

TOPO-Bench: An Open-Source Topological Mapping Evaluation Framework with Quantifiable Perceptual Aliasing

Jiaming Wang*, Diwen Liu*, Jizhuo Chen*, Harold Soh

Abstract—Topological mapping offers a compact and robust representation for navigation, but progress in the field is hindered by the lack of standardized evaluation metrics, datasets, and protocols. Existing systems are assessed using different environments and criteria, preventing fair and reproducible comparisons. Moreover, a key challenge—perceptual aliasing—remains under-quantified, despite its strong influence on system performance. We address these gaps by (i) formalizing *topological consistency* as the fundamental property of topological maps and showing that localization accuracy provides an efficient and interpretable surrogate metric, and (ii) proposing the first quantitative measure of dataset ambiguity to enable fair comparisons across environments. To support this protocol, we curate a diverse benchmark dataset with calibrated ambiguity levels, implement and release deep-learned baseline systems, and evaluate them alongside classical methods. Our experiments and analysis yield new insights into the limitations of current approaches under perceptual aliasing. All datasets, baselines, and evaluation tools are fully open-sourced to foster consistent and reproducible research in topological mapping.

I. INTRODUCTION

Topological maps represent the environment as a graph, where nodes correspond to discrete places and edges capture navigable connections between them. This discretization of the continuous world provides a compact and efficient representation, since only salient locations are stored and dense metric reconstructions are unnecessary. As a result, topological maps are far less computationally demanding than metric maps. Their efficiency is particularly relevant in the context of emerging mapless navigation systems [1], [2], which operate using sparse subgoals that can be obtained through graph search on a topological map without relying on a full metric reconstruction.

Topological navigation is not only a useful abstraction in robotics but also a behavior observed in many animal species, including humans. In everyday movement, we rarely think in terms of exact coordinates; instead, we rely on abstract notions of distance and spatial relationships, recognizing where we are without requiring fine-grained metric accuracy. As argued by Brooks [3], topological representations naturally help systems cope with uncertainty in mobile robot navigation. By replacing explicit metric information with notions of proximity and ordering, they avoid the accumulation of dead-reckoning errors that often undermine purely geometric approaches.

However, a persistent challenge in topological mapping is *perceptual aliasing* [4], which is the risk that two physically

distinct locations produce observations that appear identical to the robot’s sensors. Perceptual aliasing can lead to incorrect loop closures and navigation failures, making its mitigation a central research question.

Despite the large number of topological mapping systems proposed in the literature [5]–[7], there is currently no widely accepted or standardized benchmarking protocol, dataset, or evaluation metric. Existing works [8]–[13] use different datasets, metrics, and environmental conditions, making direct comparisons between methods difficult and hindering systematic progress.

In this paper, we address this gap by proposing a principled evaluation protocol for topological mapping systems. Our metrics are derived from the core property of *topological consistency* (Section III-B): if two nodes in the graph are connected by an edge, they should correspond to physically proximate locations along a traversable route in the real world, and conversely, such proximate locations should be connected. We then show that *localization accuracy* (Section III-C) provides an efficient and interpretable surrogate metric for topological consistency.

We further highlight that perceptual aliasing is not only a system-level challenge but also a dataset property: the number and difficulty of ambiguous observations in the test environment strongly influence performance. Without controlling for this *ambiguity factor*, comparisons across datasets are unreliable (Section IV-B). We propose a quantitative method to measure dataset ambiguity, enabling more meaningful evaluations (Section III-D).

Following this protocol, we curate a diverse set of scenes and design targeted test cases. We evaluate several baseline systems, spanning common strategies for topological localization, on our curated datasets. Our results show that methods relying solely on topological structure struggle in the presence of perceptual aliasing, highlighting the need for additional cues beyond topology to achieve reliable mapping and localization (Section IV-B).

In summary, our contributions are:

- 1) We systematically define metrics for evaluating topological mapping systems and introduce the first framework, to our knowledge, for quantifying dataset ambiguity in this context.
- 2) We curate and release a diverse, publicly available dataset and evaluation toolkit to support consistent and reproducible benchmarking.
- 3) We implement and open-source baseline systems using deep-learned models, and evaluate them alongside classical methods; our experiments and analysis yield

*Equal contribution.

Emails: jiaming@comp.nus.edu.sg, {e0905370, e0774920}@u.nus.edu, harold@comp.nus.edu.sg.

insights into the limitations of current approaches and factors affecting robustness under perceptual aliasing.

II. BACKGROUND AND RELATED WORK

Broadly speaking, topological maps can be categorized into two types: those constructed *with* a metric SLAM map and those constructed *without* one (SLAM-free). The former relies on a metrically consistent representation (e.g., occupancy grid, TSDF/ESDF volume [14], [15], or pose graph) and builds a sparse topological abstraction on top of it to enable efficient planning and high-level reasoning. The latter constructs the topological map directly from sensory data, without requiring a globally consistent metric map. In this paper, our focus is **SLAM-free topological mapping systems without a metric consistent map from SLAM**.

