

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ONLINE NAVIGATION REFINEMENT: ACHIEVING LANE- LEVEL GUIDANCE BY ASSOCIATING STANDARD- DEFINITION AND ONLINE PERCEPTION MAPS

Anonymous authors

Paper under double-blind review

ABSTRACT

Lane-level navigation is critical for geographic information systems and navigation-based tasks, offering finer-grained guidance than road-level navigation by standard definition (SD) maps. However, it currently relies on expansive global HD maps that cannot adapt to dynamic road conditions. Recently, online perception (OP) maps have become research hotspots, providing real-time geometry as an alternative, but lack the global topology needed for navigation. To address these issues, Online Navigation Refinement (ONR), a new mission is introduced that refines SD-map-based road-level routes into accurate lane-level navigation by associating SD maps with OP maps. The map-to-map association to handle many-to-one lane-to-road mappings under two key challenges: (1) no public dataset provides lane-to-road correspondences; (2) severe misalignment from spatial fluctuations, semantic disparities, and OP map noise invalidates traditional map matching. For these challenges, We contribute: (1) Online map association dataset (OMA), the first ONR benchmark with 30K scenarios and 2.6M annotated lane vectors; (2) MAT, a transformer with path-aware attention to aligns topology despite spatial fluctuations and semantic disparities and spatial attention for integrates noisy OP features via global context; and (3) NR P-R, a metric evaluating geometric and semantic alignment. Experiments show that MAT outperforms existing methods at 34 ms latency, enabling low-cost and up-to-date lane-level navigation.

1 INTRODUCTION

Lane-level navigation has emerged as a critical capability in geographic information systems (GIS) Hansson et al. (2020); Guo et al. (2025) and navigation-based tasks Peng et al. (2025),

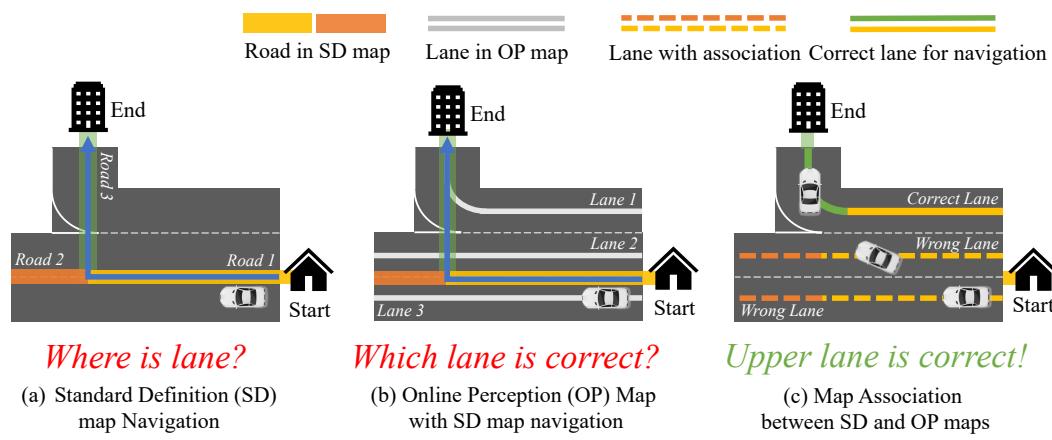


Figure 1: Motivation for online navigation refinement and map association. The roads and lanes with the same color indicate that they are interconnected: (a) Standard Definition (SD) maps offer only road-level navigation without lane details. (b) Online Perception (OP) maps offer lane-specific details, yet they are not connected to SD maps and cannot identify the correct lane. (c) Associating SD and OP maps enables lane selection, achieving online navigation refinement and lane-level navigation.

offering finer-grained guidance than traditional road-level navigation based on standard definition (SD) maps (Fig. 1 (a)) Zhang et al. (2024d). However, today’s lane-level navigation systems mostly rely on pre-built HD maps, which are expensive to create and maintain. They often miss real-world changes Elghazaly et al. (2023)—like construction, rule updates, or accidents—leading to outdated guidance and safety risks. Recently, online perception (OP) maps have become research hotspots within autonomous driving Li et al. (2022); Liao et al. (2023b). Powered by vehicle sensors, the OP maps offer current and localized lane depictions. However, as shown in Fig. 1 (b), they inherently lack global route topology and cannot support lane-level navigation Wong et al. (2020).

To achieve up-to-date and low-cost lane-level navigation, we introduce a new mission called **Online Navigation Refinement** (ONR). The core objective of the mission is to transform the road-level navigation derived from SD maps into precise lane-level navigation aligned with the OP maps. For a fast and accurate process, ONR needs a paradigm of map-to-map matching, called **map association**, as shown in Fig 1 (c). Compared with map matching (MM) that bind GPS trajectories to static SD/HD maps Chao et al. (2020) as a path-to-map matching based on the HMM Newson (2009), seq-to-seq model Feng et al. (2020); Ren et al. (2021), graph model Liu et al. (2023b) and transformers Tang et al. (2025), map association uses map-to-map matching because SD maps and OP maps are heterogeneous and do not share a one-to-one correspondence. Attempting to directly align SD paths with OP maps, or the other way around, overlooks the semantic disparities between the two.

Specifically, map association faces two key challenges. First, no public datasets provide structured, lane-to-road-level correspondences between SD maps and online perception (OP) maps. Although there exist auto-driving datasets (e.g., nuScenes Caesar et al. (2020a), OpenLaneV2 Wang et al. (2023)) offer local lane geometries, and OpenStreetMap provides road-level topologies, none establish explicit, learnable mappings between them. Secondly, SD and OP maps display inherent heterogeneity due to diverse granularity, resulting in substantial spatial and semantic differences. A robust association mechanism is required to manage spatial fluctuations (such as GPS drift and changes in scale) and semantic disparities (such as the differing number of lanes on a road). It should also be capable of processing noise issues in OP maps in real time, which include lane discontinuities, omissions, or errors in OP mapping. Figure 6 in Appendix D illustrates the complexity and challenge of map association.

To address these challenges, we make three core contributions: **(1) Dataset:** We introduce **Online Map Association Dataset (OMA)**, the first open source benchmark for the online navigation refinement mission in the paradigm of map association, derived from nuScenes Caesar et al. (2020a) and OpenStreetMap osm, which contains more than 30K scenarios, 480K road paths and 2.6M lane vectors with manually annotated associations. **(2) Baseline:** We present **MAT**, a lightweight transformer-based model for the real-time association of maps. To handle spatial and semantic differences, MAT incorporates two key modules: path-aware attention and spatial attention. Path-aware attention reorders and group vector tokens by path index, implicitly encoding topological structure to align SD and OP maps despite spatial offsets. Spatial attention sorts and groups road/centerline tokens using a spatial curve, enabling global context modeling through cross-topological feature integration. **(3) Metric:** We propose **Navigation Refinement P-R (NR P-R)**, a metric that measures alignment of the path and the correspondence through precision, recall and F1 specifically designed for map association. Traditional MM metrics only measure the accuracy of the matching, but the NR P-R evaluates both the geometric similarity and the accuracy of the matching. Moreover, NR P-R only utilizes annotation in ground-truth perception maps, rendering the benchmark suitable for the evaluation of any map generation method.

