Near EPE ↓
RANSAC	.608	13.409	22.804	.659	13.538	22.377
EA	.814	3.790	7.930	.773	4.024	<u>6.019</u>
Elastic	.522	13.559	23.291	.549	13.171	22.862
SEAMLeSS	.567	17.735	17.745	.599	18.127	18.136
GPO	.898	<u>1.780</u>	<u>6.671</u>	.918	<u>2.398</u>	6.657
GaussianDIR*	<u>.953</u>	9.463	22.098	<u>.955</u>	4.968	14.686
GED	**.967**	**.217**	**1.280**	**.971**	**.212**	**.924**

GED also has the lowest mean near-damage EPE for both damage types in Table 3. On tears, GED gives 1.420 pixels, compared with 10.338 for GPO and 31.327 for GaussianDIR*.

Results by dataset. GED has the lowest mean valid and near-damage EPE on both CREMI and AC4, as shown in Table 4.

Real damaged sections. GED has the lowest overall mean landmark error across the five real pairs in Table 2. DOLW7 contains a fold several hundred pixels wide. GED has a mean landmark error of 2.309 pixels on this pair, compared with 89.498 for GPO and 86.624 for GaussianDIR*. Figure 5 compares tissue structures near a tear in SW13L63 and a fold in ZW7L18, alongside synthetic examples. Appendix E provides NCC results and an overview of all five pairs.

4.3 Ablation Studies

Table 5 compares support geometry, field representation, and matching on the design set. All configurations use the same raw correspondence cache. In ablation tables, underlining marks the adopted setting and bold marks the best value within each block.

Adaptive support. Replacing the two decoded scales with their mean gives the isotropic variant. It increases near-damage EPE by 27%. Valid EPE and cross-boundary leakage also increase. Appendix D gives the boundary-side assignment and sampling details.

Figure 3 shows learned supports at a medium synthetic tear from CREMI. The annotations report pair-wide near-damage EPE and cross-boundary leakage. Profiles use optimization-grid pixels, and EPE uses original-resolution pixels. Ellipse colors indicate the side of the tear on which primitive centers lie; their opacity follows α_i. In this example, anisotropic primitives elongate along the tear while remaining narrow across it. The middle column shows lower weights from opposite-side primitives near the boundary. The resulting correction curve places the transition close to the ground-truth boundary. With isotropic supports, the transition shifts by several pixels.

Field representation. With the same accepted correspondences, replacing TPS by the Gaussian field reduces near-damage EPE from 6.669 to 0.988 pixels.

Motion initialization. Clustering followed by RANSAC and sequential RANSAC both have much lower near-damage EPE than single RANSAC. Their errors are close. Sequential RANSAC determines the number of hypotheses from geometric consensus.

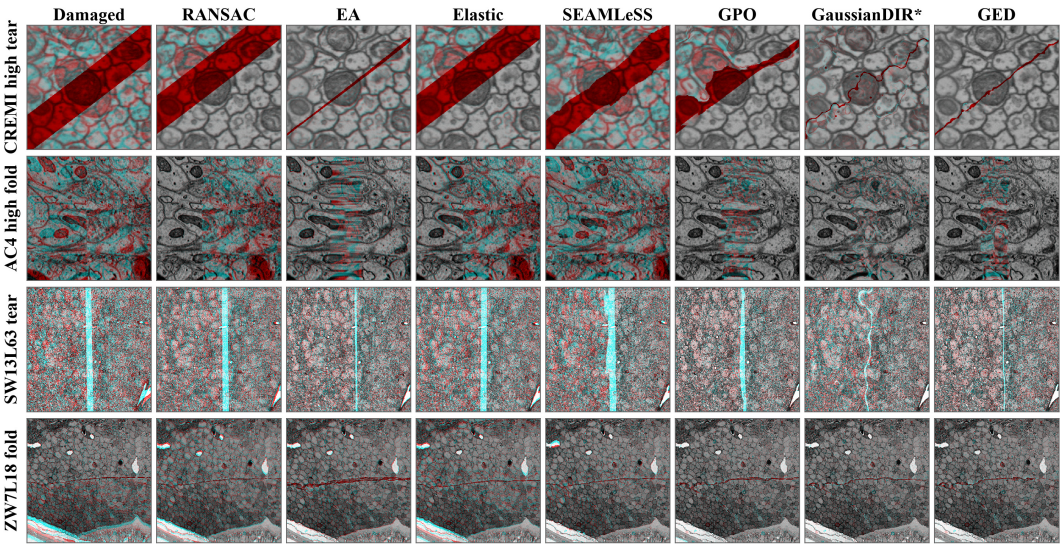

Figure 5: Local registration around folds and tears. Overlays show the reference in red and the damaged or registered image in cyan, with aligned structures appearing gray.