A. SLAM-based Topological Maps

These methods start from a metric SLAM output and derive a sparse topological graph for planning. For example, Blöchliger et al. [14] convert a TSDF map into convex free-space clusters connected by adjacency edges, enabling A* search with near-RRT* path quality but much faster computation. Oleynikova et al. [15] extract a Generalized Voronoi Diagram from an ESDF and prune it into a thin 3D skeleton for robust graph generation in MAV planning. Konolige et al. [16] overlay a navigation graph on a SLAM pose graph while maintaining only local occupancy grids for execution. Friedman et al. [17] and Beeson et al. [18] extract Voronoi skeletons from occupancy grids and segment them into meaningful places such as rooms, doorways, and corridors. While these approaches produce reliable connectivity graphs and large planning speedups, they inherit the computational and maintenance costs of building a metrically consistent map, which can become a bottleneck in large or dynamic environments.

B. SLAM-free Topological Maps

These methods construct and maintain the topological map directly from observations, avoiding a globally consistent metric map. Milford et al. [5] and Glover et al. [8] use bio-inspired pose cell dynamics to filter place-recognition proposals from visual template matching. Cummins and Newman [6] and Maddern et al. [9] adopt a Bayesian filter where the observation model is learned via a Chow–Liu tree to capture word correlations, and the motion prior models topological distance. Learning-based methods such as Savinov et al. [7] train image-retrieval networks and use a sliding window to detect related frames, while Suomela et al. [11] and Xu et al. [10] integrate Bayesian filtering for robust retrieval. However, the variety of datasets and evaluation metrics used makes direct comparison between methods difficult.

III. EVALUATION ON TOPOLOGICAL SYSTEMS

A. Preliminaries: Topological Mapping

In contrast to SLAM systems, where each location is assigned a global metric pose and evaluated using distances




Fig. 1: Illustration of topological consistency. Nodes a , b , c , and d correspond to physical locations A , B , C , and D in the environment. Location A is route-close to B and C , but far from D . In the constructed topological map, the path between a and d violates *Edge Precision*, while the missing edge between a and b violates *Edge Recall*. The connection between a and c correctly preserves topological consistency.

in Euclidean or $SE(3)$ space, topological maps do not aim to preserve global metric consistency. Instead, they represent the environment as a connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with nodes $\mathcal{V}$ denoting distinct places or observations and edges $\mathcal{E}$ denoting navigable connections.

We adopt a generic framework that reflects the core procedure of many topological mapping systems. At each timestep t , given a new observation z_t , the system follows an *update policy* P that governs how nodes and edges are added:

- The agent first attempts to *localize* in the existing graph $\mathcal{G}$. If a node $v_j \in \mathcal{V}$ is successfully retrieved (above a similarity threshold τ), an edge (v_i, v_j) is created between the current node v_i and the matched node v_j .
- If no suitable match is found, a new node v_{i+1} is created when z_t differs sufficiently from prior observations and/or the agent has moved beyond a spatial threshold from the last node. This new node is then connected to v_i with an edge (v_i, v_{i+1}) .

In this work, we consider P to be an *edge-length regular policy*, where each new edge has a bounded geodesic route length proportional to the median edge length μ_e . This mild assumption reflects common practice, as edges are typically added only between consecutive nodes once the robot has moved beyond a threshold. Focusing on edge-length regular policies remains general, since shortcut edges or other connections can always be added later through post-processing.

Based on this policy, we define an *edge creation opportunity* as a pair of nodes (u, v) that P could have connected at the evaluation scale d . The set of such pairs, denoted $\Omega_P(d)$, specifies the policy-eligible connections that should be considered when evaluating recall.

B. Topological Consistency

For a topological map to be meaningful and useful for representing the environment and enabling navigation planning,



Fig. 2: Examples of map, test, and ambiguous cases across datasets. Top row: A+P case from OpenLORIS [19], where the test differs from the map due to a time-of-day change, and the ambiguous (false-positive) image appears visually close to the target. Bottom left: P.O. case from RobotCar [20], showing a map–test pair of the same location with the building façade temporarily covered during construction. Bottom right: A.O. case from RELLIS-3D [21], where the ambiguous resembles the test except for differences in vegetation type. See Sec. III-D for case definitions.

it should satisfy two properties:

- 1) **Edge Precision:** If two nodes in the topological graph are within n hops of each other, then their corresponding locations in the environment should be within a geodesic route distance threshold d .
- 2) **Edge Recall (policy-conditioned):** If two locations in the environment are within a geodesic route distance threshold d , and their corresponding nodes form an edge creation opportunity under the update policy P (i.e., $(u, v) \in \Omega_P(d)$), then those nodes should be connected within n hops in the graph.

This concept is illustrated in Fig. 1. We use *route distance* (the shortest traversable geodesic in the environment) rather than Euclidean distance, since topological maps are intended to capture navigable paths where direct straight-line connections may not exist. In practice, ground-truth route distances are rarely available, so we approximate them using traversal distances obtained from the robot’s odometry. During dataset curation, we manually verified that no repeated zig-zag motions were present, ensuring that traversal distances do not overestimate the true geodesic distance.

