TurboReg: TurboClique for Robust and Efficient Point Cloud Registration

Supplementary Material

In this supplementary material, we first provide additional analyses to elaborate on concepts introduced in the main paper. Specifically: (1) Suppl. A details why TurboClique employs a stringent compatibility threshold. (1.1) Suppl. A.1 elaborates on pairwise compatibility-induced stability. (1.2) Suppl. A.2 and Suppl. A.3 illustrate the application of pairwise compatibility-induced stability in the design of TurboClique through experimental validation and geometric intuition, respectively. (2) Suppl. B offers a detailed numerical interpretation of the SC^2 scores. (3) Suppl. C proves the Unique Assignment Property of TurboClique in the O2Graph. (4) Suppl. D presents the Tensor-style pseudocode for PGS.

Next, we introduce foundational concepts to enhance the completeness. (1) Suppl. E.1 defines the concepts of clique and maximal clique. (2) Suppl. E.2 provides the derivation of the variance of the Least Squares Estimator.

Finally, we present additional experiments. (1) Suppl. F.1 provides a deep understanding experiment for searched TurboClique (2) Suppl. F.3 provides the runtime of TurboReg components. (3) Suppl. F.4 analyzes the failure cases of TurboReg. (4) Suppl. F.5 visualizes qualitative examples not included in the main paper.

A. Why TurboClique using Stringent Compatibility Threshold ?

In this section, we introduce the rationale behind the stringent compatibility threshold employed by TurboClique. We first elaborate on the concept of pairwise compatibility-induced stability in Suppl. A.1. Furthermore, we demonstrate how this stability principle is incorporated into the design of TurboClique through experimental validation in Suppl. A.2 and geometric intuition in Suppl. A.3.

A.1. Pairwise Compatibility-induced Stability

This section explains pairwise compatibility-induced stability by analyzing the relationship between matching noise variance and the spatial compatibility constraint. The analysis demonstrates that smaller τ values enhance pairwise compatibility-induced stability, enabling 3-cliques to achieve stability comparable to larger cliques (e.g., maximal cliques). We also provide experimental and intuitive analyses for a comprehensive understanding.

Given a matching set $\mathcal{M} = \{\boldsymbol{m}_i\}_{i=1}^{N}$, where, $\boldsymbol{m}_i = (\boldsymbol{x}_i, \boldsymbol{y}_i)$ and $\boldsymbol{x}_i, \boldsymbol{y}_i \in \mathbb{R}^3$ represent source and target keypoints, respectively, the matching relationship is defined as $\boldsymbol{y}_i = \mathbf{T}(\boldsymbol{x}_i) + \boldsymbol{r}_i$. Here, $\mathbf{T}(\cdot)$ denotes the rigid transformation (including rotation and translation), and $\boldsymbol{r}_i$ represents noise in the matching process. We assume $\boldsymbol{r}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I}_{3\times3})$, indicating that the noise follows an independent, zero-mean, isotropic Gaussian distribution with variance σ^2.

In the absence of constraints between matches, the distribution of $\boldsymbol{r}_i$ remains entirely random, with noise uniformly distributed across three-dimensional space. We then analyze the influence of introducing spatial compatibility constraints, defined as follows for any $\boldsymbol{m}_i \in \mathcal{M}$:

$$\big|\|\boldsymbol{y}_i - \boldsymbol{y}_j\| - \|\boldsymbol{x}_i - \boldsymbol{x}_j\|\big| \leq \tau, \tag{1}$$

where $\tau \geq 0$ represents the compatibility threshold, limiting the distance difference between matching pairs. This constraint reduces the randomness of $\boldsymbol{r}_i$, narrows the noise distribution, and yields an effective variance σ^2_{eff} that is potentially smaller than the initial variance σ^2.

For a precise analysis, we define $\boldsymbol{d}_{ij} = \boldsymbol{x}_i - \boldsymbol{x}_j$ and $\boldsymbol{e}_{ij} = \boldsymbol{r}_i - \boldsymbol{r}_j$. Leveraging the distance-preserving property of rigid transformations, which ensures that $\|\mathbf{T}(\boldsymbol{x}_i) - \mathbf{T}(\boldsymbol{x}_j)\| = \|\boldsymbol{x}_i - \boldsymbol{x}_j\|$, we rewrite Eq. (1) as:

$$\big|\|\boldsymbol{d}_{ij} + \boldsymbol{e}_{ij}\| - \|\boldsymbol{d}_{ij}\|\big| \leq \tau. \tag{2}$$

Initially, $\boldsymbol{e}_{ij}$ follows a Gaussian distribution $\mathcal{N}(\mathbf{0}, 2\sigma^2\mathbf{I}_{3\times3})$. However, the constraint in Eq. (2) limits the norm of $\boldsymbol{e}_{ij}$, effectively truncating the joint distribution of $\boldsymbol{r}_i$ and $\boldsymbol{r}_j$. Consequently, the variance of this truncated distribution is smaller than that of the original, resulting in an effective variance $\sigma^2_{\text{eff}} < \sigma^2$. Specifically, Eq. (2) enforces $\|\boldsymbol{d}_{ij} + \boldsymbol{e}_{ij}\|$ to lie within the interval $[\|\boldsymbol{d}_{ij}\| - \tau, \|\boldsymbol{d}_{ij}\| + \tau]$. As τ decreases, this interval narrows, further restricting the possible values of $\boldsymbol{e}_{ij}$. In the limit where $\tau \to 0$, the constraint reduces to $\|\boldsymbol{d}_{ij} + \boldsymbol{e}_{ij}\| = \|\boldsymbol{d}_{ij}\|$, which geometrically implies that $\boldsymbol{e}_{ij} \to \mathbf{0}$. This condition suggests that $\boldsymbol{r}_i \approx \boldsymbol{r}_j$, and given the zero-mean property of $\boldsymbol{r}_i$, it follows that $\boldsymbol{r}_i \to \mathbf{0}$. As a result, the noise distribution approaches a Dirac delta function, with $\sigma^2_{\text{eff}} \to 0$. Thus, smaller values of τ progressively reduce σ^2_{eff}, ultimately approaching zero.

In summary, spatial compatibility reduces the randomness of $\boldsymbol{r}_i$, with its variance decreasing to zero as τ diminishes.