Table 5: Support, field, and matching ablations on the design set.

Factor	Setting	NCC ↑	Valid EPE ↓	Near EPE ↓	Leakage ↓
Support	Anisotropic	**.973**	**.173**	**.987**	**.019**
	Isotropic	.969	.199	1.252	.023
Field	TPS	.889	.727	6.669	–
	Gaussian GED	**.973**	**.173**	**.988**	–
Matching	Single RANSAC	.792	7.169	21.733	–
	Cluster + RANSAC	.972	.186	1.021	–
	Sequential RANSAC	**.973**	**.173**	**.988**	–

We measure correspondence coverage by the fraction of occupied cells on a 32×32 grid. Both multigroup approaches cover more of the image than single RANSAC, with similar coverage.

Capacity. The 24×24 grid and four primitives per anchor give the lowest mean valid EPE in their respective blocks in Table 6. Higher capacity can reduce near-damage EPE at greater optimization cost.

Optimization length. Most of the valid EPE reduction occurs in the first 25 steps in Figure 6. The 50-step setting gives the lowest mean valid EPE in Table 6. Longer runs give higher NCC and slightly lower near-damage EPE, with slightly higher mean valid EPE.

Table 6: Capacity and optimization length on the design set. EPE is in original-resolution pixels. Time measures field optimization.

Factor	Setting	NCC ↑	Valid EPE ↓	Near EPE ↓	Time (s)
Anchor grid	16×16	.965	.199	1.251	50.3
	$\underline{24 \times 24}$	.973	**.173**	.988	51.3
	32×32	.975	.182	**.927**	52.1
	48×48	**.977**	.206	.963	54.7
Gauss./anchor	1	.964	.182	1.122	44.5
	$\underline{4}$	.973	**.173**	.988	51.3
	8	.974	.178	.996	57.0
	16	**.975**	.179	**.928**	65.9
Opt. steps	0	.881	.704	6.764	0.1
	25	.967	.193	1.148	25.5
	$\underline{50}$	.973	**.173**	.988	51.3
	100	.976	.177	.940	102.2
	500	**.980**	.198	**.910**	498.2

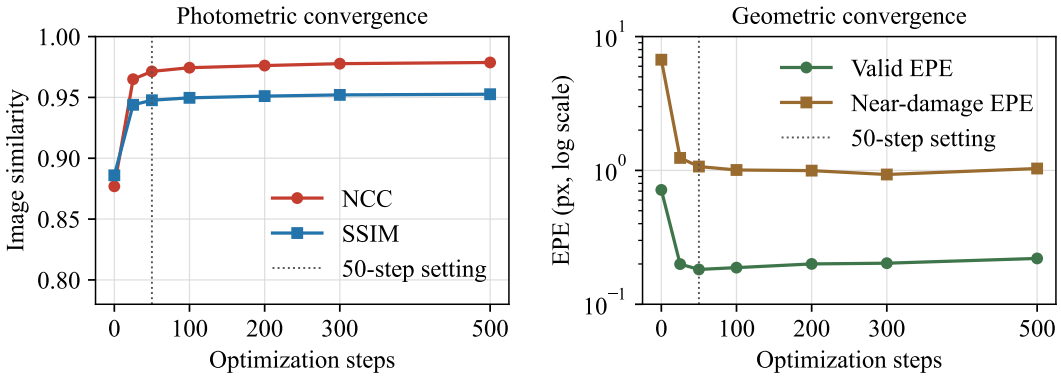

Figure 6: Image similarity and geometric error over optimization steps.

5 Limitations and Conclusion

GED learns anisotropic Gaussian supports to propagate local motions within tissue and limit their mixing near damage boundaries. The resulting field approximates abrupt motion changes without damage masks. Synthetic comparisons and manual landmarks on real sections show lower mean geometric error than the evaluated baselines. Support ablations link anisotropy to reduced cross-boundary mixing.

The evaluation covers controlled fold and tear deformations and five real section pairs. Future work will extend it to damage combined with broader nonrigid motion and stack-level alignment.