Different mapping systems produce edges of varying lengths, so fixed thresholds (n, d) may bias evaluation. We instead define a characteristic edge length using the median, which robustly captures a “typical” edge:

$$\mu_e := \text{median}\{\text{dist}_R(u, v) \mid (u, v) \in \mathcal{E}\}, \quad (1)$$

where $\text{dist}_R(u, v)$ is the geodesic route distance between the locations of nodes u and v .

Given a physical route distance threshold d , we adapt the hop threshold as $n = \max\left(1, \left\lceil \frac{\epsilon d}{\mu_e} \right\rceil\right)$, where $\epsilon > 0$ is a

tolerance factor and the floor ensures $n \geq 1$. Evaluation is thus based on the physical route distance d , as navigation requires maps to support localization and connectivity within a bounded physical distance. Since d is usually larger than the typical edge length μ_e , n often exceeds 1.

With this definition, the two properties can be formalized as:

$$\forall u, v \in \mathcal{V}, \quad \text{dist}_G(u, v) \leq n \Rightarrow \text{dist}_R(u, v) \leq d, \quad (2)$$

$$\forall (u, v) \in \Omega_P(d), \quad \text{dist}_R(u, v) \leq d \Rightarrow \text{dist}_G(u, v) \leq n, \quad (3)$$

where $\Omega_P(d)$ is the set of policy-eligible edge opportunities (Section III-A), $\text{dist}_G(\cdot, \cdot)$ is shortest-path hop distance in the graph, and $\text{dist}_R(\cdot, \cdot)$ is geodesic route distance in the environment.

Equation (2) defines *Edge Precision*: nodes that are close in the graph (within n hops) should also be geodesically close in the environment (within d). Equation (3) defines *Policy-conditioned Edge Recall*: geodesically close locations should be connected within n hops, but only if the update policy P makes such a connection feasible.

This policy-aware definition avoids penalizing systems for missing edges they would never create (e.g., between two nearby but non-consecutive nodes), while still enforcing that feasible edge opportunities preserve navigability.

C. Localization Accuracy

While edge precision and recall (Section III-B) measure structural consistency, a topological map must also support reliable *localization* for safe navigation. That is, given the current sensor observation, the system should correctly identify the robot’s current node in the existing graph.

Formally, let $\hat{v}_t$ be the predicted node at time t and v_t^* the ground-truth node. Localization is considered successful if the geodesic route distance between them is within a physical threshold d :

$$\forall t, \quad \text{dist}_R(\hat{v}_t, v_t^*) \leq d. \quad (4)$$

If the test observation lies in a novel region with no corresponding ground-truth node, then success is defined by correct rejection: accuracy is 1 if the system abstains from matching, and 0 otherwise.

We argue that if the system makes *correct localization decisions* at each step, then Edge Precision and policy-conditioned Edge Recall are preserved as the map grows, making localization accuracy a faithful surrogate for topological consistency. We provide a brief sketch proof by induction below.

Since the update policy P is *edge-length regular* (Sec. III-A), each new edge has geodesic route length at most $\kappa \mu_e$ for some $\kappa \geq 1$. Choose the hop budget $n = \max(1, \lfloor \epsilon d / \mu_e \rfloor)$ with $\epsilon \leq 1/\kappa$, so that any n -hop path has geodesic length $\leq n \kappa \mu_e \leq d$.

Suppose at time t the map $\mathcal{G}_t$ already satisfies Edge Precision and policy-conditioned Edge Recall at scale (n, d) . After processing z_{t+1} there are two cases:

(i) *Correct accept (revisit)*. The system links the current node v_i to the correctly retrieved node v_j with a new edge. For **Edge Precision**, any pair within n hops in $\mathcal{G}_{t+1}$ either (a) was already within n hops in $\mathcal{G}_t$ and thus satisfies $\text{dist}_R \leq d$ by induction, or (b) uses the new edge once; in this case the path length is at most $n \kappa \mu_e \leq d$, so $\text{dist}_R \leq d$. For **policy-conditioned Edge Recall**, all previously eligible pairs remain connected within n hops, and any new eligible pair (in particular (v_i, v_j)) is directly connected, i.e., $\text{dist}_G(v_i, v_j) = 1 \leq n$.

(ii) *Correct reject (novel)*. The system creates a new node u and a consecutive edge (v_i, u) . For **Edge Precision**, any n -hop path either existed before (hence $\text{dist}_R \leq d$) or uses the new edge once, which again yields $\text{dist}_R \leq d$. For **policy-conditioned Edge Recall**, the only new eligible pair is (v_i, u) , which is connected in one hop, while all previously eligible pairs retain their n -hop connectivity.

Thus, under correct localization decisions, both properties are invariants: if they hold at time t , they also hold at $t + 1$. Since the base map (empty or a single node) trivially satisfies them, they hold at all times. This establishes localization accuracy as a concise, interpretable, and faithful surrogate for topological consistency. In summary, localization accuracy is directly linked to both navigation success and successful map construction.

