G2M-Reg: Graph-to-Manifold Verification for Robust Point Cloud Registration

Anonymous authors
Paper under double-blind review

Abstract

Robust point cloud registration remains challenging under low overlap, where sparse true matches can be overwhelmed by locally consistent outliers that form strong compatibility structures. Consequently, an incorrect pose may receive substantial correspondence support while failing to explain the underlying common surface. This reveals a fundamental limitation of correspondence-space reasoning: evidence that is effective for generating a pose is not necessarily sufficient for verifying its geometric validity. We propose G2M-Reg, a training-free Graph-to-Manifold Verification framework that separates pose proposal from pose verification. A correspondence graph first generates a set of plausible pose hypotheses. Instead of ranking them solely by correspondence-level evidence, G2M-Reg evaluates the cross-cloud surface relation induced by each candidate pose. Bidirectional partial-overlap consensus measures whether the two fragments mutually explain a plausible common region, while structural reliability evaluates whether the supporting geometry is coherent and sufficiently informative to constrain the pose. Graph, surface, and structural evidence is subsequently integrated in rank space to select a surface-valid hypothesis, followed by constrained surface-guided refinement. Experiments across indoor and outdoor benchmarks and diverse hand-crafted and learned correspondence sources demonstrate that G2M-Reg consistently improves robust pose selection, particularly in challenging low-overlap scenarios.

1 Introduction

Point cloud registration estimates the rigid transformation between partially overlapping 3D scans and is fundamental to robotic perception, SLAM, autonomous driving, and 3D reconstruction Ullah et al. (2024); Bailey & Durrant-Whyte (2006); Li et al. (2020b); Liu et al. (2021); Jaramillo & Sipiran (2024); Xu et al. (2024). Despite advances in hand-crafted and learned descriptors, low overlap, repetitive geometry, sensor noise, and nonuniform sampling can produce correspondence sets dominated by outliers, making reliable pose estimation challenging.

Existing robust estimators mainly improve correspondence selection or pose hypothesis generation. Sampling-based methods, including RANSAC variants Fischler & Bolles (1981); Li et al. (2020a); Barath & Matas (2017); Brachmann & Rother (2019); Wei et al. (2022), generate and evaluate pose hypotheses, while graph-based methods exploit rigid-invariant compatibility to recover reliable seeds, subgraphs, cliques, cores, or higher-order structures Bai et al. (2021); Chen et al. (2022); Zhang et al. (2023); Wu et al. (2026). Recent approaches further explore k-truss structures, efficient clique search, unsupervised consensus sampling, geometric proxies, and pretrained visual priors Wu et al. (2025); Yan et al. (2025a;b); Li et al. (2026); Zheng et al. (2025). Although these methods improve correspondence reliability and hypothesis generation under high outlier ratios, pose validity is still largely assessed from compatibility, inlier support, correspondence residuals, or related statistics derived from the same putative matches Fischler & Bolles (1981); Yang et al. (2021).

Such correspondence-level evidence can be misleading under low overlap. As illustrated in Fig. 1, sparse true matches may compete with dense but locally consistent outliers, allowing

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.

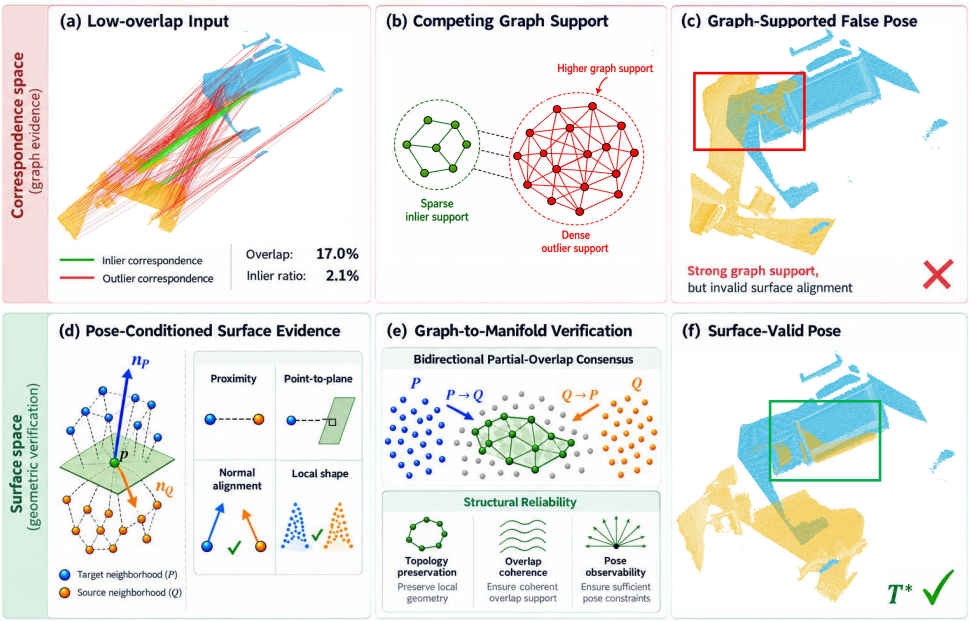

Figure 1: Motivation of G2M-Reg. Under low overlap, locally consistent outliers may receive strong graph support yet induce a surface-invalid pose. G2M-Reg therefore separates correspondence-space proposal from surface-space verification.

an incorrect pose to receive strong graph support without explaining the underlying common surface. The key issue is therefore not only whether a plausible hypothesis can be generated, but whether it is geometrically valid: correspondence support does not necessarily imply surface validity. This distinction cannot be fully resolved when pose proposal and verification remain tied to essentially the same evidence space.