A.2. Experimental Validation

The analysis above suggests that as τ approaches zero, pairwise compatibility-induced stability increases, compensating for the loss of data scaling stability in TurboClique due to a reduced number of matches. Consequently, rigid transformations derived from 3-cliques exhibit minimal differences compared to those from larger cliques, supporting TurboClique's preference for a small τ. This section empirically

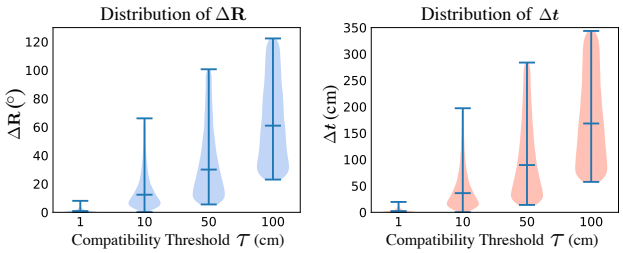

Figure 1. Distribution of discrepancies between transformations estimated from 3-clique and 10-clique configurations in terms of rotation and translation.

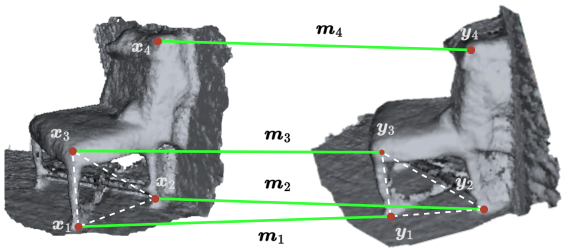

Figure 2. Demonstration of matches when $\tau = 0$.

confirms this inference by demonstrating that transformation discrepancies between 3-cliques and 10-cliques decrease as τ diminishes. Specifically, we assess transformation discrepancies between estimates derived from 3-cliques and 10-cliques on the 3DMatch+FPFH dataset across increasing τ values. The procedure is outlined as follows:

1. Set $\tau \in \{1\text{ cm}, 10\text{ cm}, 50\text{ cm}, 100\text{ cm}\}$ to construct compatibility graph $\mathbf{G}$;
2. Extract all 10-cliques from $\mathbf{G}$ and compute multiple transformations $\mathbf{T}^{(10)} = (\mathbf{R}^{(10)}, \boldsymbol{t}^{(10)})$;
3. For each 10-clique, estimate transformations $\{\mathbf{T}_k^{(3)}\}_{k=1}^{K} = \{(\mathbf{R}_k^{(3)}, \boldsymbol{t}_k^{(3)})\}$ from all $K = \binom{10}{3}$ 3-clique subsets;
4. Calculate rotation and translation errors:

$$\Delta\mathbf{R}_k = \arccos\left(\frac{\text{tr}(\mathbf{R}^{(10)\top}\mathbf{R}_k^{(3)}) - 1}{2}\right), \tag{3}$$

$$\Delta\boldsymbol{t}_k = \|\boldsymbol{t}^{(10)} - \boldsymbol{t}_k^{(3)}\|_2; \tag{4}$$

5. Visualize error distributions across τ values in Fig. 1.

Results show negligible discrepancies at $\tau = 1$ cm ($< 0.1°$ rotation, < 0.5 mm translation), with errors rising proportionally to τ. This confirms that 3-cliques achieve accuracy comparable to larger cliques under tight thresholds, while significantly reducing computational complexity.

A.3. Geometric Intuition

We further provide an intuitive analysis by examining the extreme case where $\tau = 0$. Three matches $\{\boldsymbol{m}_1, \boldsymbol{m}_2, \boldsymbol{m}_3\}$, as shown in Fig. 2, form congruent triangles $\triangle\boldsymbol{x}_1\boldsymbol{x}_2\boldsymbol{x}_3$ and $\triangle\boldsymbol{y}_1\boldsymbol{y}_2\boldsymbol{y}_3$, uniquely determining the rigid transformation $\mathbf{T}^{(3)}$. Introducing a fourth match $\boldsymbol{m}_4$ that satisfies $\tau = 0$ compatibility with the initial trio results in a transformation $\mathbf{T}^{(4)}$ identical to $\mathbf{T}^{(3)}$, as $\boldsymbol{m}_4$ must conform to the existing geometric constraints. This principle applies to additional matches: any correspondence satisfying $\tau = 0$ preserves the original transformation. Thus, under ideal compatibility conditions (i.e., $\tau = 0$), 3-cliques fully encapsulate transformation information, rendering larger cliques unnecessary. In practice, however, a small, non-zero τ is adopted to account for sensor noise and matching imperfections, justifying our use of a modest compatibility threshold.

B. Numerical Interpretation of SC² Scores

In Sec. 3.3 of the main paper, we claim that SC^2 scores quantify TurboClique density. We now provide a brief explanation. The SC^2 score between matches $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ is defined as:

$$\hat{\mathbf{G}}_{ij} = \mathbf{G}_{ij} \sum_{k=1}^{N} \mathbf{G}_{ik} \cdot \mathbf{G}_{jk}, \tag{5}$$

where $\mathbf{G}_{ij} \in \{0, 1\}$ indicates spatial compatibility between $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$. Two observations are as follows:

- If $\mathbf{G}_{ij} = 0$, then $\hat{\mathbf{G}}_{ij} = 0$, indicating no edge between $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$. Consequently, the number of TurboCliques around $(\boldsymbol{m}_i, \boldsymbol{m}_j)$ is zero.
- If $\mathbf{G}_{ij} = 1$, the summation counts nodes k where $\mathbf{G}_{ik} = \mathbf{G}_{jk} = 1$. Each such k forms a TurboClique $\{\boldsymbol{m}_i, \boldsymbol{m}_j, \boldsymbol{m}_k\}$, making $\hat{\mathbf{G}}_{ij}$ equal to the number of TurboCliques containing the edge (i, j).

Combining these cases, the value of $\hat{\mathbf{G}}_{ij}$ represents the number of TurboCliques associated with $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$.

C. Unique Assignment Property of TurboClique

In this section, we demonstrate how the O2Graph eliminates the redundant detection of TurboCliques by proving that each TurboClique can be uniquely assigned to a single pivot.

Given a TurboClique around π_z, denoeted as $\text{TC}(\pi_z) = \{\boldsymbol{m}_{z_1}, \boldsymbol{m}_{z_2}, \boldsymbol{m}_{z_3}\}$, where $z_1 < z_2 < z_3$ (without loss of generality), the O2Graph defines edge directions from lower-indexed to higher-indexed nodes. This implies:

- $\mathcal{N}(\boldsymbol{m}_{z_1}) = \{\boldsymbol{m}_{z_2}, \boldsymbol{m}_{z_3}\}$,
- $\mathcal{N}(\boldsymbol{m}_{z_2}) = \{\boldsymbol{m}_{z_3}\}$,
- $\mathcal{N}(\boldsymbol{m}_{z_3}) = \emptyset$,

where $\mathcal{N}(\cdot)$ denotes the set of neighboring nodes in the compatibility graph.