D. Ambiguity Quantification

Perceptual aliasing, where distinct locations produce highly similar sensor observations, is a core challenge in topological mapping [4]. It can cause the map to mistakenly add edges between unrelated places, leading to structural corruption. Such ambiguity arises across diverse environments, from uniform office corridors and near-identical urban




Fig. 3: Examples of evaluation scenarios with quantified ambiguity. (a) **Ambiguous + Positive (A+P)**: The test sequence revisits a mapped region, but at least one distractor subsequence appears nearly as similar as the true match, making localization uncertain. (b) **Positive Only (P.O.)**: The test sequence revisits a mapped region and the true subsequence clearly dominates all distractors, resulting in unambiguous localization. (c) **Ambiguous Only (A.O.)**: The test sequence lies in a novel, unmapped region, yet one or more mapped subsequences spuriously appear visually similar, producing false matches. Omitting the intermediate trajectory (shown in magenta) yields the kidnapped-robot variant; otherwise, the scenario corresponds to a classical loop closure.

intersections to repetitive patterns in natural trails (Fig. 2). Yet most prior work evaluates mapping performance without explicitly measuring or controlling for this phenomenon, resulting in inconsistent and potentially misleading comparisons across systems and environments. This motivates our explicit quantification of ambiguity as a central element of evaluation.

To incorporate ambiguity-aware evaluation, we adopt the following protocol. First, the system is provided with a *mapping sequence* $\mathcal{S}_m = \{z_m^j\}_{j=1}^N$ followed by a short *test sequence* $\mathcal{S}_t = \{z_t^i\}_{i=1}^L$. The test sequence may revisit previously mapped regions or explore an entirely new place. Second, we evaluate localization accuracy (Section III-C) using the final frame of $\mathcal{S}_t$. Here, z_m^j and z_t^i denote individual sensor observations (e.g., images) in the mapping and test phases, respectively. Using ground truth (e.g., GPS or motion capture), each test frame z_t^i is aligned to its corresponding mapping frame via an index function $\pi(i) \in \{1, \dots, N\}$, such that $z_m^{\pi(i)}$ is the true spatial match. If the test sequence corresponds to a novel place, we set $\pi(i) = \emptyset$. This setup accommodates both loop closure and kidnapped-robot scenarios (Fig. 3).

We define ambiguity at the level of short *sequences* of frames. For a candidate subsequence of length L , $\mathcal{S}_m^{[u]} = \{z_m^{j_u}, \dots, z_m^{j_u+L-1}\}$, we measure similarity as

$$\text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[u]}) = \frac{1}{L} \sum_{i=1}^L \text{sim}(z_t^i, z_m^{j_u+i-1}),$$

where $\text{sim}(\cdot, \cdot)$ denotes the similarity between two frames, provided by a visual place recognition (VPR) model [22]. A sequence pair is deemed visually similar if $\text{Sim}(\cdot, \cdot) \geq \tau$, with τ a predefined threshold.

To validate the VPR model choice, we curated a balanced

dataset of 200 image pairs spanning a wide range of appearance similarities. We evaluated 11 VPR methods reviewed in [23], and found a mean inter-method correlation of **0.83**. This high agreement suggests that ambiguity is an intrinsic property, largely independent of the specific method used. For dataset construction, we adopted BOQ [12] due to its efficiency and robustness.

Under this definition, each test sequence falls into one of three categories (Fig. 3):

Ambiguous + Positive (A+P):

$$\pi(\mathcal{S}_t) \neq \emptyset, \quad \frac{\max_{u \neq \pi(\mathcal{S}_t)} \text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[u]})}{\text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[\pi(\mathcal{S}_t)]})} \geq \alpha.$$

The sequence revisits a known region, but at least one distractor subsequence in the map appears almost as similar as the true match. Such cases often arise from scene changes in the test sequence (e.g., variations in lighting or time of day) that narrow the confidence gap between true and distractor matches.

Positive Only (P.O.):

$$\pi(\mathcal{S}_t) \neq \emptyset, \quad \frac{\max_{u \neq \pi(\mathcal{S}_t)} \text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[u]})}{\text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[\pi(\mathcal{S}_t)]})} < \alpha.$$

The true match clearly dominates all distractors, making localization unambiguous.

Ambiguous Only (A.O.):

$$\pi(\mathcal{S}_t) = \emptyset, \quad \max_u \text{Sim}(\mathcal{S}_t, \mathcal{S}_m^{[u]}) \geq \tau.$$

The test sequence lies in a novel region, but one or more mapped subsequences appear spuriously similar.

Discussion. Our protocol improves on loop-closure precision/recall in two ways. First, it is *ambiguity-aware*: by separating Positive-Only (P.O.), Ambiguous+Positive (A+P), and Ambiguous-Only (A.O.), it reveals failure modes that pairwise metrics conflate, such as disambiguation errors versus unsafe matches in novel regions. Second, it evaluates the robot’s *query-time decision*, that is, whether to localize or abstain, rather than candidate-pair rates. This better reflects safe navigation needs, where a single wrong match can corrupt the map. In contrast, precision/recall may look strong even when top-1 mapping decisions are unreliable, since they summarize candidate pairs rather than end-to-end reliability.

Still, comparing systems by per-case accuracies is difficult. A method may excel on P.O. yet fail on A.O. To address this, we introduce a summary metric that balances performance across all case types.