Extensive experiments indicate that MAT achieves state-of-the-art performance with 34ms latency on OMA, showing significant improvements compared to traditional map matching methods and deep learning map matching method, enabling low-cost and up-to-date lane-level navigation.

2 RELATED WORK

Map Matching. Map matching investigates how to link GPS tracks to SD or HD maps. Hansson et al. (2020) utilize a hidden Markov model (HMM) considering trajectories as observations and the corresponding road segments as states, demonstrating superiority over earlier geometric matching algorithms. In the realm of deep learning, the main approach involves sequence-to-sequence (Seq2Seq) methods Feng et al. (2020); Ren et al. (2021). With an input trajectory consisting of a sequence of geo-points, the encoder-decoder framework provides a sequence of matched road segments. Recently,

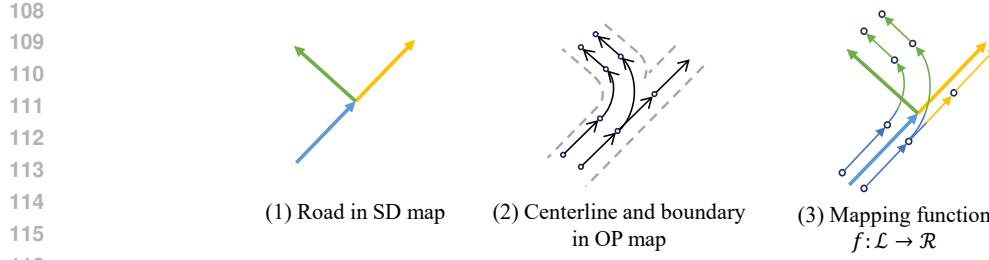


Figure 2: (a) The schema of SD map input: Road \mathcal{R} . (b) The schema of OP map input: Centerline \mathcal{L} and boundary \mathcal{B} . (c) The objective in our task: Mapping function f .

GraphMM Liu et al. (2023b) has used a graph-based approach that explicitly integrates all the correlations mentioned above, while EAM³ Tang et al. (2025) use a BERT-like transformer Alaparthi & Mishra (2020) to achieve the best performance. Compared to the map matching method, MAT uses a map-to-map paradigm to accommodate the many-to-one correspondence between lanes and roads in real time. Moreover, unlike the self-attention used in EAM³, MAT employs spatial and path-aware attention by incorporating the spatial curve and the path index, thus striking a balance between performance and latency.

Map Generation and Priors. The construction of online perception maps is a popular topic in autonomous driving Hao et al. (2024); Zhang et al. (2024a); Li et al. (2024) and is crucial for subsequent tasks Wan et al. (2024); Zhang et al. (2024b). HDMapNet Li et al. (2022) pioneered BEV-based map generation through sensor fusion, while LSS Philion & Fidler (2020) introduced depth-aware BEV transformation. VectorMapNet Liu et al. (2023a) enabled end-to-end vector prediction, and the MapTR Liao et al. (2023a) series introduced hierarchical query embeddings for instance-level construction. We adapt MapTRv2 Liao et al. (2023b) for online HD construction using the annotations of our dataset and demonstrate its compatibility with the SD-HD association. Recent works such as SMERF Luo et al. (2024), P-MapNet Jiang et al. (2024), and TopoSD Yang et al. (2024) leverage SD map priors to mitigate sensor noise and refine perception geometry. Unlike these approaches that focus on enhancing generation quality, MAT targets the explicit *association* phase, utilizing topological optimization to resolve assignment ambiguity arising from residual noise.

3 TASK DEFINITION

The goal of online navigation refinement is to transform road-topology-aligned navigation into lane-topology-aligned navigation. To achieve this, we utilize the map association as a map-to-map match paradigm. This involves correlating SD maps with online perception maps, allowing us to translate any route from an SD map into the corresponding route on an online perception map.

Standard Definition Map (SD Map). As shown in Fig.2 (a), SD maps represent road networks that use roads as primary primitives. Formally, a SD map is defined as a graph $\mathcal{G}_{\mathcal{R}} = (\mathcal{R}, \mathcal{E}_{\mathcal{R}})$, where $\mathcal{R} = \{r_1, r_2, \dots, r_m\}$ denotes a set of roads and $\mathcal{E}_{\mathcal{R}}$ encodes their topological connectivity. Each road r_j is parameterized by a sequence of directed vectors:

$$r_j = (\overrightarrow{q_{j1}q_{j2}}, \overrightarrow{q_{j2}q_{j3}}, \dots, \overrightarrow{q_{jk-1}q_{jk}}), \quad q_{jk} \in \mathbb{R}^2, \quad (1)$$

where consecutive points define road segments through uniform spatial sampling.

Online Perception Map (OP Map). As shown in Fig.2 (b), OP maps provide details of the lane level, primarily represented as a center line network. We model an OP map as a graph $\mathcal{G}_{\mathcal{L}} = (\mathcal{L}, \mathcal{E}_{\mathcal{L}})$, where $\mathcal{L} = \{l_1, l_2, \dots, l_n\}$ is a set of centerlines, each sampled at uniform intervals:

$$l_i = \overrightarrow{p_i^1 p_i^2}, \quad p_i^1, p_i^2 \in \mathbb{R}^2. \quad (2)$$

$\mathcal{E}_{\mathcal{L}}$ captures topological relations between the adjacent centerlines. In addition, we include road boundary vectors $\mathcal{B} = \{b_1, b_2, \dots, b_{m_b}\}$, which reflect the extent and shape of the actual road. Each boundary b_j is specified as:

$$b_j = (\overrightarrow{h_{j1}h_{j2}}, \overrightarrow{h_{j2}h_{j3}}, \dots, \overrightarrow{h_{jk-1}h_{jk}}), \quad h_{jk} \in \mathbb{R}^2. \quad (3)$$

Objective. As shown in Fig.2 (c), given $\mathcal{G}_{\mathcal{R}}$ and $\mathcal{G}_{\mathcal{L}}$, the task is to learn a mapping function $f: \mathcal{L} \rightarrow \mathcal{R}$ that assigns each centerline $l \in \mathcal{L}$ to its corresponding road $r_l \in \mathcal{R}$. The function satisfies

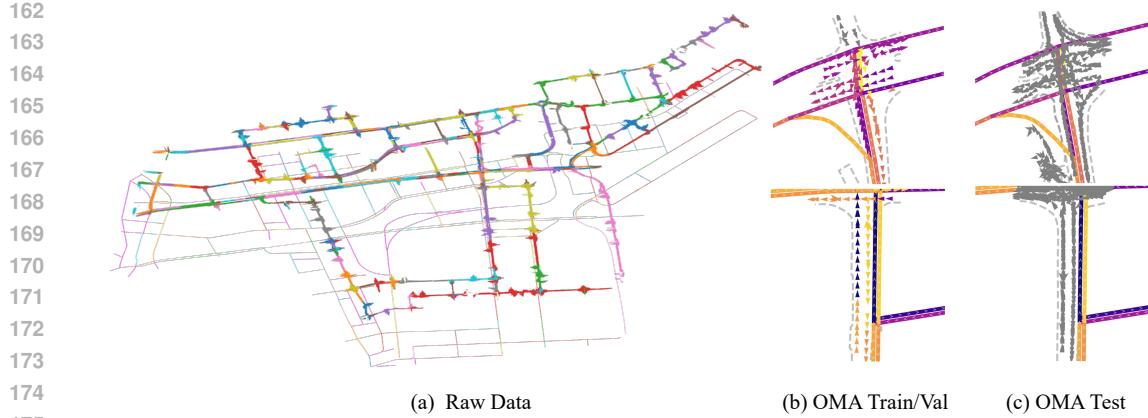


Figure 3: (a) The visualization of SD map and GT OP map with association annotations of Boston in nuScenes. The same color implies an associative pair. (b) The visualization of OMA Train/Val set. (c) The visualization of OMA Test set.

two key constraints: 1. *Uniqueness*: Each centerline l corresponds to exactly one ground truth road r_l ; 2. *Multiplicity*: A single road $r \in \mathcal{R}$ may be associated with multiple centerlines $l_1, l_2, \dots \in \mathcal{L}$.