We therefore formulate robust registration as a Graph-to-Manifold Verification problem and propose G2M-Reg, a training-free framework that separates correspondence-space pose proposal from surface-space geometric verification. A compatibility graph first generates plausible hypotheses, after which each candidate is evaluated through its induced pose-conditioned cross-cloud surface relation. Bidirectional partial-overlap consensus measures whether the two fragments mutually explain a plausible common region, while structural reliability evaluates whether the supporting geometry is coherent and sufficiently constraining. Graph, surface, and structural evidence are then integrated in rank space to select a surface-valid hypothesis, followed by constrained surface-guided refinement.

G2M-Reg operates on both hand-crafted and learned correspondences without training or retraining. Experiments on 3DMatch, 3DLoMatch, and KITTI show consistent improvements across indoor and outdoor settings, with particularly clear gains under low overlap. Controlled ablations further validate the roles of bidirectional surface consensus, structural reliability, rank-space fusion, and constrained refinement.

The main contributions are:

- We formulate robust registration as Graph-to-Manifold Verification, separating correspondence-space pose proposal from surface-space pose verification.
- We introduce pose-conditioned cross-cloud verification based on bidirectional partial-overlap consensus and structural reliability, enabling rejection of graph-supported but surface-invalid hypotheses.
- We develop a training-free realization that integrates heterogeneous evidence in rank space and validate it across indoor and outdoor benchmarks using both hand-crafted and learned correspondences.

2 Related Work

2.1 Correspondence-Based Registration

Correspondence-based registration estimates rigid transformations from putative matches between point clouds. Hand-crafted descriptors such as FPFH and SHOT Rusu et al. (2009); Salti et al. (2014), and learned features including FCGF, Predator, and GeoTransformer Choy et al. (2019); Huang et al. (2020); Qin et al. (2023), improve local matching, but repetitive geometry, noise, and limited overlap can still produce heavily contaminated correspondence sets.

Robust estimators address these outliers through sampling, optimization, or learned filtering. Representative approaches include RANSAC variants Fischler & Bolles (1981); Li et al. (2020a); Barath & Matas (2017); Wei et al. (2022), robust optimization methods such as FGR and TEASER++ Zhou et al. (2016); Yang et al. (2020), ICP-based refinement Besl & McKay (1992); Chen & Medioni (1992); Segal et al. (2009); Yang et al. (2016), and learned estimators such as DGR, PointDSC, and VBReg Choy et al. (2020); Bai et al. (2021); Jiang et al. (2023). Recent work further explores heuristic sampling, dual-space optimization, and zero-shot refinement Yan et al. (2025b); Li et al. (2026); Zheng et al. (2025). Despite these advances, pose validity is still largely assessed from correspondence-derived evidence or refinement of an existing estimate.

2.2 Graph-Based Registration

Graph-based methods represent correspondences as nodes and encode rigid-invariant compatibility as edges. SC^2-PCR introduces second-order spatial compatibility Chen et al. (2022), while MAC and FastMAC exploit maximal cliques for robust and efficient consensus search Zhang et al. (2023; 2024). Recent approaches further explore multi-hop expansion, k-truss structures, correspondence regeneration, and lightweight clique search Wu et al. (2025); Zhao et al. (2025); Wu et al. (2026); Yan et al. (2025a).

These methods establish correspondence graphs as effective hypothesis generators under high outlier ratios, and recent work also improves ranking with spatial coverage or residual quality Wu et al. (2026; 2025). However, such evidence remains primarily derived from the input correspondences or their induced inlier sets Fischler & Bolles (1981); Yang et al. (2021); Chen et al. (2023). In contrast, G2M-Reg uses the graph only for pose proposal and verifies each candidate in a distinct pose-conditioned surface evidence space.

3 Method

G2M-Reg formulates robust registration as Graph-to-Manifold Verification: a correspondence graph proposes plausible pose hypotheses, while each candidate is verified through the cross-cloud surface relation induced by that pose. As shown in Fig. 2, the framework consists of graph-supported hypothesis generation, pose-conditioned surface verification, grouped rank-space evidence fusion, and constrained surface-guided refinement.

3.1 Problem Formulation and Overview

Given source and target point clouds $P = \{p_i\}_{i=1}^{N} \subset \mathbb{R}^3$ and $Q = \{q_j\}_{j=1}^{M} \subset \mathbb{R}^3$, rigid registration seeks $T^* = (R^*, t^*) \in \mathrm{SE}(3)$ that aligns their overlapping regions. Feature matching provides a set of putative correspondences $\mathcal{C} = \{c_k = (x_k, y_k)\}_{k=1}^{N_c}$, which may be heavily contaminated by outliers under low overlap.

A second-order compatibility graph generates a compact hypothesis set $\mathcal{H} = \{T_h = (R_h, t_h)\}_{h=1}^{H}$. Each candidate T_h induces a pose-conditioned cross-cloud surface relation characterized by local neighborhood structure, geometric affinity, and high-evidence support. G2M-Reg evaluates graph evidence C_h, manifold consensus M_h, and structural reliability, integrates these heterogeneous cues in rank space, and selectively refines the highest-ranked candidates.