E. Balanced Localization Accuracy

The *Balanced Localization Accuracy* (BLA) metric provides a unified measure of performance across the three case types defined in Section III-D. It is computed as the geometric mean of localization accuracies:

$$\text{BLA} = (\mathcal{L}_{A+P} \mathcal{L}_{P.O.} \mathcal{L}_{A.O.})^{1/3},$$

where $\mathcal{L}$ denotes accuracy on the respective case type. The geometric mean penalizes imbalance: a system that excels on

TABLE I: Summary of datasets with environment counts, types (In./Out.), and difficulty-controlled test cases.

	[19]	[20]	[24]	[25]	[21]	[26]	Total
Env (No.)	5	1	1	16	1	1	25
Env (Type)	In.	Out.	In.	In.	Out.	Out.	–
A+P	30	0	16	0	0	5	51
P.O.	74	116	90	74	0	30	384
A.O.	52	0	14	3	117	8	194

one case type but fails on another will see its overall score reduced, analogous to how the F -score balances precision and recall.

To avoid degenerate cases when any accuracy is exactly zero, we apply *Jeffreys prior smoothing*, a Bayesian adjustment with $\text{Beta}(\frac{1}{2}, \frac{1}{2})$ prior. Each empirical rate $r = k/n$ is replaced with $\tilde{r} = \frac{k+0.5}{n+1}$, which prevents collapse and has negligible effect when n is large.

Each accuracy $\mathcal{L}$ depends on the decision threshold τ , which controls whether a retrieved candidate is accepted or rejected. Consequently, the BLA also varies with τ . Typically, increasing τ improves rejection of ambiguous observations (raising $\mathcal{L}_{A.O.}$) but reduces acceptance of correct matches (lowering $\mathcal{L}_{P.O.}$ and $\mathcal{L}_{A+P}$).

In our experiment we report the accuracy $\mathcal{L}$ for each test case and the corresponding BLA, with the threshold τ chosen in two ways. First, $\mathcal{L}_{A.O.}@p$ denotes the operating point where $\mathcal{L}_{A.O.}(\hat{\tau}) \approx p$, within a tolerance of $\pm 3\%$ and preferring the conservative choice $\mathcal{L}_{A.O.}(\hat{\tau}) \geq p$, reflecting performance at a desired reliability level (e.g., $p = 0.90$ or 0.99). Second, $\text{BLA}_{\max}$ is obtained by selecting the threshold $\hat{\tau}$ that maximizes BLA across all values, representing the best performance achievable under ideal tuning. In both cases, the threshold $\hat{\tau}$ is determined on a smaller validation set and then applied for evaluation on the full dataset, so the resulting $\mathcal{L}$ and BLA values may differ between the two stages.

F. Dataset Collection

To evaluate robustness across diverse environments, sensing setups, and robot embodiments, we compose a benchmark from datasets that provide time-synchronized observations and ground-truth trajectories (required for the correspondence function $\pi(\cdot)$ in Sec. III-D). We include both indoor and outdoor scenes, and select datasets that naturally contain visual ambiguities through repeated traversals under varying conditions.

We consider the following datasets:

- 1) **OpenLORIS** [19]: Multiple traversals of everyday indoor environments with variations in activity, lighting, and object placement causing perceptual aliasing.
- 2) **Oxford RobotCar** [20]: over 100 traversals of a fixed route through Oxford, across a year under diverse weather, traffic, and construction.
- 3) **Rawseeds** [24]: Long outdoor and indoor trajectories with structurally similar areas.

- 4) **Habitat-Sim** [25]: Simulated indoor environments with curated perceptual ambiguities.
- 5) **RELLIS-3D** [21]: Off-road trajectories with dense vegetation, rough terrain, and aliasing.
- 6) **ROVER** [26]: Seasonal outdoor traversals capturing changes in vegetation, terrain, and lighting.

Summary statistics for the benchmark are provided in Table I.

IV. BENCHMARKING BASELINES

A. Baseline Systems

A number of recent works have introduced SLAM-free topological mapping systems (Section II-B), which differ primarily in how they address the *perceptual aliasing problem*. To study these strategies in depth, we implement three representative approaches that capture their core algorithmic principles.

a) Sequence Matching (SM).: Instead of matching a single observation, the system matches short sequences of consecutive observations to improve robustness against perceptual aliasing and viewpoint variation. This temporal context helps reduce false matches caused by transient visual similarities. Representative examples include [7], [13], [27]. A candidate shortcut edge (v_i, v_j) is accepted if the aggregated similarity score over a temporal window satisfies

$$f(\text{sim}(z_{v_i-h}, z_{v_j-h}), \dots, \text{sim}(z_{v_i+h}, z_{v_j+h})) \geq \tau,$$

where h is the half-window size in frames and $f\{\cdot\}$ is an aggregation function applied to the similarity scores in the window.