Next, we analyze three possible pivot cases to show that only one case detects $\text{TC}(\pi_z)$:

Algorithm 1: Pivot-Guided Search Algorithm (Tensor-style)

Input: Weighted graph: $\bar{\mathbf{G}} \in \mathbb{R}^{N \times N}$; number of pivots $K_1 \in \mathbb{N}^+$; number of TurboCliques for each pivot $K_2 \in \mathbb{N}^+$
Output: TurboClique set $\mathbf{C} \in \{1, \ldots, N\}^{K_1 K_2 \times 3}$
% Select top-K_1 edges as pivots
$\mathbf{P} \leftarrow \text{TopKEdges}(\bar{\mathbf{G}}, K_1)$
% Common neighbors (mask) for each pivot
$\mathbf{M} \leftarrow (\bar{\mathbf{G}}[\mathbf{P}[:, 0]] > 0) \odot (\bar{\mathbf{G}}[\mathbf{P}[:, 1]] > 0)$
% TurboClique weights for each TurboClique
$\mathbf{S} \leftarrow \bar{\mathbf{G}}[\mathbf{P}[:, 0], \mathbf{P}[:, 1]] + (\bar{\mathbf{G}}[\mathbf{P}[:, 0]] + \bar{\mathbf{G}}[\mathbf{P}[:, 1]])$
$\mathbf{S}' \leftarrow \mathbf{S} \odot \mathbf{M}$
% Top-K_2 TurboCliques for each pivot
$\mathbf{Z} \leftarrow \text{ColumnTopK}(\mathbf{S}', K_2)$
% Assemble TurboCliques: (pivots, third matches)
$\mathbf{C} \leftarrow \text{zeros}(K_1 \cdot K_2, 3)$
for $i \leftarrow 0$ *to* K_2 **do**
    % Assign first two matches
    $\mathbf{C}[(i \times K_1) : ((i+1) \times K_1), : 2] \leftarrow \mathbf{P}$
    % Assign third match
    $\mathbf{C}[(i \times K_1) : ((i+1) \times K_1), 2] \leftarrow \mathbf{Z}$
end
return $\mathbf{C}$

- **Case 1:** $\pi_z = (\boldsymbol{m}_{z_2}, \boldsymbol{m}_{z_3})$: Since $\boldsymbol{m}_{z_1} \notin \mathcal{N}(\boldsymbol{m}_{z_2})$ and $\boldsymbol{m}_{z_1} \notin \mathcal{N}(\boldsymbol{m}_{z_3})$, this pivot cannot detect TC(π_z).
- **Case 2:** $\pi_z = (\boldsymbol{m}_{z_1}, \boldsymbol{m}_{z_3})$: Since $\boldsymbol{m}_{z_2} \notin \mathcal{N}(\boldsymbol{m}_{z_3})$, this pivot cannot form TC(π_z) with $\boldsymbol{m}_{z_2}$.
- **Case 3:** $\pi_z = (\boldsymbol{m}_{z_1}, \boldsymbol{m}_{z_2})$: Here, $\boldsymbol{m}_{z_3} \in \mathcal{N}(\boldsymbol{m}_{z_1}) \cap \mathcal{N}(\boldsymbol{m}_{z_2})$, enabling the formation of TC(π_z).

Since any three matches can form at most the three above pivot configurations, and only the pivot consisting of the two lowest-indexed nodes detects a TurboClique, this proves that each TurboClique is uniquely assigned to a single pivot.

D. Tensor-style Pseudo-code of PGS

We present the Tensor-style pseudo-code of the PGS algorithm in Algorithm 1.

E. Supporting Theorems and Derivations

To ensure the completeness of this paper, this section provides foundational theorems and derivations that support the main analysis.

E.1. Definition of Clique and Maximal Clique

Figure Fig. 3 depicts a graph with 7 vertices, denoted as $\mathcal{G}$. A clique is a complete subgraph $\mathcal{C} \subseteq \mathcal{G}$ where every pair of distinct vertices is adjacent:

$$\forall \boldsymbol{u}, \boldsymbol{v} \in \mathcal{C}, (\boldsymbol{u}, \boldsymbol{v}) \in E(\mathcal{G}), \tag{6}$$

where $E(\mathcal{G})$ represents the edge set of $\mathcal{G}$. For example, the vertices $\{1, 3, 5\}$ are fully connected, forming a 3-clique. Similarly, $\{2, 4, 6, 7\}$ constitutes a 4-clique. For example, the vertices $\{1, 3, 5\}$ are fully connected, forming a 3-clique. Similarly, $\{2, 4, 6, 7\}$ forms a 4-clique.

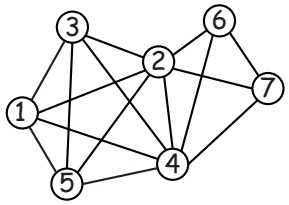

Figure 3. Undirected graph for demonstration.

The maximal clique is defined as a clique that cannot be extended by including any adjacent vertex:

$$\nexists \boldsymbol{w} \in \mathcal{G} \setminus \mathcal{C} \text{ such that } \mathcal{C} \cup \{\boldsymbol{w}\} \text{ forms a clique.} \tag{7}$$

For instance, the 5-clique $\{1, 2, 3, 4, 5\}$ is maximal because the remaining vertices $\{6, 7\}$ cannot be added to form a larger clique. Similarly, $\{2, 4, 6, 7\}$ is a maximal 4-clique.

E.2. Variance of LS Estimator

This section derives the variance of the least squares (LS) estimator in a standard linear regression framework. Consider the linear regression model:

$$Y = X\beta + \epsilon, \tag{8}$$

where Y is the response variable, X is the design matrix, β is the coefficient vector, and ϵ is the error term. We assume the errors satisfy:

$$\mathbb{E}[\epsilon|X] = 0, \quad \text{Var}(\epsilon|X) = \sigma^2 I_n. \tag{9}$$

The ordinary least squares (OLS) estimator for β is:

$$\hat{\beta} = (X'X)^{-1}X'Y. \tag{10}$$

The variance of $\hat{\beta}$ is computed as:

$$\text{Var}(\hat{\beta}|X) = \text{Var}\left((X'X)^{-1}X'Y \mid X\right). \tag{11}$$

Substituting $Y = X\beta + \epsilon$ and applying variance properties:

$$\text{Var}(\hat{\beta}|X) = (X'X)^{-1}X'\text{Var}(\epsilon|X)X(X'X)^{-1}. \tag{12}$$

Given $\text{Var}(\epsilon|X) = \sigma^2 I_n$, this simplifies to:

$$\text{Var}(\hat{\beta}|X) = \sigma^2(X'X)^{-1}. \tag{13}$$

This result indicates that the variance of the LS estimator depends on the noise variance σ^2 and the design matrix X. Notably, a smaller σ^2 or a larger sample size (reflected in X) reduces the variance.