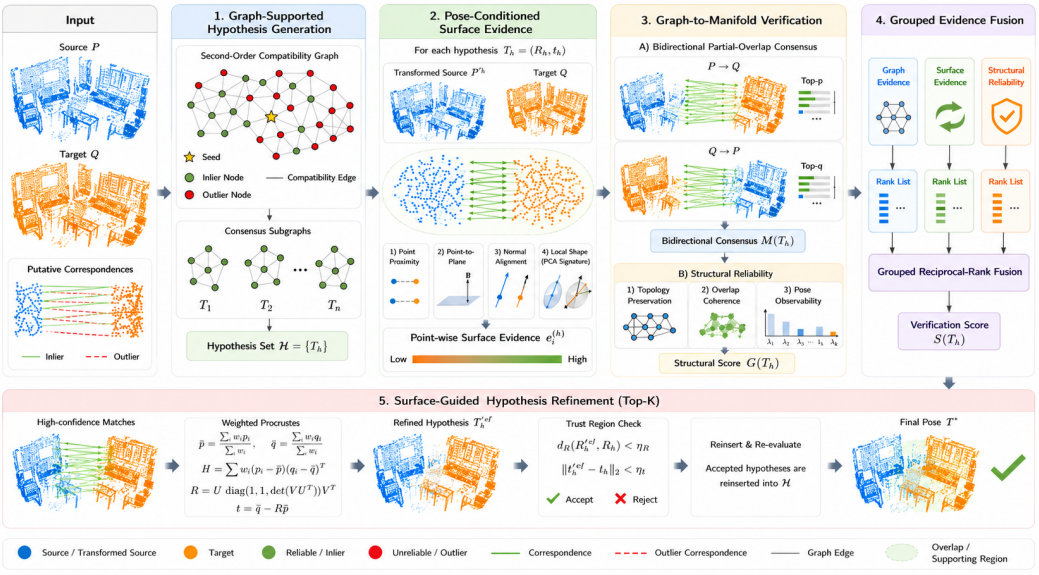

Figure 2: Overview of G2M-Reg. A correspondence graph proposes pose hypotheses, which are verified by pose-conditioned surface consensus and structural reliability. Graph, surface, and structural evidence are fused in rank space, followed by selective surface-guided refinement.

3.2 Graph-Supported Hypothesis Generation

G2M-Reg uses the correspondence graph only to propose plausible poses. Following SC^2-PCR Chen et al. (2022), correspondences are connected according to rigid-invariant geometric compatibility, and common-neighbor support identifies high-confidence consensus groups. Weighted SVD Kabsch (1976) then generates the hypothesis set $\mathcal{H}$, while the aggregated second-order support of each hypothesis defines its graph evidence C_h.

Because C_h is derived from the same correspondence structure used for hypothesis generation, it provides useful proposal evidence but does not by itself establish geometric validity. G2M-Reg therefore verifies each candidate in a distinct surface evidence space.

3.3 Pose-Conditioned Cross-Cloud Manifold

For each point, we precompute a surface normal and a compact PCA-based shape descriptor, denoted by (n_i^P, g_i^P) and (n_j^Q, g_j^Q). Under hypothesis $T_h = (R_h, t_h)$, the source position and normal become

$$\hat{p}_i^{(h)} = R_h p_i + t_h, \qquad \hat{n}_i^{(h)} = R_h n_i^P.$$

For each transformed source point $\hat{p}_i^{(h)}$, a local neighborhood $\mathcal{N}_h(i) \subset Q$ defines candidate cross-cloud surface relations.

Point-wise Surface Evidence. For each pair $(\hat{p}_i^{(h)}, q_j)$ with $j \in \mathcal{N}_h(i)$, we define

$$E_{ij}^{(h)} = \alpha_d \frac{\|\hat{p}_i^{(h)} - q_j\|_2}{\sigma_d} + \alpha_\pi \frac{\left|(n_j^Q)^\top (\hat{p}_i^{(h)} - q_j)\right|}{\sigma_\pi} + \alpha_n \left(1 - \left|(\hat{n}_i^{(h)})^\top n_j^Q\right|\right) + \alpha_g \frac{\|g_i^P - g_j^Q\|_1}{\sigma_g}. \tag{1}$$

The four terms respectively measure point proximity, point-to-plane consistency, normal alignment, and local shape compatibility. The residual is converted to surface affinity by

$W_{ij}^{(h)} = \exp(-E_{ij}^{(h)})$. Rather than using a hard nearest-neighbor assignment, neighboring responses are softly aggregated:

$$e_i^{P\to Q}(T_h) = \sum_{j\in\mathcal{N}_h(i)} \pi_{ij}^{(h)} W_{ij}^{(h)},$$
$$\pi_{ij}^{(h)} = \frac{\exp\left(-E_{ij}^{(h)}/\tau_a\right)}{\sum_{k\in\mathcal{N}_h(i)} \exp\left(-E_{ik}^{(h)}/\tau_a\right)}. \tag{2}$$

Thus, $e_i^{P\to Q}(T_h)$ measures how well the target surface explains the transformed source point. The reverse direction analogously produces $e_j^{Q\to P}(T_h)$.

3.4 Graph-to-Manifold Verification

A geometrically valid pose should satisfy two complementary requirements: the two fragments should mutually explain a plausible partial-overlap region, and the supporting geometry should be spatially coherent and sufficiently informative to constrain the pose. G2M-Reg evaluates these properties through bidirectional partial-overlap consensus and structural reliability.