AI USE STATEMENT

We used generative AI tools for language editing, literature discovery, and assistance with experiment and analysis code. The authors take responsibility for the final manuscript and artifacts.

REPRODUCIBILITY STATEMENT

The supplementary archive contains anonymized data-generation and preprocessing code, generation settings, and two example pairs. Appendix B describes the matching, model, optimization, baselines, and evaluation settings. The full method and experiment code, sample splits, and analysis scripts will be released upon acceptance.

REFERENCES

Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski. A Database and Evaluation Methodology for Optical Flow. *International Journal of Computer Vision*, 92(1):1–31, March 2011. ISSN 0920-5691, 1573-1405. doi: 10.1007/s11263-010-0390-2.

F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989. ISSN 01628828. doi: 10.1109/34.24792.

Albert Cardona, Stephan Saalfeld, Johannes Schindelin, Ignacio Arganda-Carreras, Stephan Preibisch, Mark Longair, Pavel Tomancak, Volker Hartenstein, and Rodney J. Douglas. TrakEM2 Software for Neural Circuit Reconstruction. *PLoS ONE*, 7(6):e38011, June 2012. ISSN 1932-6203. doi: 10.1371/journal.pone.0038011.

Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010. ISSN 0920-5691, 1573-1405. doi: 10.1007/s11263-009-0275-4.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24 (6):381–395, June 1981. ISSN 0001-0782, 1557-7317. doi: 10.1145/358669.358692.

Wei Huang, Chang Chen, Zhiwei Xiong, Yueyi Zhang, Dong Liu, and Feng Wu. Learning to Restore ssTEM Images from Deformation and Corruption. In Adrien Bartoli and Andrea Fusiello (eds.), *Computer Vision – ECCV 2020 Workshops*, volume 12535, pp. 394–410. Springer International Publishing, Cham, 2020. ISBN 978-3-030-66414-5 978-3-030-66415-2. doi: 10.1007/978-3-030-66415-2_26.

Narayanan Kasthuri, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, Amelio Vázquez-Reina, Verena Kaynig, Thouis Raymond Jones, Mike Roberts, Josh Lyskowski Morgan, Juan Carlos Tapia, H. Sebastian Seung, William Gray Roncal, Joshua Tzvi Vogelstein, Randal Burns, Daniel Lewis Sussman, Carey Eldin Priebe, Hanspeter Pfister, and Jeff William Lichtman. Saturated Reconstruction of a Volume of Neocortex. *Cell*, 162(3):648–661, July 2015. ISSN 00928674. doi: 10.1016/j.cell.2015.06.054.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4):1–14, August 2023. ISSN 0730-0301, 1557-7368. doi: 10.1145/3592433.

Khaled Khairy, Gennady Denisov, and Stephan Saalfeld. Joint Deformable Registration of Large EM Image Volumes: A Matrix Solver Approach, April 2018.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

Kisuk Lee, Nicholas Turner, Thomas Macrina, Jingpeng Wu, Ran Lu, and H Sebastian Seung. Convolutional nets for reconstructing neural circuits from brain images acquired by serial section electron microscopy. *Current Opinion in Neurobiology*, 55:188–198, April 2019. ISSN 09594388. doi: 10.1016/j.conb.2019.04.001.

John P. Lewis. Fast normalized cross-correlation. In *Vision Interface*, volume 10, pp. 120–123, 1995.

Jihe Li, Xiang Liu, Fabian Zhang, Xixin Cao, Joachim M. Buhmann, Ye Zhang, and Xia Li. Gaussian primitives for deformable image registration. *Physics and Imaging in Radiation Oncology*, 35:100821, July 2025. ISSN 24056316. doi: 10.1016/j.phro.2025.100821.

Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 17627–17638, 2023.

Xinzhao Liu, Yueyi Zhang, Shenglong Zhou, Zhiwei Xiong, and Xiaoyan Sun. Electron Microscopy Image Registration Using Correlation Volume. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, Cartagena, Colombia, April 2023. IEEE. ISBN 978-1-6654-7358-3. doi: 10.1109/ISBI53787.2023.10230498.

Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3431–3440, Boston, MA, USA, June 2015. IEEE. ISBN 978-1-4673-6964-0. doi: 10.1109/CVPR.2015.7298965.

Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-GS: Structured 3D Gaussians for View-Adaptive Rendering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20654–20664, Seattle, WA, USA, June 2024. IEEE. ISBN 979-8-3503-5300-6. doi: 10.1109/CVPR52733.2024.01952.