	Metrics	RR (%)	TQRR (%)	ICRR (%)	TKRR (%)			
					@2	@3	@5	@50
FPFH	IN	84.10	93.72	70.48	85.97	86.50	87.77	92.38
	MSE	82.99	99.94	68.90	83.77	84.43	85.34	90.04
	MAE	83.43	99.94	69.26	84.22	84.87	85.79	90.51
FCGF	IN	93.59	97.66	90.24	94.06	94.45	94.91	96.66
	MSE	93.47	99.94	89.73	93.73	93.86	94.24	95.86
	MAE	93.35	99.94	89.84	93.85	93.98	94.37	95.98

Table 1. Ranking-based Registration Recall Evaluation on the 3DMatch Dataset. (1) TQRR evaluates whether the best transformation outperforms the ground truth transformation under the corresponding metrics. (2) ICRR assesses whether the best TurboClique hypothesis consists of three inliers. (3) TKRR determines whether the top-K hypotheses include a successfully registered rigid transformation.

Device	Methods	O2Graph Construction	PGS	Model Estimation	Total
CPU	Ours (0.5)	276.05 (88.33%)	30.33 (9.70%)	6.13 (1.96%)	312.50
	Ours (2K)	277.26 (67.02%)	73.26 (17.71%)	63.17 (15.27%)	413.68
GPU	Ours (0.5)	0.04 (0.25%)	11.36 (71.71%)	4.44 (28.05%)	15.84
	Ours (2K)	0.05 (0.25%)	11.88 (60.80%)	7.61 (38.95%)	19.54

Table 2. Average consumed time (ms) per point cloud pair on the 3DMatch+FPFH dataset across CPU and GPU implementations.

#hypotheses	3DMatch				3DLoMatch			
	FPFH		FCGF		FPFH		FCGF	
	3DMAC	Ours	3DMAC	Ours	3DMAC	Ours	3DMAC	Ours
100	50.67	**78.39**	61.92	**90.67**	12.22	**23.19**	30.47	**52.11**
200	89.27	**151.12**	119.20	**178.99**	17.59	**37.34**	55.57	**97.87**
500	162.41	**346.03**	269.06	**429.54**	23.32	**45.43**	109.32	**206.49**
1000	217.32	**598.01**	456.18	**777.29**	26.02	**63.33**	156.11	**316.24**
2000	254.13	**770.39**	669.32	**1034.39**	29.31	**78.34**	202.12	**362.05**

Table 3. Comparison of correct hypothesis counts

F. More Experiments

F.1. Understanding the Searched TurboCliques

To better understand TurboReg, we propose three ranking-based registration recall metrics to analyze the K_1K_2 TurboCliques identified by PGS. Specifically, we first rank the K_1K_2 transformation hypotheses based on inlier number (IN), mean absolute error (MAE), and mean squared error (MSE). The three metrics are defined as follows: (1) Transformation Quality Registration Recall (TQRR): The proportion of cases where the top-1 hypothesis achieves a score equal to or exceeding the ground-truth transformation. Inlier-Clique Registration Recall (ICRR): The proportion of cases where the top-1 hypothesis clique contains only inliers. (2) Top-K Hypothesis Registration Recall (TKRR): The proportion of cases where at least one valid transformation exists among the top-K hypotheses. Results are summarized in Table 1.

Discussion of TQRR. From Tab. 1, TQRR consistently exceeds RR by significant margins. For instance, when ranked by IN, TQRR surpasses RR by 9.62%, indicating that erroneous rigid transformations with higher consistency scores than the ground truth are frequently selected during model estimation. This suggests that the ground-truth transformation does not always align with the maximum consistency assumption, potentially due to: (1) Sampling Error: Discrete keypoint sampling or insufficient sampling density causing deviations between the ground truth and maximum consistency transformations. (2) Scene Ambiguity: Repetitive structures (e.g., identical objects) leading to ambiguous alignments.

Notably, TurboReg achieves 99.94% (1622/1623) TQRR under MAE and MSE metrics, demonstrating its ability to prioritize highly consistent hypotheses over the ground truth in nearly all cases.

Discussion of ICRR. Tab. 1 reveals that ICRR is consistently lower than RR, indicating that many cliques containing outliers still produce successful registrations. These findings demonstrate that even cliques with outliers can yield correct registrations.

Discussion of TKRR. As K increases, TKRR improves significantly and eventually surpasses RR (Table 1). This indicates that the correct transformation is more likely to reside among the top-K transformations rather than exclusively in the top-1. However, conventional methods typically select the top-1 transformation based on ranking, implying that our model selection strategy may impose performance limitations. This reliance on top-1 selection often overlooks potentially correct transformations within the broader top-K set, highlighting a key bottleneck in registration.

Summary. In summary, we demonstrate that TurboReg excels at identifying TurboCliques with a high inlier ratio, characterized by high IN and lower MSE/MAE. However, the correct transformation may not always be selected due to the inherent limitation of choosing only the top candidate, which constrains the overall performance of the registration algorithm despite its ability to generate high-scoring cliques.

F.2. Comparison with MAC hypotheses.

Following [1], we evaluate the quality of the generated hypotheses by comparing those produced by MAC and TurboReg against the ground truth transformation. The results, shown in Tab. 3, indicate that under the same number of hypotheses, our method yields a higher proportion of correct hypotheses.

F.3. Runtime of TurboReg Components

This experiment investigates the temporal characteristics of TurboReg modules on the 3DMatch+FPFH dataset. Average execution times (ms) for CPU and GPU implementations are presented in Tab. 2.