This formulation casts the alignment task as a multi-to-one classification problem, where the number of classes equals the number of $|\mathcal{R}|$, and each centerline acts as an input sample. The goal is to maximize classification accuracy while preserving topological consistency between SD and OP maps. After we form this association, any path on the SD map allows us to locate all matching paths on the OP map through a single topological sorting.

4 DATASET AND METRIC

4.1 DATASET OVERVIEW

The source of ground-truth OP maps is nuScenes Caesar et al. (2020b), which includes locations in Boston and Singapore, featuring centerline geometries scanned with LiDAR. The SD maps were obtained from OpenStreetMap (OSM). We manually annotated the data to establish the association between the ground truth of the OP map and the SD map. The visualization is shown in Fig. 3 (a), which shows the global SD and ground-truth OP map of Boston with an association annotation.

The dataset is divided into training, validation, and test sets (see Fig. 3 (b) and (c)). Within the nuScenes framework, we employ the pon-split of nuScene. This approach designates distinct areas for training and validation/testing datasets, effectively preventing data leakage. The validation and test data sets are generated from the same scenes, but they have a variety of OP maps. For the OP map, both training and validation use ground-truth maps with manual correspondence annotations. For testing, the OP map is generated by MapTRv2 and its annotation is the same as the scene in the val set. It is essential to explain that the validation set and test set in OMA have equal significance. The purpose of the validation set is to assess the association capability of the OP map under noise-free conditions, whereas the test set evaluates this capability when the OP map experiences significant noise. This distinction sets apart the validation/test split in OMA from that in typical datasets. The specific analysis is provided in the Appendix D.

4.2 EVALUATION METRIC

In map matching, evaluation typically involves metrics such as Trajectory-Level Accuracy and the Ratio of Longest Common Subsequence, which presume a one-to-one mapping between trajectories and paths, concentrating only on the accuracy and sequence of path matching. Conversely, in map association, to assess various map generation methods, our goal is to use the association annotations of the ground-truth OP map to evaluate the results of the various OP maps with noise. Drawing inspiration from Reachability Precision Recall (P-R) Lu et al. (2023), we developed the Navigation Refinement P-R metric. This allows for the assessment of any map generation method’s perception map performance for the map association task, relying solely on ground-truth OP map annotations.

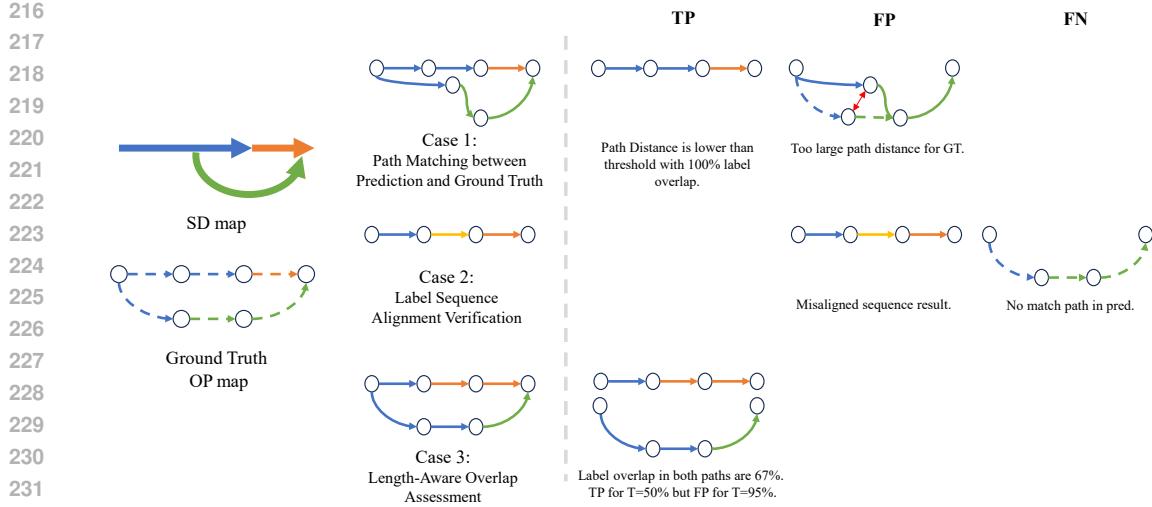


Figure 4: Example of TP, FP and FN for evaluate Navigation Refinement Precision-Recall.

Navigation Refinement P-R: Fig. 4 provides a detailed description of typical scenarios encountered during the Navigation Refinement Precision Recall (P-R) evaluation. The examination protocol involves three consecutive stages:

1. *Path Matching between Prediction and Ground Truth:* In the first stage, a connection is drawn between the predicted paths and GT paths within the OP map. Notably, only GT paths come with association annotations, making this matching essential for the following evaluations. A match is deemed successful if the bidirectional Chamfer distance is less than the distance threshold of 1 meter, consistent with Reachability P-R Lu et al. (2023). Predictions that do not meet this threshold are labeled FP, thereby excluding them from further SD-HD analysis. As illustrated in Row 1 of Fig. 4, Case 1 (left) shows that the prediction with minimal deviation successfully aligns with the GT paths. While Case 1 (middle) represents the heavily-deviated prediction results in an FP label.

2. *Label Sequence Alignment Verification:* Successful matches are converted into simplified label sequences, indicating SD map link traversals. Failures in sequence alignment at this point lead to FP/FN labels. As shown in Case 2 (middle) of Fig. 4, the case is a misaligned sequence result which is an FP, while Case 2 (right) shows that the GT without a match leads to a false negative (FN).

3. *Length-Aware Overlap Assessment:* For aligned sequences, we quantify the semantic consistency by computing the Label Overlap score. Specifically, this score is defined as the ratio of the length of correctly matched segments (where the predicted road ID matches the ground truth road ID) to the total path length. We compare the length proportions of road labels between the predicted and GT sequences to determine these matches. The true positive (TP) classification is determined based on thresholds. As shown in Case 3 (left) of Fig. 4, the case is a 67% overlap, which is TP at $T = 50\%$ but FP at $T=95\%$.