3.4.1 Bidirectional Partial-Overlap Consensus

Non-overlapping regions should not contribute to pose verification. For a candidate overlap ratio ρ, G2M-Reg retains only the strongest $\lceil \rho N \rceil$ source-to-target responses:

$$M_\rho^{P\to Q}(T_h) = \frac{1}{\lceil \rho N \rceil} \sum_{i\in \mathrm{Top}_{\lceil \rho N \rceil}(e^{P\to Q})} e_i^{P\to Q}(T_h). \tag{3}$$

Since the true overlap ratio is unknown, evidence is aggregated over a set of candidate ratios $\mathcal{R}$:

$$M^{P\to Q}(T_h) = \sum_{\rho\in\mathcal{R}} \omega_\rho M_\rho^{P\to Q}(T_h). \tag{4}$$

The reverse direction gives $M^{Q\to P}(T_h)$. We then define bidirectional manifold consensus by harmonic fusion:

$$M_h = \frac{2M^{P\to Q}(T_h)M^{Q\to P}(T_h)}{M^{P\to Q}(T_h) + M^{Q\to P}(T_h) + \varepsilon}. \tag{5}$$

Requiring support in both directions suppresses one-sided pseudo-alignments caused by repetitive or locally similar geometry.

3.4.2 Structural Reliability

High manifold consensus alone may still arise from fragmented, spatially inconsistent, or geometrically degenerate support. We therefore identify the high-evidence region

$$\Omega_h = \mathrm{Top}_K\left(\{e_i^{P\to Q}(T_h)\}_{i=1}^N\right),$$
$$\pi_h(i) = \arg\max_{j\in\mathcal{N}_h(i)} W_{ij}^{(h)}, \qquad i \in \Omega_h, \tag{6}$$

where $\pi_h(i)$ denotes the most strongly supported target point. Structural reliability evaluates three complementary properties: topology preservation, overlap coherence, and pose observability.

Topology preservation. A valid cross-cloud relation should preserve local source geometry. For local pairs $(i,k) \in \mathcal{P}_h$ sampled from Ω_h, we define

$$\begin{aligned} \delta_{ik}^{(h)} &= \left| \|p_i - p_k\|_2 - \|q_{\pi_h(i)} - q_{\pi_h(k)}\|_2 \right|, \\ G_{\text{top}}(T_h) &= \frac{1}{|\mathcal{P}_h|} \sum_{(i,k)\in\mathcal{P}_h} \exp\left(-\frac{\delta_{ik}^{(h)}}{\tau_t}\right). \end{aligned} \tag{7}$$

This term favors hypotheses whose induced target relations preserve the local geometric organization of the source support.

Overlap coherence. Reliable surface support should form a coherent region rather than a set of isolated high-affinity points. We measure

$$G_{\text{coh}}(T_h) = \frac{1}{|\Omega_h|} \sum_{i\in\Omega_h} \frac{|\mathcal{N}_P(i) \cap \Omega_h|}{|\mathcal{N}_P(i)| + \varepsilon}. \tag{8}$$

A higher score indicates that high-evidence points form spatially coherent support on the source surface.

Pose observability. Even geometrically consistent support may be unreliable if it cannot sufficiently constrain all six rigid-pose dimensions. We first center and scale-normalize the support as $\tilde{p}_i^{(h)} = (\hat{p}_i^{(h)} - \mu_h)/s_h$, where μ_h and s_h denote the centroid and RMS radius, respectively. The normalized point-to-plane Jacobian and information matrix are

$$\begin{aligned} \tilde{J}_i^{(h)} &= \left[(\tilde{p}_i^{(h)} \times n_{\pi_h(i)}^Q)^\top \quad (n_{\pi_h(i)}^Q)^\top \right], \\ \tilde{\mathcal{I}}_h &= \frac{1}{|\Omega_h|} \sum_{i\in\Omega_h} e_i^{P\to Q}(T_h) (\tilde{J}_i^{(h)})^\top \tilde{J}_i^{(h)}. \end{aligned} \tag{9}$$

Let $0 \le \xi_1 \le \cdots \le \xi_6$ denote the eigenvalues of $\tilde{\mathcal{I}}_h$, and let $\bar{\xi}_h = \frac{1}{6}\sum_{r=1}^{6} \xi_r$. We define pose observability as

$$G_{\text{obs}}(T_h) = \frac{\xi_1}{\bar{\xi}_h + \tau_{\text{obs}}}. \tag{10}$$

The score approaches zero for rank-deficient support and increases when all six pose directions are sufficiently constrained.

3.5 Grouped Evidence Fusion

Graph, manifold, and structural evidence have different numerical scales and distributions, making direct raw-score fusion sensitive to calibration. G2M-Reg therefore integrates heterogeneous evidence in rank space. For evidence group $r \in \{c, m, g\}$ with measures $\mathcal{U}_r$, we define

$$\Phi_r(T_h) = \sum_{u\in\mathcal{U}_r} \frac{1}{\kappa + \text{rank}_u(T_h)}, \tag{11}$$

where κ stabilizes reciprocal ranks. The final verification score is

$$S(T_h) = \sum_{r \in \{c,m,g\}} \lambda_r \Phi_r(T_h). \tag{12}$$

Rank-space integration places heterogeneous evidence on comparable scales while preserving within-group ordering, thereby reducing sensitivity to raw-score calibration.