Significant differences exist between CPU and GPU implementations. Focusing first on the CPU variant, the O2Graph Construction module dominates the processing time under both 0.5K and 2K pivot configurations. The PGS

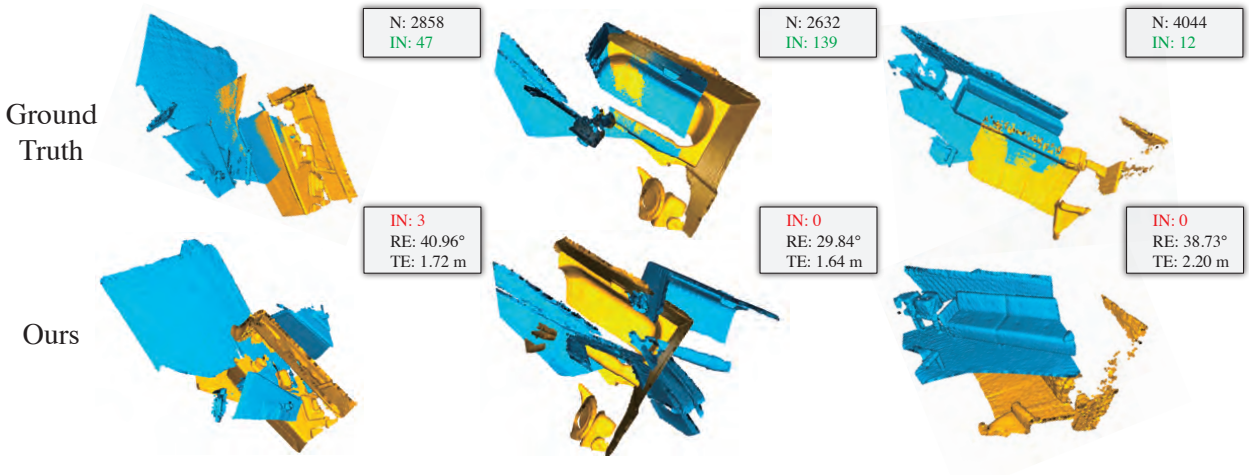

Figure 4. **Insufficient Consensus Correspondences.** Red indicates lower IN values, while green denotes higher IN values.

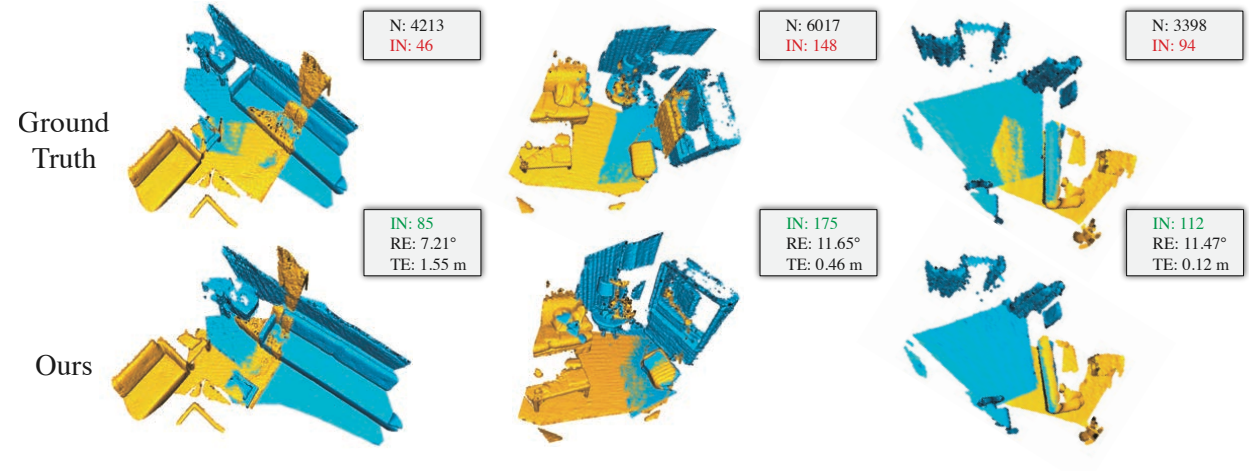

Figure 5. **Larger Consensus Set with Small Errors.** Red indicates lower IN values, while green denotes higher IN values.

and Model Estimation modules exhibit a positive correlation between K_1 and runtime, since an increase in K_1 leads to a higher number of TurboCliques.

In GPU implementations, the O2Graph Construction time decreases drastically (e.g., merely 0.25% of total runtime) due to parallel computation capabilities. Furthermore, the parallelized TurboClique search enables the PGS module to maintain near-constant execution time, resulting in approximately 12 ms for both $K_1 = 500$ and 2000. Conversely, the Model Estimation module demonstrates a linear scaling trend with K_1, as additional TurboCliques necessitate incremental transformation estimations.

F.4. Failure Case Analysis

In this section, we analyze the failure cases of TurboReg. We first review the definition of successful registration: registration is successful if the error between the estimated rigid transformation and the ground truth rigid transformation falls below specific thresholds. For 3DMatch and 3DLoMatch, the requirements are RE $\leq 15^\circ$ and TE ≤ 30 cm. For the KITTI dataset, the requirements are RE $\leq 5^\circ$ and TE ≤ 60 cm.

Next, we note that the estimated rigid transformation is selected based on the inlier number (IN), under the assumption that the correct rigid transformation corresponds to the maximum consensus set.

We classify instances that do not meet the successful registration criteria into three categories:

1. **Insufficient Consensus Correspondences**: TurboReg fails to identify a sufficiently large set of consensus correspondences. This occurs in scenarios with extremely low overlap or strong symmetry, as illustrated in Fig. 4.
2. **Larger Consensus Set but Incorrect Transformation**: The algorithm identifies a larger IN than that of the ground truth transformation, yet the result remains in-