Following the mAP conventions Lin et al. (2014), we use $T = [0.5 : 0.05 : 0.95]$ (10 thresholds) and report mean P-R and F1 scores. To mitigate path-length bias, we separately compute metrics across 15 length intervals $L = [[0, 5), [5, 10), \dots, [70, +\infty)]$ before aggregation. The ablation study of distance threshold and length intervals is shown in the appendix G.

5 METHOD

5.1 OVERALL ARCHITECTURE

As depicted in Fig. 5, the Map Association Transformer (MAT) is a transformer specifically designed for real-time map association. All inputs are vectorized representations $\mathcal{V} = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_N\}$, where each vector \vec{v}_i is parameterized by two endpoints and direction: $\vec{v}_i = [p_{i1}^x, p_{i1}^y, p_{i2}^x, p_{i2}^y, \theta_i]$, with $\theta_i = \arctan\left(\frac{p_{i2}^x - p_{i1}^x}{p_{i2}^y - p_{i1}^y}\right)$ and $p_{i1}, p_{i2} \in \mathbb{R}^2$ being the start/end points. The input maps are composed of an SD map, an OP map, and a boundary. The SD map (\mathcal{G}_R) comprises road vectors $\mathcal{R} = \{r_1, \dots, r_{m_r}\}$, which form a graph with topological edges \mathcal{E}_R . Each road r_j is transformed

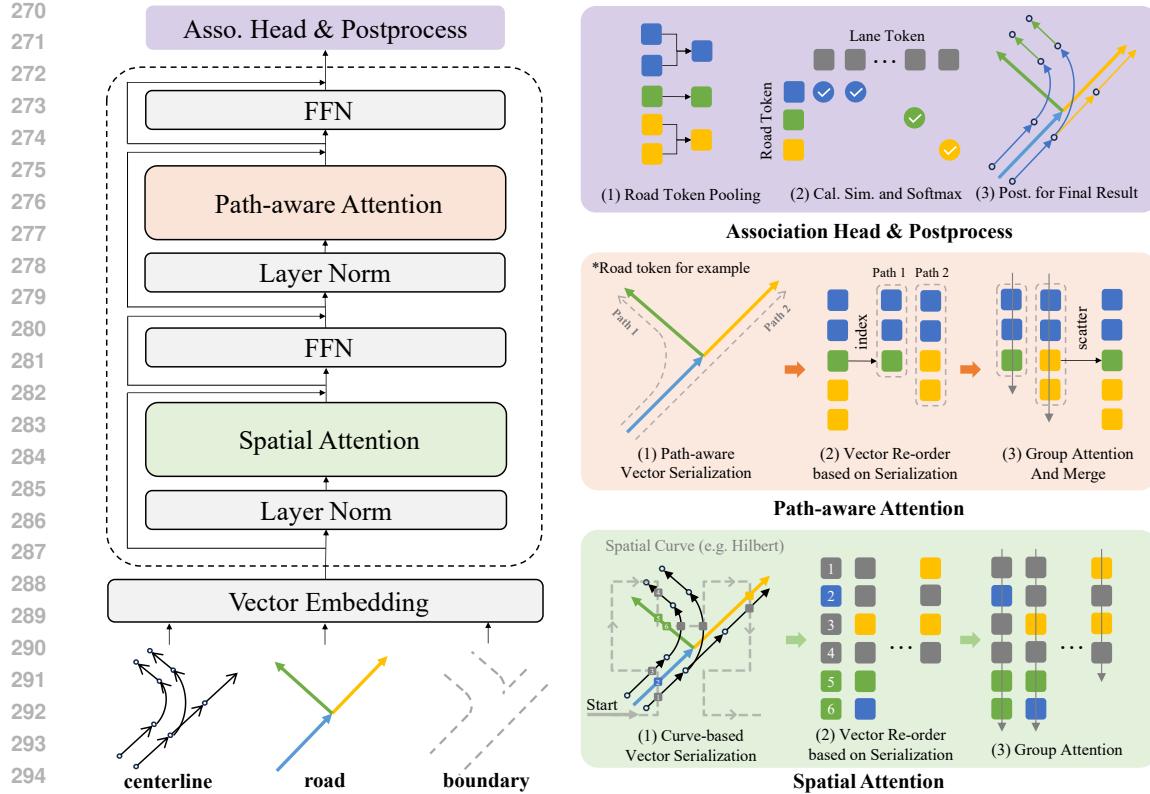


Figure 5: **Overview of Map Association Transformer (MAT).** The framework processes vectorized roads in SD map and centerlines/boundaries in OP map through N stacked layers containing Spatial Attention (for global context via curve-based serialization) and Path-Aware Attention (for topological alignment via path indexing). The Association Head then aggregates road features and calculates association probabilities with centerline tokens to generate the final navigation refinement result.

into an ordered sequence of vectors $\mathcal{R}_j = \overrightarrow{q_{j1}q_{j2}}, \overrightarrow{q_{j2}q_{j3}}, \dots$ through its parameterized segments. The OP map (\mathcal{G}_L) consists of centerline vectors $\mathcal{L} = \{l_1, \dots, l_{m_l}\}$ that represent the centerlines. Each centerline l_i is transformed into vectors $\mathcal{L}_i = \overrightarrow{p_i^1p_i^2}, \overrightarrow{p_i^2p_i^3}, \dots$ based on consecutive points p_i^j . The boundary (\mathcal{B}) includes the boundary vectors $\mathcal{B} = \{b_1, \dots, b_{m_b}\}$, converted similarly to the roads: $b_j \rightarrow \mathcal{B}_j = \overrightarrow{h_{j1}h_{j2}}, \overrightarrow{h_{j2}h_{j3}}, \dots$. These vectors are processed by the vector embedding module, which maps each 5D vector \vec{v}_i to a high-dimensional feature $F_{\vec{v}_i} \in \mathbb{R}^C$ via a two-layer MLP. The outputs are aggregated into feature matrices: $F_{\text{road}} \in \mathbb{R}^{N_r \times C}$, $F_{\text{centerline}} \in \mathbb{R}^{N_l \times C}$, and $F_{\text{boundary}} \in \mathbb{R}^{N_b \times C}$, where $N_{(.)}$ denotes the total number of vectors (e.g. $N_r = \sum_{j=1}^{m_r} \text{length}(\mathcal{R}_j)$).

Subsequently, the vector tokens are input into a transformer network. MAT consists of stacked MAT blocks to extract hierarchical features, each block containing Path-Aware Attention (PA), Spatial Attention (SA), and feed-forward network (FFN). In the association head, each road token requires the pooling to obtain a representative token \bar{F}_{road} for the current road. The association of the centerline with the road is calculated by combining the attention between \bar{F}_{road} and $F_{\text{centerline}}$, generating the probability distribution of the associations of the SD-OP map. The association probabilities are further refined by a post-processing method to enforce topological constraints. The details of the implementation and the module are provided in Appendix F.

5.2 PATH-AWARE ATTENTION

Path-Aware Attention (PA) is designed to extract locally stable geometric features that preserve the structural arrangement of elements.