Yanan Lv, Haoze Jia, Xi Chen, Haiyang Yan, and Hua Han. Aligning and Restoring Imperfect ssEM Images for Continuity Reconstruction. In Marius George Linguraru, Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker, Karim Lekadir, and Julia A. Schnabel (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, pp. 543–552, Cham, 2024. Springer Nature Switzerland. ISBN 978-3-031-72069-7. doi: 10.1007/978-3-031-72069-7_51.

Eric Mitchell, Stefan Keselj, Sergiy Popovych, Davit Buniatyan, and H. Sebastian Seung. Siamese Encoding and Alignment by Multiscale Learning with Self-Supervision, 2019.

Sergiy Popovych, J. Alexander Bae, and H. Sebastian Seung. Caesar: Segment-Wise Alignment Method for Solving Discontinuous Deformations. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pp. 1214–1218, April 2020. doi: 10.1109/ISBI45749.2020.9098363.

Sergiy Popovych, Thomas Macrina, Nico Kemnitz, Manuel Castro, Barak Nehoran, Zhen Jia, J. Alexander Bae, Eric Mitchell, Shang Mu, Eric T. Trautman, Stephan Saalfeld, Kai Li, and H. Sebastian Seung. Petascale pipeline for precise alignment of images from serial section electron microscopy. *Nature Communications*, 15(1):289, January 2024. ISSN 2041-1723. doi: 10.1038/s41467-023-44354-0.

Stephan Saalfeld, Albert Cardona, Volker Hartenstein, and Pavel Tomančák. As-rigid-as-possible mosaicking and serial section registration of large ssTEM datasets. *Bioinformatics*, 26(12):i57–i63, June 2010. ISSN 1367-4811, 1367-4803. doi: 10.1093/bioinformatics/btq219.

Stephan Saalfeld, Richard Fetter, Albert Cardona, and Pavel Tomancak. Elastic volume reconstruction from series of ultra-thin microscopy sections. *Nature Methods*, 9(7):717–720, July 2012. ISSN 1548-7105. doi: 10.1038/nmeth.2072.

Louis K. Scheffer, Bill Karsh, and Shiv Vitaladevun. Automated Alignment of Imperfect EM Images for Neural Reconstruction, 2013.

Xin Tian, Jiazheng Wang, Yuxi Zhang, Xiang Chen, Renjiu Hu, Gaolei Li, Min Liu, and Hang Zhang. Gaussian Primitive Optimized Deformable Retinal Image Registration. In James C. Gee, Daniel C. Alexander, Jaesung Hong, Juan Eugenio Iglesias, Carole H. Sudre, Archana Venkataraman, Polina Golland, Jong Hyo Kim, and Jinah Park (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, volume 15963, pp. 218–228. Springer Nature Switzerland, Cham, 2025. ISBN 978-3-032-04964-3 978-3-032-04965-0. doi: 10.1007/978-3-032-04965-0_21.

Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April 2004. ISSN 1941-0042. doi: 10.1109/TIP.2003.819861.

Tong Xin, Lijun Shen, Linlin Li, Xi Chen, and Hua Han. Expected affine: A registration method for damaged section in serial sections electron microscopy. *Frontiers in Neuroinformatics*, 16, September 2022. ISSN 1662-5196. doi: 10.3389/fninf.2022.944050.

Tong Xin, Yanan Lv, Haoran Chen, Linlin Li, Lijun Shen, Guangcun Shan, Xi Chen, and Hua Han. A novel registration method for long-serial section images of EM with a serial split technique based on unsupervised optical flow network. *Bioinformatics*, 39(8):btad436, August 2023. ISSN 1367-4811. doi: 10.1093/bioinformatics/btad436.

Inwan Yoo, David G. C. Hildebrand, Willie F. Tobin, Wei-Chung Allen Lee, and Won-Ki Jeong. ssEMnet: Serial-section Electron Microscopy Image Registration using a Spatial Transformer Network with Learned Features. *CoRR*, abs/1707.07833, 2017.

Zhenbang Zhang, Hongjia Li, Zhiqiang Xu, Wenjia Meng, and Renmin Han. A Gaussian Filter-Based 3D Registration Method for Series Section Electron Microscopy. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(1):1156–1164, April 2025. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v39i1.32103.