In our experiments, we evaluate two variants of sequence matching:

- **SM-Med**: $f\{\cdot\}$ is the median of the similarity scores, which provides robustness to outliers.
- **SM-All**: $f\{\cdot\}$ requires that *all* similarity scores in the window exceed the threshold τ , enforcing stricter consistency across the entire sequence.

b) Probabilistic belief update (PBU).: To maintain temporal consistency, the system maintains a posterior belief $p_t(v)$ over the current node v and updates it at each timestep. Between observations, the belief is propagated using a topological motion model that constrains allowable transitions:

$$p(v_i | v_j) \propto \begin{cases} \alpha & \text{if } \text{dist}_G(v_i, v_j) \leq w_u, \\ \beta & \text{otherwise,} \end{cases}$$

where $\text{dist}_G(\cdot, \cdot)$ is the hop distance in the graph, w_u is the maximum allowed hop movement per step, and α, β are constants controlling transition likelihoods.

After motion propagation, the belief is updated using the similarity scores as observation likelihoods:

$$p_t(v) \propto g(\text{sim}(z_t, z_v)) \cdot \sum_{u \in \mathcal{V}} p(v | u) p_{t-1}(u),$$

where $g(\cdot)$ is a likelihood function mapping similarity scores to measurement probabilities.

This probabilistic filtering suppresses spurious matches that violate motion constraints, improving robustness to perceptual aliasing. Representative examples include [10], [11].

c) Greedy Matching (GM).: We additionally implement a simple greedy matching baseline, which localizes by selecting the highest-scoring match above a fixed similarity threshold; if no match exceeds the threshold, the localization result remains unchanged.

In all our baseline implementations, we use the same VPR model MegaLoc [28] for fairness, in contrast to prior works that often train environment-specific retrieval models, which may not generalize well to unseen environments. This setup also enables a fair comparison of different topological mapping strategies under identical retrieval performance.

Finally, we also include specialized strategies that have demonstrated strong performance in large-scale datasets:

- **FAB-MAP** [6]: Each observation Z_k is represented as a bag-of-visual-words from SURF features. A Bayesian filter maintains a posterior over location nodes, with likelihoods estimated via a Chow–Liu tree. A motion prior favors adjacent locations, while a Monte Carlo approximated “new-place” hypothesis accounts for novel scenes. Data association selects the most probable location, and top matches are verified with RANSAC.
- **RatSLAM** [5]: A biologically inspired system where each node (experience) $v_i = \{P^i, V^i, \mathbf{p}_i\}$ stores the pose code P^i , local view code V^i , and map location $\mathbf{p}_i$. Each edge $e_{ij} \in E$ encodes the odometry-based displacement $\Delta \mathbf{p}_{ij}$ between connected experiences. Upon receiving a new image observation, RatSLAM performs visual template matching to activate corresponding local view cells, applies pose cell filtering with attractor dynamics to integrate visual and self-motion cues, and finally selects the best-matching vertex in G for localization.

B. Benchmarking Results

We evaluated the baseline systems on our ambiguity-quantified dataset, reporting per-case localization accuracies (A+P, P.O., A.O.) as well as the aggregated BLA metric. Following Section III-E, we highlight two complementary operating points: $\mathcal{L}_{A.O.}@90$, which anchors evaluation at a safety-critical A.O. reliability level, and BLA_{max} , which captures the best achievable balance across thresholds. Note that $\mathcal{L}_{A.O.}@90$ is determined by selecting the threshold τ on a smaller validation set; consequently, the measured $\mathcal{L}_{A.O.}$ on the full dataset may differ. Numerical results are summarized in Table II, while Fig. 4 plots localization accuracy as a function of threshold for each case type, illustrating how different baselines trade recall against safety.

a) Overall trends.: All evaluated methods fall short of the robustness required for real-world deployment. When thresholds are set conservatively so that A.O. cases achieve high rejection accuracy (e.g., 90% in Table II), performance on A+P and P.O. cases drops sharply. At even stricter settings ($\mathcal{L}_{A.O.}@99$), every baseline fails completely on A+P and

TABLE II: Localization accuracy of each baseline method, reported at two operating points: the decision threshold τ corresponding to $\mathcal{L}_{A.O.}@90$ and the threshold yielding $BLA_{\max}$. For both settings, per-case accuracies (A+P, P.O., A.O.) and the corresponding BLA are shown (see Section III-E).

Method	$\mathcal{L}_{A.O.}@90$				$BLA_{\max}$			
	$\mathcal{L}_{A+P} (\uparrow)$	$\mathcal{L}_{P.O.} (\uparrow)$	$\mathcal{L}_{A.O.} (\uparrow)$	$BLA (\uparrow)$	$\mathcal{L}_{A+P} (\uparrow)$	$\mathcal{L}_{P.O.} (\uparrow)$	$\mathcal{L}_{A.O.} (\uparrow)$	$BLA (\uparrow)$
FAB-MAP [6]	0.0	0.036	0.964	0.058	0.0	0.078	0.732	0.067
RatSLAM [5]	0.0	0.081	0.438	0.057	0.0	0.081	0.438	0.057
GM	0.020	0.402	0.577	0.178	0.078	0.479	0.423	0.256
SM-Med	0.039	0.311	0.778	0.220	0.078	0.489	0.412	0.255
SM-All	0.020	0.229	0.784	0.164	0.039	0.388	0.469	0.200
PBU	0.020	0.402	0.577	0.178	0.235	0.646	0.165	0.295