Topological Ordering for Efficient Attention. To achieve real-time inference, our framework employs Group Attention, which reduces computational complexity from $O(N^2)$ to linear $O(N)$

324 by restricting interactions to local windows. However, this efficiency comes with a constraint: the
 325 effectiveness of group attention depends heavily on the semantic meaningfulness of token order.
 326 PA addresses this by introducing **Topological Ordering** as an inductive bias. Unlike random
 327 ordering, which would fragment the graph into unrelated segments, PA explicitly models long-range
 328 dependencies by constructing paths from root to leaf nodes. This ensures that topologically connected
 329 predecessor and successor nodes are placed adjacently in the sequence.

330 **Mechanism.** Specifically, we identify all valid paths from a starting point to an endpoint. We then
 331 reorder vector tokens to align with these path indices. This strict ordering allows the subsequent
 332 grouped attention (with size k) to focus solely on topologically relevant neighbors. Finally, tokens are
 333 reversed to their original sequence, and features of tokens appearing in multiple paths are averaged.
 334

335 5.3 SPATIAL ATTENTION

336 While PA captures topological connectivity, it may miss interactions between geometrically adjacent
 337 but topologically distant segments (e.g., parallel lanes or disconnected road boundaries). The Spatial
 338 Attention (SA) mechanism complements PA by capturing instance-level interactions across a wider
 339 spatial scope via **Spatial Ordering**.
 340

341 **Motivation: Vector Serialization as Geometric Clustering.** Similar to PA, SA relies on a specific
 342 ordering strategy to maximize the efficacy of group attention. However, instead of following graph
 343 connectivity, SA utilizes **Vector Serialization** to cluster tokens based on geometric proximity. This
 344 serves as a spatial inductive bias: it forces physically adjacent entities—even if they belong to
 345 different map layers or are disconnected—to be placed adjacently in the 1D sequence. Consequently,
 346 when the sequence is sliced into groups, highly correlated tokens naturally fall into the same attention
 347 bucket, enabling the model to handle GPS offsets and map alignment robustly.
 348

349 **Mechanism: Attention with Vector Serialization.** As illustrated in Fig. 5 green part, we implement
 350 this spatial ordering through the following steps:

351 1) *Coordinate Discretization*: Each vector token \vec{v}_i is encoded into a 3D discrete coordinate (x, y, r) ,
 352 representing the quantized grid location and orientation of the vector.

353 2) *Serialization via Space-Filling Curves*: We employ a space-filling curve function φ^{-1} (e.g., the
 354 Hilbert curve Hilbert (1935)) to map the 3D coordinates to a single 1D index. This mapping is crucial
 355 as it preserves spatial locality in the 1D domain better than simple row-major scanning.

356 3) *Grouped Attention & Restoration*: Tokens are reordered based on their 1D indices. Self-attention
 357 is performed within each group to aggregate spatial context, followed by an inverse operation to
 358 restore the original order.

359 5.4 ASSOCIATION AND LOSS FUNCTION

360 **Association.** The association between the roads and the centerline is calculated through a mechanism
 361 of cross-attention. For each road j , we first aggregate its token features $\{F_{j1}^{road}, \dots, F_{jN}^{road}\}$ into
 362 a representative feature $\bar{F}_j^{road} = \frac{1}{N} \sum_{n=1}^N F_{jn}^{road}$, where N denotes the number of road tokens on
 363 roads r_j and $F_{jn}^{road} \in \mathbb{R}^d$. The association probability $Prob_{ij}$ between the centerline i and the road
 364 j is then calculated as:

$$365 \quad Prob_{ij} = \exp \left(\frac{F_i^{cl} \cdot \bar{F}_j^{road}}{\sqrt{d}} \right) / \sum_{k=1}^K \exp \left(\frac{F_i^{cl} \cdot \bar{F}_k^{road}}{\sqrt{d}} \right), \quad (4)$$

366 where $F_i^{cl} \in \mathbb{R}^d$ is the token feature of the centerline, d is the dimension of the feature, and K
 367 represents the total number of roads. This formulation normalizes the similarity scores in all roads
 368 for each centerline i , ensuring a valid probability distribution.

369 **Loss Function.** Followed by Liu et al. (2023b); Ren et al. (2021), we optimize the model using a
 370 combination of cross-entropy loss (CE) and connection temporal classification (CTC) loss. The total
 371 loss is a weighted sum:

$$372 \quad \mathcal{L}_{\text{total}} = \alpha \cdot \mathcal{L}_{\text{CE}} + \beta \cdot \mathcal{L}_{\text{CTC}}, \quad (5)$$

373 with hyperparameters α and β balancing the two objectives. In practice, $\alpha = 1, \beta = 0.01$.
 374

378 Table 1: Result on OMA Val set. La. means latency. MM, PM, GM, MA means map matching, graph
 379 matching, point matching and map association method. P-M and M-M means the model is trained
 380 and inferred by path-to-map or map-to-map.

382 Methods	383 Present	384 Type	385 Paradigm	386 Val set	
				387 NR-F1 ^{50:95}	388 La./ms
HMM Newson (2009)	SIGSPATIAL'09	MM	P-M	70.1	465
DeepMM Feng et al. (2020)	TMC'20	MM	P-M	70.6	328
MTrajRec Ren et al. (2021)	KDD'21	MM	P-M	71.2	849
GraphMM Liu et al. (2023b)	TKDE'23	MM	P-M	72.3	469
EAM ³ Tang et al. (2025)	TITS'25	MM	P-M	72.9	345
KNN Cochran (1965)	–	GM	M-M	68.6	299
GMT He et al. (2024)	TPAMI'24	GM	M-M	68.8	93
FastMAC Zhang et al. (2024e)	CVPR'24	PM	M-M	70.3	62
MAT-T (Ours)	–	MA	M-M	78.2	34
MAT-L (Ours)	–	MA	M-M	78.7	70

395 Table 2: Result on OMA Test set. La. means latency. MM, PM, GM, MA means map matching,
 396 graph matching, point matching and map association method. P-M and M-M means the model is
 397 trained and inferred by path-to-map or map-to-map. SGG. means SeqGrowGrpah Xie et al. (2025a).

399 Methods	400 Present	401 Type	402 Paradigm	403 NR-F1 ^{50:95} on Test Sets			404 La./ms
				405 MapTR	406 MapTRv2	407 SGG.	
HMM Newson (2009)	SIGSPATIAL'09	MM	P-M	33.1	36.0	46.5	561
DeepMM Feng et al. (2020)	TMC'20	MM	P-M	32.8	33.5	43.9	733
MTrajRec Ren et al. (2021)	KDD'21	MM	P-M	34.1	36.8	48.9	1593
GraphMM Liu et al. (2023b)	TKDE'23	MM	P-M	35.9	38.7	49.2	889
EAM ³ Tang et al. (2025)	TITS'25	MM	P-M	36.3	39.1	50.8	679
KNN Cochran (1965)	–	GM	M-M	31.7	34.6	43.5	313
GMT He et al. (2024)	TPAMI'24	GM	M-M	30.1	33.4	42.3	105
FastMAC Zhang et al. (2024e)	CVPR'24	PM	M-M	32.0	35.8	45.0	81
MAT-T (Ours)	–	MA	M-M	41.5	44.8	54.8	35
MAT-L (Ours)	–	MA	M-M	41.9	45.0	54.9	74

412 5.5 TOPOLOGY POST-PROCESS

414 We formalize topological decoding as a structured prediction on the entire path of the centerline
 415 \mathcal{P}_j , $j \in [1, \dots, K]$. K is the number of total paths. The two-stage decoding process operates as
 416 follows:

417 **Token Initialization.** For each centerline path \mathcal{P}_j , we select the initial centerline T_{\max} via:

$$421 T_{\max} = \operatorname{argmax}_{l \in \mathcal{P}_j} \max_{r \in \mathcal{R}} P(l, r) \quad (6)$$

424 where $P(l, r)$ is the probability of association from centerline l to road r .