Xiaoming Zhao, Xingming Wu, Weihai Chen, Peter C. Y. Chen, Qingsong Xu, and Zhengguo Li. ALIKED: A Lighter Keypoint and Descriptor Extraction Network via Deformable Transformation. *IEEE Transactions on Instrumentation and Measurement*, 72:1–16, 2023. ISSN 1557-9662. doi: 10.1109/TIM.2023.3271000.

Shenglong Zhou, Zhiwei Xiong, Chang Chen, Xuejin Chen, Dong Liu, Yueyi Zhang, Zheng-Jun Zha, and Feng Wu. Fast and Accurate Electron Microscopy Image Registration with 3D Convolution. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, volume 11764, pp. 478–486. Springer International Publishing, Cham, 2019. ISBN 978-3-030-32238-0 978-3-030-32239-7. doi: 10.1007/978-3-030-32239-7_53.

A Synthetic Data Generation

We generate folds and tears using piecewise translations across approximately straight boundaries with random orientations. Folds move tissue toward the boundary and create overlaps. Tears move the two sides apart and create gaps. Displacement magnitude determines the low, medium, and high severity levels. Smooth controls use thin-plate-spline deformations. Each pair includes the ground-truth deformation and a validity mask. The supplementary code provides the generation procedure and parameter settings.

B Implementation Details

B.1 GED settings

Matching and motion decomposition. ALIKED-N16 extracts up to 4096 keypoints at the original resolution with detection threshold 0.01. Images are matched in 1024×1024 patches with 128-pixel overlap; LightGlue uses filter threshold 0.1. In the sequential decomposition of Section 3.2, the consensus tolerance is $\tau_m = 6$ pixels, the procedure stops after at most 20 hypotheses or when no consensus set with at least 10 inliers remains, and unassigned correspondences are discarded. Each OpenCV RANSAC fit uses at most 2000 iterations, confidence 0.99, and 50 refinement iterations.

Anchors and initialization. The grid covers the optimization image; an anchor is activated when its cell contains an accepted correspondence. Each active anchor generates four primitives with initial offsets right, left, down, and up by half a cell. Their initial displacement is the mean reference-minus-damaged displacement of the cell's correspondences. Local residual parameters are initialized from the fitted motion groups. Both log-scale vectors are initialized with the logarithm of the distance to the nearest active anchor. If only one anchor is active, we use the larger grid spacing instead of this distance. Anchor features are drawn from $\mathcal{N}(0, 0.01^2)$.

MLPs. Following Scaffold-GS (Lu et al., 2024), each of the N anchors has a fixed location $\mathbf{a}_j$, a learnable feature $\mathbf{f}_j \in \mathbb{R}^{32}$, and log-scale vectors $\boldsymbol{\rho}_j^{\mu}, \boldsymbol{\rho}_j^{\Sigma} \in \mathbb{R}^2$. Each anchor generates K primitives, giving $P = NK$ in total. We use $i \in \{1, \dots, P\}$ as the flattened index of primitive (j, k). Its center offset $\mathbf{o}_{jk}$, displacement offset $\mathbf{t}_{jk}$, and local residual parameter ϕ_{jk} are optimized directly. Shared two-layer MLPs F_{Σ}, F_{α} decode its scale $\mathbf{s}_{jk}$, orientation ψ_{jk}, and opacity α_{jk}. Each MLP has 32 GELU hidden units and decodes the primitive parameters as follows.

$$\begin{aligned} \boldsymbol{\mu}_{jk} &= \mathbf{a}_j + \exp(\boldsymbol{\rho}_j^{\mu}) \odot \mathbf{o}_{jk}, \quad [\bar{\mathbf{s}}_{jk}, \psi_{jk}] = F_{\Sigma}([\mathbf{f}_j, \mathbf{o}_{jk}]), \\ \mathbf{s}_{jk} &= \exp(\boldsymbol{\rho}_j^{\Sigma}) \odot \sigma(\bar{\mathbf{s}}_{jk}), \qquad\qquad \alpha_{jk} = \sigma\big(F_{\alpha}([\mathbf{f}_j, \mathbf{o}_{jk}])\big), \end{aligned} \tag{5}$$

where σ is the sigmoid and $[\cdot, \cdot]$ denotes concatenation; scales are clamped below at one pixel. Both decoders take $[\mathbf{f}_j, \mathbf{o}_{jk}] \in \mathbb{R}^{34}$.

Rendering. Gaussian responses are evaluated on the 512-pixel optimization grid. Gaussian candidates are selected for each tile using a 3σ criterion; a tile with no candidate keeps the primitive nearest to its center.