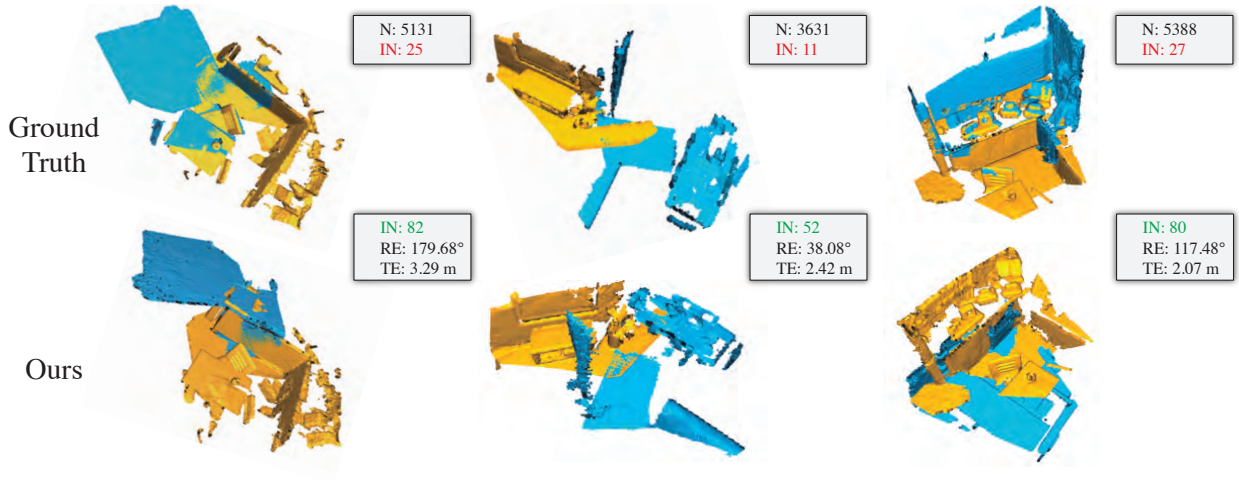

Figure 6. **Larger Consensus Set with Large Errors.** Red indicates lower IN values, while green denotes higher IN values.

correct. This contradicts the maximum consensus set assumption. We categorize this scenario into two subcategories:

(a) **Small Errors**: The estimated rigid transformation closely approximates the true rigid transformation, suggesting that registration is feasible, albeit with slightly larger errors. Due to the limited number of correct matches, the result is sensitive to noise, as illustrated in Fig. 5.

(b) **Large Errors**: The algorithm identifies a rigid transformation with a larger inlier set that still aligns visually, as shown in Fig. 6. This may occur because the matching pairs conform to an underlying geometric structure.

F.5. Qualitative Visualizations

Figs. 7-9 illustrate qualitative visualizations of challenging registration pairs. 3DMAC and SC^2-PCR fail to achieve registration, whereas TurboReg successfully completes the registration task.

References

[1] Xiyu Zhang, Jiaqi Yang, Shikun Zhang, and Yanning Zhang. 3d registration with maximal cliques. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17745–17754, 2023. 4