426 **Topological-constraint Beam Search.** Based on beam search, topological constraint beam search
 427 makes the following two improvements: 1. Modify the one-way search to implement a bidirectional
 428 search starting at T_{\max} . 2. When generating new predictions, instead of using the approach of
 429 taking the maximum value from all roads, we decode under the constraint of connectivity provided in
 430 the road network \mathcal{E}_r , thus ensuring that the connectivity of the road sequence corresponding to the
 431 lane path in the decoding result is consistent with the representation of the road network. Detailed
 expressions of the topological constraints beam search, including formula descriptions, are included
 in the appendix F.

432 Table 3: Ablation study of structure. Post. means post Table 4: Ablation study of loss and pool method.
 433 process. Bd. means Boundary. La. means latency. La. means latency.

435	SA	PA	Bd.	Post.	NR-F1 ^{50:95}	La./ms	CE	CTC	Avg.	Max	NR-F1 ^{50:95}	La./ms
436				Baseline (PTv3)	61.8	59	✓		✓		78.4	70
437	✓				62.1	77		✓	✓		67.7	70
438		✓			74.1	61	✓	✓		✓	78.5	70
439	✓	✓			77.8	64					78.7	70
440	✓	✓	✓		78.5	69						
441	✓	✓	✓	✓	78.7	70						

444 6 EXPERIMENT

445 6.1 IMPLEMENTATION DETAILS

446 We train our models from scratch for a total of 50 epochs using the AdamW optimizer. A cosine-
 447 decay learning rate scheduler is employed, incorporating a linear warm-up phase of two epochs.
 448 The initial learning rate, weight decay, and batch size are set to 0.0001, 0.05, and 128, respectively.
 449 All experiments are conducted on NVIDIA A6000 GPUs. The latency is measured on an NVIDIA
 450 A6000 GPU paired with an Intel(R) Xeon(R) Platinum 8369B CPU. Note that MAT-T and MAT-L
 451 share identical architectural components and training configurations, differing only in the number of
 452 Transformer blocks to balance real-time efficiency and model capacity, which details are shown in
 453 Tab. 22 in Appendix.

454 6.2 RESULT

455 To evaluate the performance of different methods on the OMA dataset, we categorize existing
 456 approaches into map matching methods Newson (2009); Feng et al. (2020); Ren et al. (2021); Liu
 457 et al. (2023b); Tang et al. (2025), graph matching methods He et al. (2024), and point matching
 458 methods Zhang et al. (2024e). For map matching, we traverse every path on the OP map as a GPS
 459 trajectory to match against the SD map. For graph and point matching, we treat the SD and OP maps
 460 as separate graphs or point clouds.

461 **Val set.** As shown in Table 1, MAT achieves the optimal balance between precision and efficiency on
 462 the validation set. Specifically, compared to traditional, Seq2Seq, and graph-based map matching
 463 methods, MAT-T demonstrates significant improvements in NR-F1^{50:95}, outperforming HMM New-
 464 son (2009) by 8.1%, DeepMM Feng et al. (2020) by 7.6%, MTrajRec Ren et al. (2021) by 7.0%,
 465 and the recent EAM³ Tang et al. (2025) by 5.3%, all while maintaining a low inference latency of
 466 34 ms. Moreover, MAT-T achieves superior performance compared to graph matching techniques
 467 like GMT He et al. (2024) and point matching approaches such as FastMAC Zhang et al. (2024e),
 468 proving its effectiveness in handling the heterogeneity between SD and OP maps which typically
 469 challenges traditional matching or point cloud algorithms.

470 **Test set.** To further verify the model’s robustness against varying noise patterns crucial for real-world
 471 deployment, we extended the evaluation on the Test set to include OP maps generated by three distinct
 472 methods: MapTR Liao et al., MapTRv2 Liao et al. (2023b), and SeqGrowGraph Xie et al. (2025a),
 473 as summarized in Table 2. The results demonstrate that our proposed MAT method consistently
 474 outperforms all state-of-the-art baselines across these disparate generators without requiring specific
 475 fine-tuning. On the SeqGrowGraph part, MAT-L achieves an NR-F1^{50:95} of 54.9%, surpassing the
 476 strongest baseline EAM³ Tang et al. (2025) by 4.1%, while on the more challenging MapTR dataset,
 477 which exhibits different topological error patterns, MAT-L maintains a significant lead of 5.6%
 478 over EAM³. This consistent superiority confirms that MAT’s architecture effectively generalizes
 479 across distinct noise distributions—ranging from fragmentation in MapTR to connectivity issues in
 480 growing-based methods—thereby validating the strong cross-generator generalization ability of the
 481 proposed Online Navigation Refinement framework.

482 6.3 ABLATION STUDY

483 The ablation study experiment is conducted with MAT-L in the val set of OMA, with latency measured
 484 using a NVIDIA A6000.

486 Table 5: Cross-validation between Boston and Singapore in the OMA dataset
487

488 City	489 Val	490 Test
490 Singapore → Singapore	491 77.5	492 49.5
491 Boston → Singapore	492 77.3	493 48.7
492 Boston → Boston	493 79.4	494 42.9
493 Singapore → Boston	494 79.0	495 42.6

494 Table 6: Data Efficiency experiment on OMA dataset.
495

496 Ratio	497 1%	498 2%	499 5%	500 10%	501 20%	502 50%	503 100%
500 Val / NR-F1 ^{50:95}	501 24.9	502 64.4	503 77.1	504 77.7	505 77.3	506 78.3	507 78.7
507 Test / NR-F1 ^{50:95}	508 22.3	509 41.5	510 44.5	511 44.6	512 46.6	513 44.7	514 45.0

501 **Structure.** As shown in Tab. 3, ablating path-aware (PA) and spatial attention (SA) reveals that
502 PA+SA achieves the highest NR-F1^{50:95} (+3.7% vs. PA-only, +15.7% vs. SA-only). PA-only
503 outperforms SA-only (74.1% vs. 62.1%), confirming the critical role of topological awareness.
504 Boundary improve +0.7% with a 5ms latency cost. Post-processing further enhances accuracy
505 (+0.2%) without efficiency trade-offs.

506 **Loss Function and Road Pooling Method.** Tab. 4 shows an ablation study on loss functions and
507 road pooling methods. The model exhibits strong robustness, with consistent performance across
508 different settings, indicating low sensitivity to these components. Combining CE and CTC losses
509 improves NR-F1 by +0.3% over CE alone and by +11.0% over CTC alone, while all variants perform
510 competitively. Average pooling outperforms maximum pooling by only +0.2%, further confirming the
511 model’s insensitivity to pooling strategy. These results suggest that the effectiveness of the method is
512 due to its inherently robust design, rather than specific loss or grouping choices.