Fig. 4: Localization accuracy for A+P, P.O., and A.O. test cases, and BLA, plotted against the decision threshold for all baseline methods. For SM-Med, GM, and PBU the threshold is directly tied to detection scores, while for FabMap and RatSLAM it is defined by more complex scoring functions. Results show that all baseline methods perform poorly in the A+P test case, highlighting the need for benchmarks that explicitly quantify ambiguity and for more sophisticated approaches to resolve perceptual aliasing. In contrast, methods generally perform well in the P.O. test case at low threshold values, since in the absence of perceptual aliasing this setting mainly tests the system’s ability to detect candidates (analogous to recall). For the A.O. test case, methods achieve high accuracy in the high-threshold regime, as the correct outcome is consistently to reject. Finally, the BLA case provides a balanced view across the three scenarios. Overall, these results suggest that current state-of-the-art methods cannot adequately handle scenarios where perceptual aliasing is present.

P.O., yielding zero localization accuracy. This reveals a fundamental limitation: current methods address perceptual aliasing only by rejecting more aggressively, but lack robust mechanisms to *disambiguate* visually similar places. As a result, they remain safe only by refusing nearly all revisits. In practice, both failure modes are problematic; for example, a return to the same room may be treated as a new place, fragmenting the map and forcing inefficient routes. These results underscore the lack of robust methods for aliasing and motivate the need for explicitly ambiguity-quantified datasets to expose such trade-offs.

b) Classical vs. deep-learned VPR.: Compared with classical approaches such as FAB-MAP [6] and RatSLAM [5], baselines using deep-learned VPR descriptors achieve higher $\mathcal{L}_{A+P}$ and $\mathcal{L}_{P.O.}$ while maintaining a similar level of $\mathcal{L}_{A.O.}$, as reflected in the $\mathcal{L}_{A.O.}@90$ results in Table II. This highlights that deep learning–based retrieval offers a stronger foundation for place recognition than hand-engineered pipelines based on feature detection, matching, and geometric verification.

c) Sequence matching (SM).: Matching sequences rather than single frames was expected to improve robustness, but in practice the benefit is limited and highly sensitive

to the choice of window size h and aggregation function $f(\cdot)$. If h is too short, the method collapses to greedy single-frame matching and fails on ambiguous revisits. If h is too long, test and mapping sequences may diverge due to natural viewpoint changes—for example, the robot may look rightward during testing but straight ahead during mapping—leading the aggregated similarity to fall below threshold even for true revisits. Aggregation with the median is somewhat more robust to outlier mismatches, but this only helps in cases where distractors dominate the top similarity scores; when the true positives are strong, median aggregation can suppress them. As a result, sequence matching baselines do not consistently outperform greedy matching, and their performance curves (see Fig. 4) show large variance depending on parameter settings.

d) Probabilistic belief update (PBU).: Incorporating a topological motion prior via probabilistic filtering yields only modest improvements. The prior is helpful when the system is already well localized and enters an ambiguous segment, as belief can then be reinforced sequentially and spurious matches suppressed. In practice, however, the model assumes consistent hop distances and motion patterns, which rarely hold. For instance, a mapping sequence may contain

nodes every 1 m, while an ambiguous corridor encountered during testing may produce visually similar nodes spaced 2 m apart. Without precise metric motion information, the prior incorrectly treats these sequences as consistent. This mismatch explains why PBU struggles in practice: the coarse topological prior provides only limited support for disambiguating visually similar places.

V. CONCLUSION

In this paper, we proposed an ambiguity-aware evaluation framework for topological mapping systems. We introduced *localization accuracy* as a surrogate for topological consistency and defined *balanced localization accuracy* (BLA) to aggregate performance across distinct test cases. To support fair and reproducible evaluation, we curated a new dataset with explicit ambiguity quantification.

Using this dataset, we evaluated both classical approaches and modern systems with deep-learned VPR models. The results yield three main insights: (i) ambiguity-aware evaluation reveals failure modes hidden by conventional loop-closure metrics, as tuning for safe rejection (high $\mathcal{L}_{A.O.}$) causes localization accuracy on revisits (A+P and P.O.) to collapse, showing that current methods cannot truly resolve aliasing; (ii) classical methods remain reliable only for rejecting novel places but fail even on unambiguous revisits, whereas deep-learned retrieval achieves higher revisit accuracy at similar safety levels, reflected in stronger $\mathcal{L}_{A.O.}@90$; and (iii) sequence matching and probabilistic filtering add complexity without consistent gains.

Overall, these findings underscore the need for more robust disambiguation strategies. Our evaluation metrics and ambiguity-quantified benchmark provide a consistent and reproducible basis for assessing topological mapping systems, fostering further development of reliable and scalable approaches. To ensure reproducibility, we will release all code in a public GitHub repository, which also documents implementation details that cannot be fully included in the paper due to space constraints.