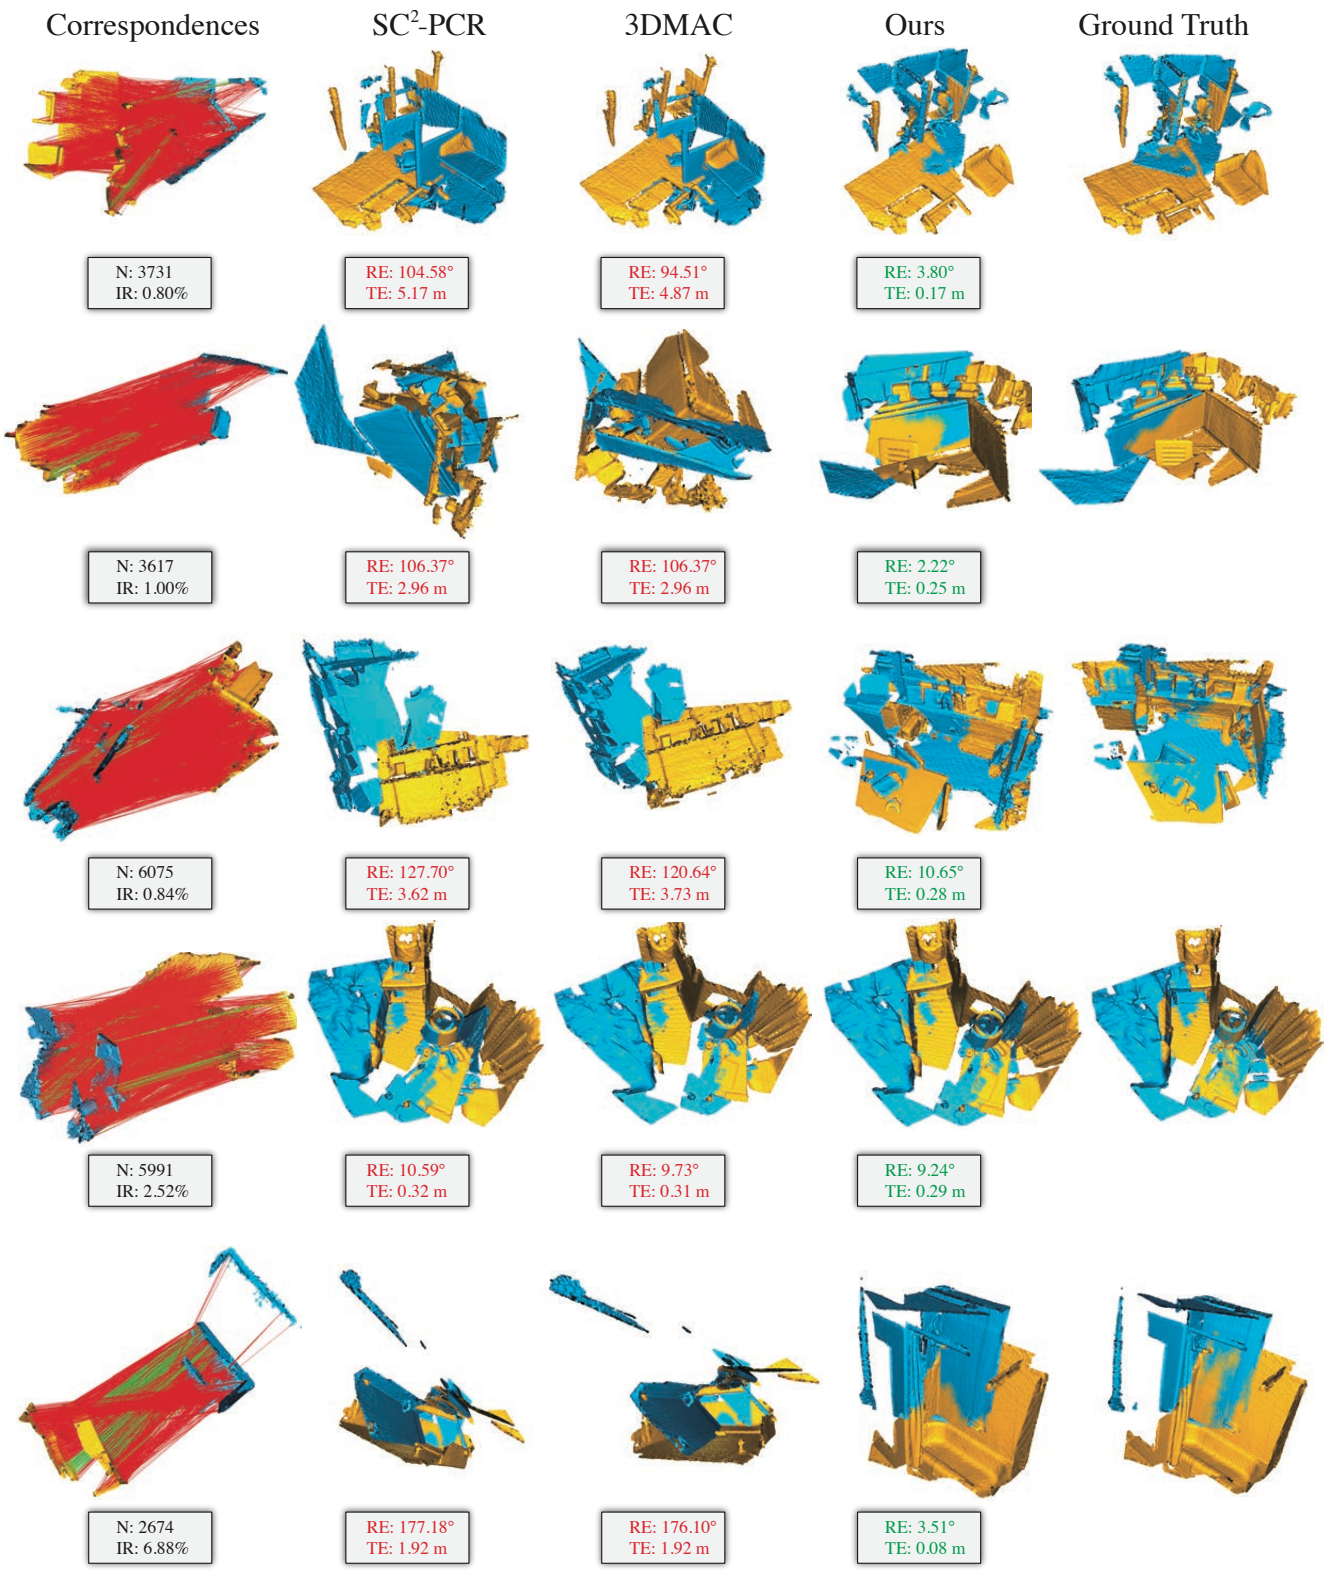

Figure 7. **Qualitative Comparison on 3DMatch.** Red and green represent failed and successful registrations, respectively.