3.6 Surface-Guided Hypothesis Refinement

Refinement is applied only after verification has identified promising alignment modes. For each top-ranked T_h, high-evidence pairs $\{(p_i, q_{\pi_h(i)}) \mid i \in \Omega_h\}$ are weighted by surface evidence and used to solve

$$T_h^{\text{ref}} = \arg \min_{R \in \text{SO}(3),\, t \in \mathbb{R}^3} \sum_{i \in \Omega_h} w_i^{(h)} \left\| Rp_i + t - q_{\pi_h(i)} \right\|_2^2. \tag{13}$$

The weighted problem admits a closed-form SVD solution. To prevent refinement from switching to a different alignment mode, an update is accepted only if

$$d_R(R_h^{\text{ref}}, R_h) < \eta_R, \qquad \|t_h^{\text{ref}} - t_h\|_2 < \eta_t. \tag{14}$$

Accepted refinements are reinserted into $\mathcal{H}$ and re-evaluated using the same verification criterion. The refined and original hypotheses therefore compete under a unified scoring function, and the highest-scoring candidate is returned as the final registration.

4 Experiments

We first introduce the experimental setup and compare G2M-Reg with state-of-the-art robust registration methods. We then analyze the contributions of its major components and present qualitative results on challenging low-overlap cases. Additional analyses are provided in the supplementary material.

4.1 Experimental Setup

We evaluate G2M-Reg on the indoor 3DMatch Zeng et al. (2016) and 3DLoMatch Huang et al. (2020) datasets and the outdoor KITTI Geiger et al. (2012) dataset. 3DLoMatch contains point cloud pairs with only 10%–30% overlap and therefore provides a particularly challenging low-overlap setting. FPFH Rusu et al. (2009) and FCGF Choy et al. (2019) are used to generate initial correspondences on all three datasets.

We compare G2M-Reg with representative learning-based and learning-free robust registration methods. The learning-based methods include PointDSC Bai et al. (2021), VBReg Jiang et al. (2023), and MHopReg Wu et al. (2026). The learning-free methods include RANSAC Fischler & Bolles (1981), TEASER++ Yang et al. (2020), SC^2-PCR Chen et al. (2022), MAC Zhang et al. (2023), FastMAC Zhang et al. (2024), TCF Shi et al. (2024), MAC++ Zhang et al. (2025), TurboReg Yan et al. (2025a), PointTruss Wu et al. (2025), and DualReg Li et al. (2026). We report registration recall (RR), rotation error (RE), translation error (TE), and runtime.

4.2 Comparisons with State-of-the-Art Methods

As shown in Table 1, G2M-Reg achieves strong performance across all three benchmarks and both correspondence sources. On 3DMatch, it obtains the highest RR of 88.91%/94.33% and the lowest TE of 6.00/5.93 cm with FPFH/FCGF, respectively. The gains are most pronounced on 3DLoMatch, where G2M-Reg reaches 47.89%/65.69% RR, outperforming the corresponding second-best results by 3.31/4.05 percentage points while achieving 9.15/9.49 cm TE. On KITTI, where RR is already near saturation, G2M-Reg achieves

Table 1: Performance comparison on the 3DMatch, 3DLoMatch, and KITTI datasets. Bold and underline denote the best and second-best accuracy results, respectively.