513 **Cross-validation.** To evaluate the generalization of the model in geographically and behaviorally
514 distinct driving environments, we conducted cross-validation in Tab. 5 using data from Boston and
515 Singapore, two cities that differ markedly in road layout, traffic density, and regulations. The model
516 maintains robust performance without fine-tuning: in the val set, the Boston-Singapore transfer
517 incurred only a 0.4-point drop (79.4→79.0) and Singapore-Boston a 0.2-point drop (77.5→77.3);
518 on the test set, the drops were similarly minor: 0.8 points (49.5→48.7) and 0.3 points (42.9→42.6),
519 respectively.

520 **Data Efficiency.** To investigate the label efficiency of MAT and validate its potential for semi-
521 supervised learning, we conducted experiments using varying subsets of the training data (ranging
522 from 1% to 100%) on MAT-L. The results, summarized in Tab. 6, demonstrate exceptional data
523 efficiency. In particular, with only 5% of the annotated training data, the model achieves a *NR* –
524 *F1*^{50:95} of 77.1% in the validation set and 44.5% in the test set. These results are comparable to
525 the performance achieved using 100% of the data (78.7% / 45.0%). This saturation in a low-data
526 regime indicates that MAT effectively captures robust topological features without requiring massive
527 amounts of dense annotations.

528 **More Ablation Study, Visualization and Failed cases.** Appendix G has shown more study of
529 ablation study of hyperparameters of PA and SA, input size of SD map, beam width of post-process,
530 visualizations, and failed cases.

532 7 CONCLUSION

534 We propose Online Navigation Refinement (ONR), a new mission that fuses static SD maps with
535 real-time perception for accurate, low-cost lane-level navigation. To achieve ONR, we make three
536 core contributions: We release OMA, the first public benchmark for map associations, and MAT, a
537 transformer-based model for aligning noisy maps. We also introduce Navigation Refinement P-R, a
538 metric evaluating both geometry and association accuracy. For limitations, we will explore end-to-end
539 integration with motion planners, cross-domain adaptation, and unsupervised association learning to
reduce annotation dependence.

540
541
ETHICS STATEMENT

542 This work uses only publicly available autonomous driving datasets (e.g., nuScenes) and Open-
 543 StreetMap (OSM), following their licenses and usage terms. We rely solely on map-level and
 544 sensor-derived geometric information and do not use or release any personally identifiable data such
 545 as faces, license plates, or raw trajectories linked to individuals. All additional annotations were
 546 produced by trained annotators on de-identified map representations. Our method and dataset are
 547 intended for research purposes to improve road safety and navigation; any real-world deployment
 548 should include thorough testing, safety checks, and compliance with relevant regulations to avoid
 549 potential harms from incorrect or misleading guidance.

550
551
REPRODUCIBILITY STATEMENT
552

553 This work is based entirely on publicly available datasets, specifically the nuScenes autonomous
 554 driving dataset and OpenStreetMap (OSM). We describe our model architectures, training objectives,
 555 and optimization hyperparameters, as well as dataset preprocessing, train/validation/test splits, and
 556 evaluation metrics in the main text and appendix to enable independent verification. Upon publication,
 557 we will release our code, including scripts for data preprocessing, map association construction,
 558 model training, and evaluation, together with configuration files specifying all hyperparameters,
 559 random seeds, and implementation details (e.g., framework and library versions). We will also release
 560 the processed map-association annotations introduced in this paper, together with the code of our
 561 models and trained checkpoints.

562
563
REFERENCES

564 Openstreetmap. <https://github.com/openmaptiles/openmaptiles>.
 565

566 Shivaji Alaparthi and Manit Mishra. Bidirectional encoder representations from transformers (bert):
 567 A sentiment analysis odyssey. *arXiv preprint arXiv:2007.01127*, 2020.

568 Xiao Bai, Pengcheng Zhang, Xiaohan Yu, Jin Zheng, Edwin R Hancock, Jun Zhou, and Lin Gu.
 569 Learning from human attention for attribute-assisted visual recognition. *IEEE Transactions on
 570 Pattern Analysis and Machine Intelligence*, 2024.

571 Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liang, Qiang Xu, Anush
 572 Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for
 573 autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
 574 recognition*, pp. 11621–11631, 2020a.

575 Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liang, Qiang Xu, Anush
 576 Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for
 577 autonomous driving. In *CVPR*, 2020b.

578 Pingfu Chao, Yehong Xu, Wen Hua, and Xiaofang Zhou. A survey on map-matching algorithms. In
 579 *Australasian database conference*, pp. 121–133. Springer, 2020.

580 Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu,
 581 and Xinggang Wang. Vadv2: End-to-end vectorized autonomous driving via probabilistic planning.
 582 *arXiv preprint arXiv:2402.13243*, 2024.

583 William G Cochran. The planning of observational studies of human populations. *Journal of the
 584 Royal Statistical Society. Series A (General)*, 128(2):234–266, 1965.

585 Patrick Dendorfer, Aljosa Osep, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, Stefan
 586 Roth, and Laura Leal-Taixé. Motchallenge: A benchmark for single-camera multiple target
 587 tracking. *International Journal of Computer Vision*, 129:845–881, 2021.

588 Gamal Elghazaly, Raphaël Frank, Scott Harvey, and Stefan Safko. High-definition maps: Compre-
 589 hensive survey, challenges, and future perspectives. *IEEE Open Journal of Intelligent Transporta-
 590 tion Systems*, 4:527–550, 2023.

594 Jie Feng, Yong Li, Kai Zhao, Zhao Xu, Tong Xia, Jinglin Zhang, and Depeng Jin. Deepmm: Deep
 595 learning based map matching with data augmentation. *IEEE Transactions on Mobile Computing*,
 596 21(7):2372–2384, 2020.

597

598 Yuan Guo, Ziyue Tian, Bijun Li, Jian Zhou, Zhangcai Yin, Quanhua Dong, and Shen Ying. Lane-level
 599 map matching for vehicles using mask-based raster high-definition maps. *Expert Systems with*
 600 *Applications*, pp. 128195, 2025.

601 Anders Hansson, Ellen Korsberg, Roza Maghsoud, Eliza Nordén, et al. Lane-level map matching
 602 based on hmm. *IEEE Transactions on Intelligent Vehicles*, 6(3):430–439, 2020.

603

604 Xiaoshuai Hao, Ruihai Li, Hui Zhang, Dingzhe Li, Rong Yin, Sangil Jung, Seung-In Park, ByungIn
 605 Yoo, Haimei Zhao, and Jing Zhang. Mapdistill: Boosting efficient camera-based hd map construc-
 606 tion via camera-lidar fusion model distillation. In *European Conference on Computer Vision*, pp.
 607 166–183. Springer, 2024.

608

609 Jiawei He, Zehao Huang, Naiyan Wang, and Zhaoxiang Zhang. Learnable graph matching: A
 610 practical paradigm for data association. *IEEE Transactions on Pattern Analysis and Machine*
611 Intelligence, 46(7):4880–4895, 2024.

612 David Hilbert. Über die stetige abbildung einer linie auf ein flächenstück. *Dritter Band: Analysis-*
613 Grundlagen der Mathematik. Physik Verschiedenes: Nebst Einer Lebensgeschichte, pp. 1–2, 1935.