Optimization. All GED runs use seed 42 for Python, NumPy, PyTorch, and OpenCV RANSAC. We run Adam with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Anchor features, log-scales, F_{Σ}, and local residual parameters use a learning rate of 3×10^{-2}. The rates for F_{α}, center offsets, and displacement offsets are 5×10^{-3}, 5×10^{-2}, and 2×10^{-1}, respectively.

B.2 Baseline settings

RANSAC and EA use the SuperPoint features and radius matching of the EA implementation (Xin et al., 2022). Elastic uses `layerScale` $= 0.20$ at a 512-pixel short edge to preserve the internal matching resolution of its 1024-pixel default. SEAMLeSS uses the Corgie pipeline (Popovych et al., 2024).

GPO uses the official DCN defaults with normalized grayscale images and raw subpixel ALIKED+LightGlue matches as input. We adjust the match count to its input requirements and scale Gaussian radii to the optimization grid.

Using the default GaussianDIR settings, we convert coordinates, Gaussian parameters, and spatial operations to two dimensions for section registration.

B.3 Evaluation and runtime

All methods use bilinear image warping with replicated borders. Instance labels use nearest-neighbor interpolation and zero padding. Undefined pixels in Elastic maps are filled from the nearest valid pixel; they account for 2.9% of pixels on average.

Timing. Mean measured wall time over the primary evaluation set is 1.0 s for RANSAC, 15.6 s for EA, 25.9 s for Elastic, 16.4 s for SEAMLeSS, and 68.5 s for GED. These measurements include matcher initialization, matching, optimization, and output. GPO and GaussianDIR* take 18.2 s and 16.0 s per pair, respectively, for image loading, optimization, and output. GPO uses precomputed matches.

Environment. Ubuntu 18.04.6, Python 3.10.17, PyTorch 1.13.1, CUDA 11.6, cuDNN 8.3.2, OpenCV 4.9.0, NumPy 1.26.4; two Xeon Gold 6142 CPUs, 512 GB memory, two Tesla V100 32 GB GPUs. GPO uses Python 3.9.22 and PyTorch 2.0.1 with CUDA 11.7.

C Additional Synthetic Results

C.1 Damage severity

Figure 7 reports NCC and valid EPE by damage type and severity. Tables 7 and 8 give the values, while Table 3 in the main text reports near-damage EPE by damage type.

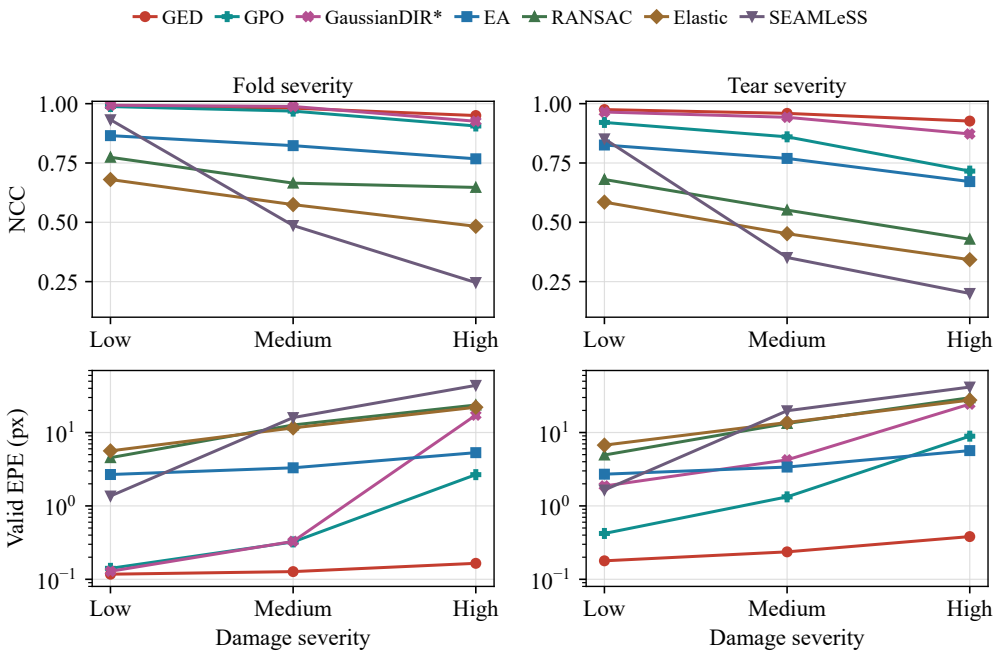

Figure 7: NCC and valid EPE by damage type and severity. EPE is shown on a logarithmic scale.