REFERENCES

- [1] J. Wang, D. Liu, J. Chen, J. Da, N. Qian, T. M. Man, and H. Soh, “Genie: A generalizable navigation system for in-the-wild environments,” *arXiv preprint arXiv:2506.17960*, 2025.
- [2] A. Sridhar, D. Shah, C. Glossop, and S. Levine, “Nomad: Goal masked diffusion policies for navigation and exploration,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 63–70.
- [3] R. A. Brooks, “Intelligence without representation,” *Artificial intelligence*, vol. 47, no. 1-3, pp. 139–159, 1991.
- [4] J. Boal, A. Sánchez-Mirallès, and A. Arranz, “Topological simultaneous localization and mapping: a survey,” *Robotica*, vol. 32, no. 5, pp. 803–821, 2014.
- [5] M. J. Milford and G. F. Wyeth, “Mapping a suburb with a single camera using a biologically inspired slam system,” *IEEE Transactions on Robotics*, vol. 24, no. 5, pp. 1038–1053, 2008.
- [6] M. Cummins and P. Newman, “Highly scalable appearance-only slam-fab-map 2.0,” *Robotics: Science and Systems V*, 2009.
- [7] N. Savinov, A. Dosovitskiy, and V. Koltun, “Semi-parametric topological memory for navigation,” *arXiv preprint arXiv:1803.00653*, 2018.
- [8] A. J. Glover, W. P. Maddern, M. J. Milford, and G. F. Wyeth, “Fab-map+ ratslam: Appearance-based slam for multiple times of day,” in *2010 IEEE international conference on robotics and automation*. IEEE, 2010, pp. 3507–3512.
- [9] W. Maddern, M. Milford, and G. Wyeth, “Cat-slam: probabilistic localisation and mapping using a continuous appearance-based trajectory,” *The International Journal of Robotics Research*, vol. 31, no. 4, pp. 429–451, 2012.
- [10] M. Xu, N. Snderhauf, and M. Milford, “Probabilistic visual place recognition for hierarchical localization,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 311–318, 2020.
- [11] L. Suomela, J. Kalliola, H. Edelman, and J.-K. Kämäräinen, “Placenv: Topological navigation through place recognition,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 5205–5213.
- [12] A. Ali-Bey, B. Chaib-draa, and P. Giguere, “Boq: A place is worth a bag of learnable queries,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17 794–17 803.
- [13] X. Meng, N. Ratliff, Y. Xiang, and D. Fox, “Scaling local control to large-scale topological navigation,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 672–678.
- [14] F. Blöchliger, M. Fehr, M. Dymczyk, T. Schneider, and R. Siegwart, “Topomap: Topological mapping and navigation based on visual slam maps,” in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 3818–3825.
- [15] H. Oleynikova, Z. Taylor, R. Siegwart, and J. Nieto, “Sparse 3d topological graphs for micro-aerial vehicle planning,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [16] K. Konolige, E. Marder-Eppstein, and B. Marthi, “Navigation in hybrid metric-topological maps,” in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 3041–3047.
- [17] S. Friedman, H. Pasula, and D. Fox, “Voronoi random fields: Extracting the topological structure of indoor environments via place labeling,” in *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2007.
- [18] P. Beeson, N. K. Jong, and B. Kuipers, “Towards autonomous topological place detection using the extended voronoi graph,” in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2005, pp. 4373–4379.
- [19] X. Shi, D. Li, P. Zhao, Q. Tian, Y. Tian, Q. Long, C. Zhu, J. Song, F. Qiao, L. Song, Y. Guo, Z. Wang, Y. Zhang, B. Qin, W. Yang, F. Wang, R. H. M. Chan, and Q. She, “Are we ready for service robots? the OpenLORIS-Scene datasets for lifelong SLAM,” in *2020 International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3139–3145.
- [20] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 Year, 1000km: The Oxford RobotCar Dataset,” *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017. [Online]. Available: <http://dx.doi.org/10.1177/0278364916679498>
- [21] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, “Rellis-3d dataset: Data, benchmarks and analysis,” *arXiv preprint arXiv:2011.12954*, 2020.
- [22] S. Lowry, N. Sönderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, “Visual place recognition: A survey,” *IEEE transactions on robotics*, vol. 32, no. 1, pp. 1–19, 2015.
- [23] G. Berton, G. Trivigno, B. Caputo, and C. Masone, “Eigenplaces: Training viewpoint robust models for visual place recognition,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 11 080–11 090.
- [24] G. Fontana, M. Matteucci, and D. G. Sorrenti, “Rawseeds: Building a benchmarking toolkit for autonomous robotics,” in *Methods and experimental techniques in computer engineering*. Springer, 2013, pp. 55–68.
- [25] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, “Habitat: A Platform for Embodied AI Research,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [26] F. Schmidt, J. Daubermann, M. Mitschke, C. Blessing, S. Meyer, M.ENZweiler, and A. Valada, “Rover: A multiseason dataset for visual slam,” *IEEE Transactions on Robotics*, vol. 41, pp. 4005–4022, 2025.
- [27] M. J. Milford and G. F. Wyeth, “Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights,” in *2012 IEEE international conference on robotics and automation*. IEEE, 2012, pp. 1643–1649.
- [28] G. Berton and C. Masone, “Megaloc: One retrieval to place them all,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 2861–2867.