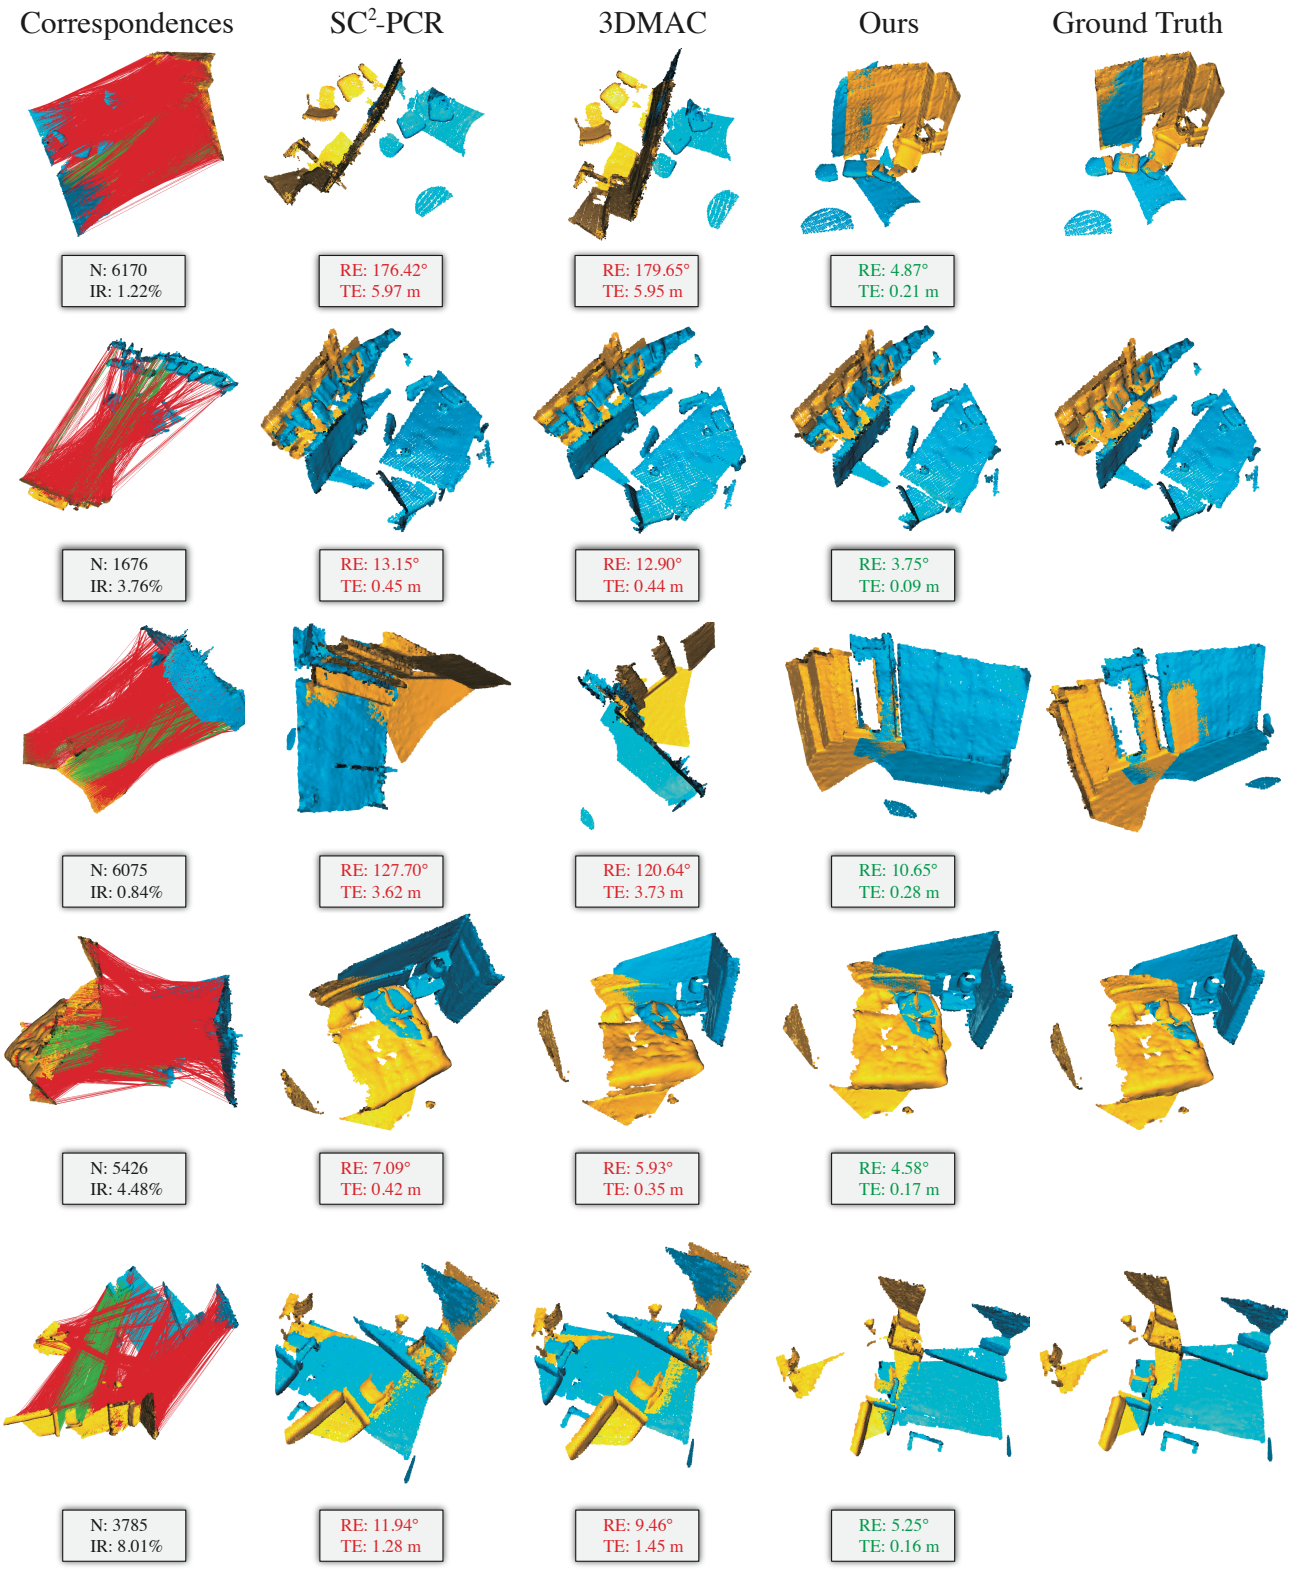

Figure 8. **Qualitative Comparison on 3DLoMatch.** Red and green represent failed and successful registrations, respectively.

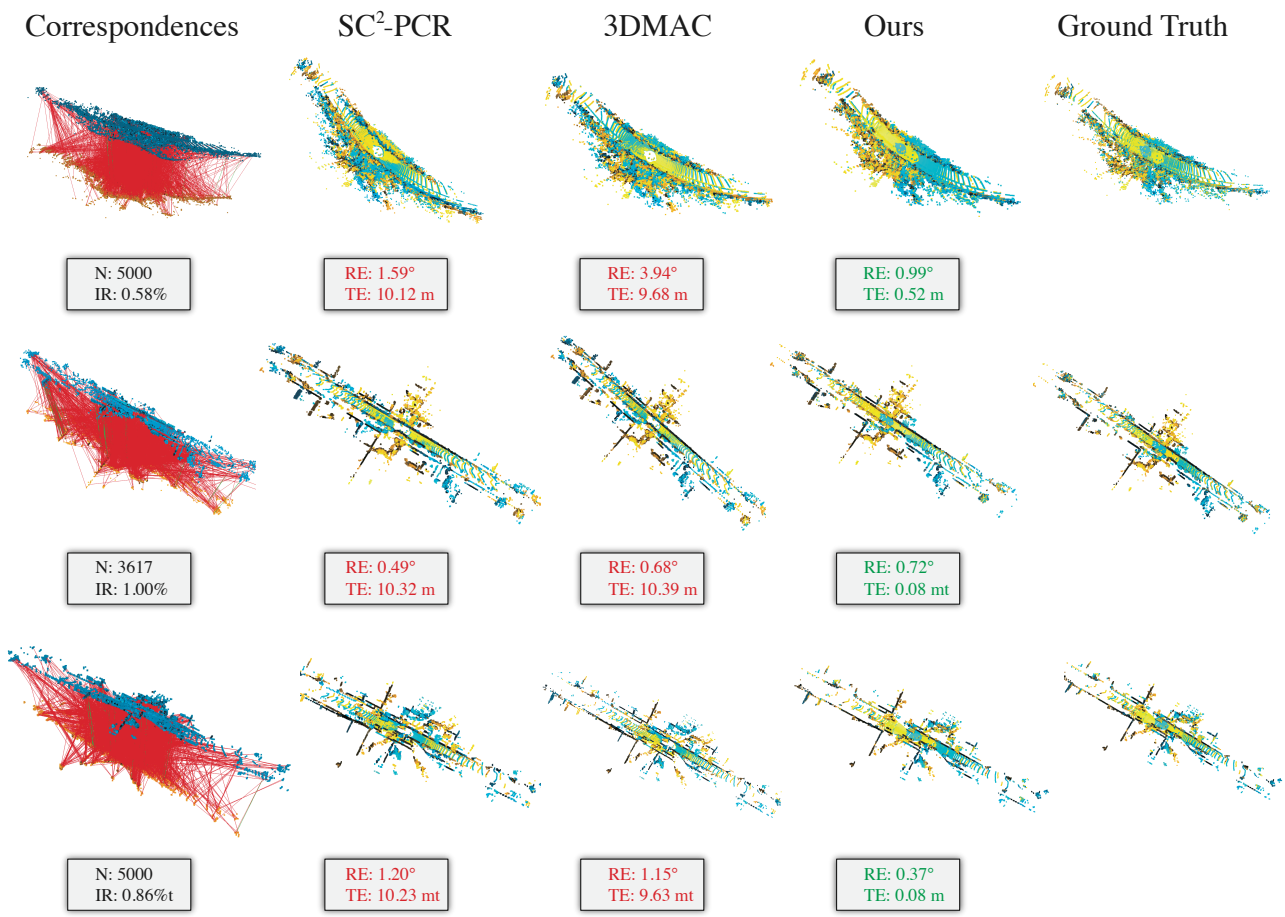

Figure 9. **Qualitative comparison on KITTI.** Red and green represent failed and successful registrations, respectively.