Table 7: Fold results by severity. EPE in original-resolution pixels.

	Low		Medium		High	
Method	NCC ↑	Valid EPE ↓	NCC ↑	Valid EPE ↓	NCC ↑	Valid EPE ↓
RANSAC	.774	4.544	.665	12.669	.647	23.709
EA	.865	2.678	.823	3.306	.768	5.317
Elastic	.680	5.620	.575	11.559	.483	22.117
SEAMLeSS	.932	1.364	.487	15.981	.246	43.889
GPO	.988	.141	.969	.323	.906	2.667
GaussianDIR*	**.993**	.129	**.988**	.329	.926	17.267
GED	.992	**.117**	.981	**.127**	**.950**	**.165**

Of the 560 damaged pairs, 495 contain valid ground-truth pixels in the near-damage band. The other 65 are folds with no valid ground-truth pixels in that band.

Table 8: Tear results by severity. EPE in original-resolution pixels.

	Low		Medium		High	
Method	NCC ↑	Valid EPE ↓	NCC ↑	Valid EPE ↓	NCC ↑	Valid EPE ↓
RANSAC	.680	4.937	.551	13.299	.429	29.783
EA	.826	2.705	.769	3.384	.672	5.671
Elastic	.585	6.731	.452	13.698	.342	27.654
SEAMLeSS	.852	1.648	.352	19.801	.200	41.647
GPO	.921	.421	.860	1.328	.716	8.920
GaussianDIR*	.965	1.868	.943	4.241	.872	24.408
GED	**.975**	**.179**	**.959**	**.237**	**.927**	**.382**

D Support Geometry Analysis

D.1 Support evaluation

The isotropic variant replaces the two decoded scales with their mean and sets $\psi_i = 0$.

Cross-boundary leakage. Each pixel is assigned to a side of the damage boundary by $\mathrm{sign}(\mathbf{u}_{\mathrm{gt}} \cdot \mathbf{n})$, where $\mathbf{n}$ is the generator's boundary normal, keeping pixels with $|\mathbf{u}_{\mathrm{gt}} \cdot \mathbf{n}| > 0.25$ original-resolution pixels. Remaining pixels and primitive centers receive the side of the nearest labeled pixel. The near-damage band is the original-resolution damage mask dilated by 32 original-resolution pixels, and its pixel centers are indexed onto the optimization grid. Leakage is evaluated on valid band pixels, with subsampling where needed.

D.2 Transition-width derivation

Consider the normal slice through a damage boundary. Two primitives have centers at $s = \pm a$ with equal normal scale σ_n and equal opacity. Their unnormalized weights are $w_{\pm}(s) = \exp[-(s \mp a)^2/(2\sigma_n^2)]$, and their normalized mixture weight is

$$\pi_+(s) = \frac{w_+(s)}{w_+(s) + w_-(s)} = \mathrm{sigmoid}\left(\frac{2as}{\sigma_n^2}\right). \tag{6}$$

Setting $\pi_+ = 0.1$ and $\pi_+ = 0.9$ gives $2as/\sigma_n^2 = \mp \ln 9$, so the 10–90 transition width is

$$W_{10\text{–}90} = \frac{\sigma_n^2 \ln 9}{a}. \tag{7}$$

For fixed center separation, the normal scale controls this width independently of the tangential scale. Figure 8 illustrates anisotropic support geometry near a fold junction.

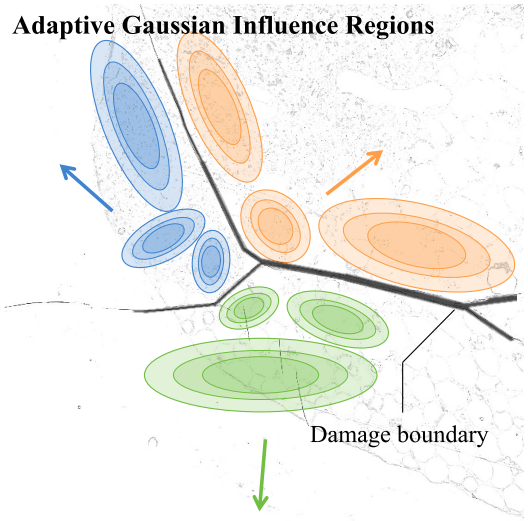

Figure 8: Schematic of Gaussian primitives near a fold junction. Ellipses show the 1σ–3σ level sets of primitives associated with three neighboring tissue regions; arrows show their local motions.