	Method	FPFH				FCGF			
		RR↑ (%)	RE↓ (°)	TE↓ (cm)	Time↓ (s)	RR↑ (%)	RE↓ (°)	TE↓ (cm)	Time↓ (s)
3DMatch	PointDSC	78.00	2.08	6.48	0.04	92.85	2.04	6.50	0.05
	VBReg	81.39	2.12	6.64	0.18	93.47	2.04	6.49	0.19
	MHopReg	88.56	**1.67**	6.24	0.10	93.96	1.70	6.16	0.10
	RANSAC10K	68.76	3.86	10.73	1.58	91.56	2.73	8.31	1.55
	RANSAC50K	76.71	3.34	9.46	8.16	92.05	2.52	7.84	7.99
	TEASER++	78.74	2.52	7.55	3.02	85.64	2.75	9.26	8.53
	SC^2-PCR	83.80	2.12	6.60	0.89	92.98	2.05	6.47	0.89
	FastMAC	82.44	2.11	6.68	1.27	92.61	2.01	6.43	1.26
	MAC	83.92	2.11	6.79	2.10	93.72	2.03	6.54	1.82
	TCF	78.13	2.37	7.54	0.08	88.79	2.19	7.24	0.80
	MAC++	83.73	2.11	6.78	4.28	93.22	2.10	6.83	5.19
	TurboReg	84.10	2.19	6.81	0.03	92.98	2.05	6.53	0.03
	PointTruss	84.70	1.80	6.22	0.20	93.84	1.70	6.13	0.20
	DualReg	84.41	1.75	6.19	0.14	93.65	**1.65**	6.13	0.45
	G2M-Reg (ours)	**88.91**	1.77	**6.00**	0.31	**94.33**	1.74	**5.93**	0.34
3DLoMatch	PointDSC	28.86	3.92	10.27	0.05	56.65	3.79	10.42	0.05
	VBReg	37.00	4.04	10.20	0.19	59.74	3.80	10.46	0.18
	MHopReg	–	–	–	–	61.59	3.23	9.75	–
	RANSAC10K	18.75	6.53	14.35	1.22	49.02	4.74	12.36	1.39
	RANSAC50K	26.56	5.55	12.92	6.74	54.80	4.51	12.03	7.32
	TEASER++	28.30	4.75	11.64	1.80	40.43	4.82	13.68	2.47
	SC^2-PCR	39.53	4.11	10.39	0.85	58.06	3.75	10.44	0.86
	FastMAC	38.18	4.12	10.45	1.03	57.72	3.80	10.62	0.70
	MAC	41.10	4.03	10.46	1.85	60.13	3.75	10.59	1.97
	TCF	32.57	4.26	11.38	0.06	48.91	4.16	11.25	0.12
	MAC++	44.58	4.28	11.38	3.29	61.09	3.99	11.39	3.72
	TurboReg	39.19	4.03	10.36	0.02	58.34	3.76	10.35	0.03
	PointTruss	43.96	**2.89**	**8.93**	–	61.64	3.30	9.72	–
	DualReg	41.66	3.08	9.36	0.11	60.86	**3.06**	9.61	0.16
	G2M-Reg (ours)	**47.89**	3.16	9.15	0.30	**65.69**	3.16	**9.49**	0.33
KITTI	PointDSC	97.12	0.45	9.80	0.04	97.30	0.41	9.83	0.05
	VBReg	97.12	0.43	9.79	0.20	95.86	0.41	9.73	0.21
	MHopReg	**99.82**	0.39	8.70	0.12	98.74	0.30	17.73	0.12
	RANSAC10K	12.97	1.86	30.63	1.92	65.05	1.48	24.57	1.95
	RANSAC50K	26.49	1.72	24.28	9.37	76.58	1.26	20.90	9.48
	TEASER++	90.63	1.11	18.21	0.07	94.77	0.86	14.90	0.08
	SC^2-PCR	97.48	0.43	9.57	0.90	96.94	0.41	9.66	0.90
	FastMAC	93.69	0.82	13.20	0.63	71.89	0.92	14.22	0.67
	MAC	97.48	0.41	8.62	3.91	97.12	0.36	7.99	0.99
	TCF	97.84	0.53	11.85	0.07	96.22	0.61	13.01	0.39
	MAC++	98.02	0.77	15.11	9.49	95.86	0.45	9.97	2.86
	TurboReg	97.48	0.45	8.67	0.03	97.84	0.38	7.80	0.03
	PointTruss	99.64	0.43	**5.31**	0.21	**99.10**	0.59	11.06	0.21
	DualReg	98.20	0.28	7.59	0.12	96.58	**0.27**	**6.65**	0.36
	G2M-Reg (ours)	**99.82**	**0.24**	7.51	0.47	98.74	**0.27**	15.98	0.44

99.82% RR and 0.24° RE with FPFH, and 98.74% RR and 0.27° RE with FCGF. Runtime remains below 0.5 s across all settings. Overall, the largest gains on 3DLoMatch support the advantage of separating correspondence-space hypothesis proposal from surface-space pose verification under challenging low-overlap conditions.

4.3 Ablation Study

Table 2 evaluates the five key designs of G2M-Reg on 3DLoMatch using FPFH and FCGF correspondences. Introducing the pose-conditioned cross-cloud manifold improves RR from 43.18% to 44.96% with FPFH and from 62.21% to 63.36% with FCGF. Bidirectional consensus further increases RR to 45.82% and 64.08%, respectively, demonstrating the benefit

Table 2: Component-wise ablation of G2M-Reg on 3DLoMatch using FPFH and FCGF correspondences. Each pair of rows forms a controlled comparison.

No.	Design	Variant	FPFH			FCGF		
			RR↑ (%)	RE↓ (°)	TE↓ (cm)	RR↑ (%)	RE↓ (°)	TE↓ (cm)
1	Cross-Cloud Manifold	Graph	43.18	3.42	9.42	62.21	3.35	9.70
		Manifold	44.96	3.31	9.18	63.36	3.26	9.54
2	Bidirectional Consensus	One-way	44.96	3.31	9.18	63.36	3.26	9.54
		Bidirectional	45.82	3.27	9.10	64.08	3.28	9.57
3	Structural Reliability	w/o	46.72	3.10	8.98	64.47	3.12	9.34
		w/	47.89	3.16	9.15	65.69	3.16	9.49
4	Grouped Evidence Fusion	Raw	46.86	3.29	9.31	64.58	3.24	9.62
		Grouped	47.89	3.16	9.15	65.69	3.16	9.49
5	Surface-Guided Refinement	w/o	47.16	3.28	9.27	65.02	3.22	9.61
		w/	47.89	3.16	9.15	65.69	3.16	9.49

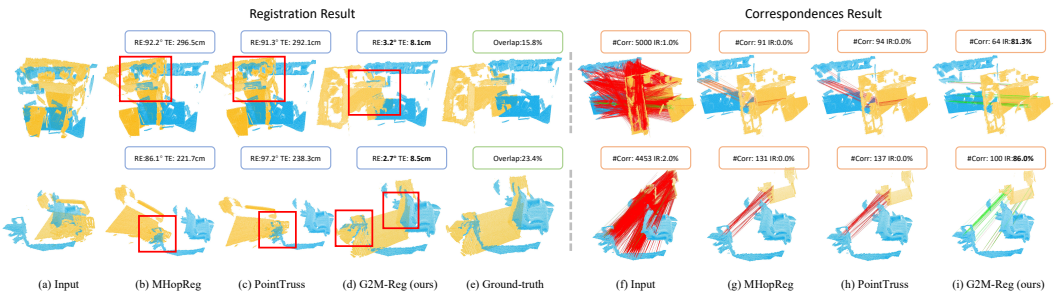

Figure 3: Qualitative results on challenging 3DLoMatch pairs. Correspondence inliers and outliers are shown in green and red, respectively. G2M-Reg rejects graph-supported but surface-invalid hypotheses and selects surface-valid poses.

of mutual surface verification over one-way evidence. Structural reliability provides additional gains of 1.17 and 1.22 percentage points in RR under FPFH and FCGF, respectively, indicating that coherent and well-constrained geometric support is important for robust pose selection. Grouped rank-space fusion also consistently outperforms raw-score fusion, showing the advantage of integrating heterogeneous evidence in rank space. Finally, surface-guided refinement further improves both registration recall and pose accuracy. Overall, the complete G2M-Reg achieves the highest RR of 47.89% and 65.69% under FPFH and FCGF, confirming the complementary roles of the proposed components.