E Additional Real-Section Results

Landmarks are selected on recognizable corresponding structures. They are placed near damage in damaged sections and distributed evenly across the image in sections without visible damage. All methods use the same 80 correspondences. We bilinearly sample each saved backward map at the reference landmarks and measure distances to their annotated counterparts. Elastic returned no registration for DOLW7, so its five-pair mean is unavailable.

Table 9 reports NCC on the same five pairs. NCC and landmark error give different method rankings on SW13L63, where GPO has higher image similarity and GED has lower geometric error.

Table 9: NCC on the real section pairs.

Pair	RANSAC	EA	Elastic	SEAMLeSS	GPO	GaussianDIR*	**GED**
DOLW7	.101	.168	—	.130	.189	**.380**	.356
SS2L30	.629	.546	.543	.567	.810	**.847**	.814
SW13L63	.164	.148	.129	.144	**.445**	.289	.416
ZW35L414	.720	.696	.685	.712	.871	**.918**	.887
ZW7L18	.835	.818	.819	.846	.946	**.950**	.942

Figure 9 shows the real section pairs and registration overlays corresponding to Table 2. DOLW7 contains an image-spanning fold, SW13L63 a tear, ZW35L414 and ZW7L18 shallow and branching folds, and SS2L30 nonlinear deformation without visible damage.

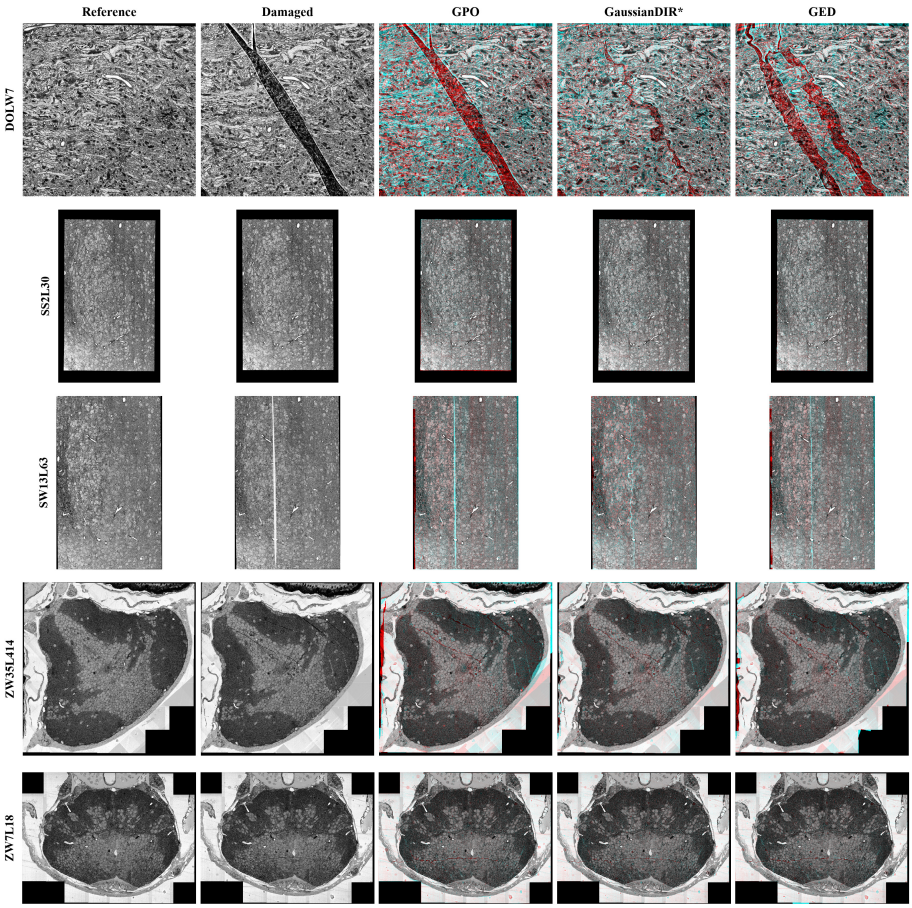

Figure 9: Registration overlays on real section pairs.