4.4 Qualitative Results

Figure 3 compares G2M-Reg with strong learning-based and learning-free baselines on challenging 3DLoMatch pairs. Under repetitive low-overlap geometry, the baselines can select locally plausible but incorrect poses, whereas G2M-Reg rejects such surface-invalid hypotheses and recovers alignments close to the ground truth.

5 Conclusion

We presented G2M-Reg, a training-free Graph-to-Manifold Verification framework that separates correspondence-space pose proposal from surface-space pose verification. G2M-Reg combines pose-conditioned surface evidence, bidirectional partial-overlap consensus, structural reliability, and rank-space evidence fusion to select surface-valid hypotheses, followed by constrained refinement. Experiments across indoor and outdoor benchmarks and diverse correspondence sources demonstrate consistent gains, particularly under challenging low-overlap conditions.

AI use statement

Generative AI tools were used solely to assist with language polishing and improving the clarity of the manuscript. All AI-assisted text was reviewed and revised by the authors. The authors take full responsibility for the final content of this work.

Ethics statement

This work uses publicly available point cloud benchmarks and does not involve human subjects, personal data, or other sensitive information. We are not aware of any specific ethical concerns associated with this work.

Reproducibility statement

Implementation details, hyperparameter settings, evaluation protocols, and additional experiments are provided in the supplementary material. The anonymized source code, including configuration files and evaluation scripts, is also included in the supplementary submission.

References

Xuyang Bai, Zixin Luo, Lei Zhou, Hongkai Chen, Lei Li, Zeyu Hu, Hongbo Fu, and Chiew-Lan Tai. Pointdsc: Robust point cloud registration using deep spatial consistency. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 15854–15864, 2021. URL https://api.semanticscholar.org/CorpusID:232168842.

Tim Bailey and Hugh Durrant-Whyte. Simultaneous localization and mapping (slam): Part ii. IEEE robotics & automation magazine, 13(3):108–117, 2006.

Daniel Barath and Jiri Matas. Graph-cut ransac. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6733–6741, 2017. URL https://api.semanticscholar.org/CorpusID:32145255.

Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In Sensor fusion IV: control paradigms and data structures, volume 1611, pp. 586–606. Spie, 1992.

Eric Brachmann and Carsten Rother. Neural-guided ransac: Learning where to sample model hypotheses. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4322–4331, 2019.

Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. Image and vision computing, 10(3):145–155, 1992.

Zhi Chen, Kun Sun, Fan Yang, and Wenbing Tao. Sc2-pcr: A second order spatial compatibility for efficient and robust point cloud registration. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 13211–13221, 2022. URL https://api.semanticscholar.org/CorpusID:252571579.

Zhi Chen, Kun Sun, Fan Yang, Lin Guo, and Wenbing Tao. Sc2-pcr++: Rethinking the generation and selection for efficient and robust point cloud registration. IEEE transactions on pattern analysis and machine intelligence, 45(10):12358–12376, 2023.

Christopher Choy, Wei Dong, and Vladlen Koltun. Deep global registration. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 2514–2523, 2020.

Christopher Bongsoo Choy, Jaesik Park, and Vladlen Koltun. Fully convolutional geometric features. 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 8957–8965, 2019. URL https://api.semanticscholar.org/CorpusID:201686687.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM, 24:381–395, 1981. URL https://api.semanticscholar.org/CorpusID:972888.

Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 3354–3361, 2012. URL https://api.semanticscholar.org/CorpusID:6724907.

Shengyu Huang, Zan Gojcic, Mikhail (Misha) Usvyatsov, Andreas Wieser, and Konrad Schindler. Predator: Registration of 3d point clouds with low overlap. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4265–4274, 2020. URL https://api.semanticscholar.org/CorpusID:227209593.

Pablo Jaramillo and Ivan Sipiran. Cultural heritage 3d reconstruction with diffusion networks. ArXiv, abs/2410.10927, 2024. URL https://api.semanticscholar.org/CorpusID:273350603.

Haobo Jiang, Zheng Dang, Zhen Wei, Jin Xie, Jian Yang, and Mathieu Salzmann. Robust outlier rejection for 3d registration with variational bayes. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1148–1157, 2023. URL https://api.semanticscholar.org/CorpusID:257921545.

Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. Acta Crystallographica Section A, 32:922–923, 1976. URL https://api.semanticscholar.org/CorpusID:97383637.

Jiayi Li, Yuxin Yao, Qiuhang Lu, and Juyong Zhang. Dualreg: Dual-space filtering and reinforcement for rigid registration. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2026.

Jiayuan Li, Qingwu Hu, and Mingyao Ai. Gesac: Robust graph enhanced sample consensus for point cloud registration. ISPRS Journal of Photogrammetry and Remote Sensing, 2020a. URL https://api.semanticscholar.org/CorpusID:225238083.

Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A Chapman, Dongpu Cao, and Jonathan Li. Deep learning for lidar point clouds in autonomous driving: A review. IEEE Transactions on Neural Networks and Learning Systems, 32(8):3412–3432, 2020b.

Guanze Liu, Yu Rong, and Lu Sheng. Votehmr: Occlusion-aware voting network for robust 3d human mesh recovery from partial point clouds. Proceedings of the 29th ACM International Conference on Multimedia, 2021. URL https://api.semanticscholar.org/CorpusID:239011881.

Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, Slobodan Ilic, Dewen Hu, and Kai Xu. Geotransformer: Fast and robust point cloud registration with geometric transformer. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45:9806–9821, 2023. URL https://api.semanticscholar.org/CorpusID:257646283.

Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. 2009 IEEE International Conference on Robotics and Automation, pp. 3212–3217, 2009. URL https://api.semanticscholar.org/CorpusID:15022990.

Samuele Salti, Federico Tombari, and Luigi di Stefano. Shot: Unique signatures of histograms for surface and texture description. Comput. Vis. Image Underst., 125:251–264, 2014. URL https://api.semanticscholar.org/CorpusID:17799553.

Aleksandr V. Segal, Dirk Hähnel, and Sebastian Thrun. Generalized-icp. In Robotics: Science and Systems Conference, 2009. URL https://api.semanticscholar.org/CorpusID:231748613.

Pengcheng Shi, Shaocheng Yan, Yilin Xiao, Xinyi Liu, Yongjun Zhang, and Jiayuan Li. Ransac back to sota: A two-stage consensus filtering for real-time 3d registration. IEEE Robotics and Automation Letters, 9(12):11881–11888, 2024. doi: 10.1109/LRA.2024.3502056.

Inam Ullah, Deepak Adhikari, Habib Khan, Muhammad Shahid Anwar, Shabir Ahmad, and Xiaoshan Bai. Mobile robot localization: Current challenges and future prospective. Comput. Sci. Rev., 53:100651, 2024. URL https://api.semanticscholar.org/CorpusID:271020621.

Tong Wei, Yash J. Patel, Alexander Shekhovtsov, Juan E. Sala Matas, and Daniel Barath. Generalized differentiable ransac. 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 17603–17614, 2022. URL https://api.semanticscholar.org/CorpusID:257557687.

Yue Wu, Jun Jiang, Yongzhe Yuan, Maoguo Gong, Qiguang Miao, Hao Li, Mingyang Zhang, and Wenping Ma. Pointtruss: K-truss for point cloud registration. In Advances in Neural Information Processing Systems, volume 38, 2025.

Yue Wu, Feng Xiao, Yongzhe Yuan, Hao Li, Kaiyuan Feng, Maoguo Gong, Qiguang Miao, and Wenping Ma. Mhopreg: Efficient hierarchical multi-hop graph search for point cloud registration. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 24217–24226, 2026.

Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. ArXiv, abs/2403.14621, 2024. URL https://api.semanticscholar.org/CorpusID:268554137.

Shaocheng Yan, Pengcheng Shi, Zhenjun Zhao, Kaixin Wang, Kuang Cao, Ji Wu, and Jiayuan Li. Turboreg: Turboclique for robust and efficient point cloud registration. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 26371–26381, 2025a.

Shaocheng Yan, Yiming Wang, Kaiyan Zhao, Pengcheng Shi, Zhenjun Zhao, Yongjun Zhang, and Jiayuan Li. Hemora: Unsupervised heuristic consensus sampling for robust point cloud registration. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1363–1373, 2025b.

Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. IEEE Transactions on Robotics, 37(2):314–333, 2020.

Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38:2241–2254, 2016. URL https://api.semanticscholar.org/CorpusID:689166.

Jiaqi Yang, Zhiqiang Huang, Siwen Quan, Qian Zhang, Yanning Zhang, and Zhiguo Cao. Toward efficient and robust metrics for ransac hypotheses and 3d rigid registration. IEEE Transactions on Circuits and Systems for Video Technology, 32(2):893–906, 2021.

Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas A. Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 199–208, 2016. URL https://api.semanticscholar.org/CorpusID:11446141.

Xiyu Zhang, Jiaqi Yang, Shikun Zhang, and Yanning Zhang. 3d registration with maximal cliques. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 17745–17754, 2023. URL https://api.semanticscholar.org/CorpusID:258762823.

Xiyu Zhang, Yanning Zhang, and Jiaqi Yang. Mac++: Going further with maximal cliques for 3d registration. In 2025 International Conference on 3D Vision (3DV), pp. 261–275. IEEE, 2025. doi: 10.1109/3DV66043.2025.00029.

Yifei Zhang, Hao Zhao, Hongyang Li, and Siheng Chen. Fastmac: Stochastic spectral sampling of correspondence graph. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 17857–17867, 2024. URL https://api.semanticscholar.org/CorpusID:268379338.

Guiyu Zhao, Sheng Ao, Ye Zhang, Kai Xu, and Yulan Guo. Progressive correspondence regenerator for robust 3d registration. In Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 1210–1219, 2025.

Chengyu Zheng, Jin Huang, Honghua Chen, and Mingqiang Wei. Rare: Refine any registration of pairwise point clouds via zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 26549–26558, 2025.

Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In European conference on computer vision, pp. 766–782. Springer, 2016.