614

615 Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu,
 616 Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous
 617 driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp.
 618 8340–8350, 2023.

619

620 Zhou Jiang, Zhenxin Zhu, Pengfei Li, Huan-ang Gao, Tianyuan Yuan, Yongliang Shi, Hang Zhao,
 621 and Hao Zhao. P-mapnet: Far-seeing map generator enhanced by both sdmap and hdmap priors.
622 IEEE Robotics and Automation Letters, 2024.

623

624 Qi Li, Yue Wang, Yilun Wang, and Hang Zhao. Hdmapnet: An online hd map construction and
 625 evaluation framework. In *2022 International Conference on Robotics and Automation (ICRA)*, pp.
 4628–4634. IEEE, 2022.

626

627 Ruihai Li, Hao Shan, Han Jiang, Jianru Xiao, Yizhuo Chang, Yifan He, Haiyang Yu, and Yilong Ren.
 628 E-mlp: Effortless online hd map construction with linear priors. In *2024 IEEE Intelligent Vehicles*
629 Symposium (IV), pp. 1008–1014. IEEE, 2024.

630

631 Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang
 632 Huang. Maptr: Structured modeling and learning for online vectorized hd map construction.
In The Eleventh International Conference on Learning Representations.

633

634 Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang
 635 Huang. Maptr: Structured modeling and learning for online vectorized HD map construction. In
636 ICLR, 2023a.

637

638 Bencheng Liao, Shaoyu Chen, Yunchi Zhang, Bo Jiang, Qian Zhang, Wenyu Liu, Chang Huang, and
 639 Xinggang Wang. Maptrv2: An end-to-end framework for online vectorized hd map construction.
arXiv preprint arXiv:2308.05736, 2023b.

640

641 Bencheng Liao, Shaoyu Chen, Haoran Yin, Bo Jiang, Cheng Wang, Sixu Yan, Xinbang Zhang,
 642 Xiangyu Li, Ying Zhang, Qian Zhang, et al. Diffusiondrive: Truncated diffusion model for
 643 end-to-end autonomous driving. In *Proceedings of the Computer Vision and Pattern Recognition*
644 Conference, pp. 12037–12047, 2025.

645

646 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 647 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision-
 648 ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings,
 649 part v 13*, pp. 740–755. Springer, 2014.

648 Yicheng Liu, Tianyuan Yuan, Yue Wang, Yilun Wang, and Hang Zhao. Vectormapnet: End-to-end
 649 vectorized hd map learning. In *International Conference on Machine Learning*, pp. 22352–22369.
 650 PMLR, 2023a.

651 Yu Liu, Qian Ge, Wei Luo, Qiang Huang, Lei Zou, Haixu Wang, Xin Li, and Chang Liu. Graphmm:
 652 Graph-based vehicular map matching by leveraging trajectory and road correlations. *IEEE Trans-*
 653 *actions on Knowledge and Data Engineering*, 36(1):184–198, 2023b.

654 Jiachen Lu, Renyuan Peng, Xinyue Cai, Hang Xu, Hongyang Li, Feng Wen, Wei Zhang, and Li Zhang.
 655 Translating images to road network: A non-autoregressive sequence-to-sequence approach. In
 656 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 23–33, 2023.

657 Katie Z Luo, Xinshuo Weng, Yan Wang, Shuang Wu, Jie Li, Kilian Q Weinberger, Yue Wang, and
 658 Marco Pavone. Augmenting lane perception and topology understanding with standard definition
 659 navigation maps. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp.
 660 4029–4035. IEEE, 2024.

661 Paul Newson. Hidden markov map matching through noise and sparseness. In Divyakant Agrawal,
 662 Walid G. Aref, Chang-Tien Lu, Mohamed F. Mokbel, Peter Scheuermann, Cyrus Shahabi, and
 663 Ouri Wolfson (eds.), *17th ACM SIGSPATIAL International Symposium on Advances in Geographic
 664 Information Systems, ACM-GIS 2009, November 4-6, 2009, Seattle, Washington, USA, Proceedings*,
 665 pp. 336–343. ACM, 2009. doi: 10.1145/1653771.1653818. URL <https://doi.org/10.1145/1653771.1653818>.

666 Qucheng Peng, Chen Bai, Guoxiang Zhang, Bo Xu, Xiaotong Liu, Xiaoyin Zheng, Chen Chen, and
 667 Cheng Lu. Navigscene: Bridging local perception and global navigation for beyond-visual-range
 668 autonomous driving. *arXiv preprint arXiv:2507.05227*, 2025.

669 Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by
 670 implicitly unprojecting to 3d. In *ECCV*, 2020.

671 Huimin Ren, Sijie Ruan, Yanhua Li, Jie Bao, Chuishi Meng, Ruiyuan Li, and Yu Zheng. Mtrajrec:
 672 Map-constrained trajectory recovery via seq2seq multi-task learning. In *Proceedings of the 27th
 673 ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 1410–1419, 2021.

674 Thomas Roddick and Roberto Cipolla. Predicting semantic map representations from images using
 675 pyramid occupancy networks. In *Proceedings of the IEEE/CVF conference on computer vision
 676 and pattern recognition*, pp. 11138–11147, 2020.

677 Ziying Song, Caiyan Jia, Lin Liu, Hongyu Pan, Yongchang Zhang, Junming Wang, Xingyu Zhang,
 678 Shaoqing Xu, Lei Yang, and Yadan Luo. Don’t shake the wheel: Momentum-aware planning in
 679 end-to-end autonomous driving. In *Proceedings of the Computer Vision and Pattern Recognition
 680 Conference*, pp. 22432–22441, 2025.

681 Wenchao Sun, Xuewu Lin, Yining Shi, Chuang Zhang, Haoran Wu, and Sifa Zheng. Sparsedrive:
 682 End-to-end autonomous driving via sparse scene representation. In *2025 IEEE International
 683 Conference on Robotics and Automation (ICRA)*, pp. 8795–8801. IEEE, 2025.

684 Jie Tang, Sunjian Zheng, Bo Yu, and Xue Liu. Elevation-aware map matching model leveraging
 685 transfer learning in sparse data conditions. *IEEE Transactions on Intelligent Transportation
 686 Systems*, 2025.

687 Jiaxu Wan, Hong Zhang, Ziqi He, Qishu Wang, Ding Yuan, and Yifan Yang. Sp2t: Sparse proxy
 688 attention for dual-stream point transformer. *arXiv preprint arXiv:2412.11540*, 2024.

689 Huijie Wang, Tianyu Li, Yang Li, Li Chen, Chonghao Sima, Zhenbo Liu, Bangjun Wang, Peijin
 690 Jia, Yuting Wang, Shengyin Jiang, Feng Wen, Hang Xu, Ping Luo, Junchi Yan, Wei Zhang, and
 691 Hongyang Li. Openlane-v2: A topology reasoning benchmark for unified 3d hd mapping. In
 692 *NeurIPS*, 2023.

693 Kelvin Wong, Yanlei Gu, and Shunsuke Kamijo. Mapping for autonomous driving: Opportunities
 694 and challenges. *IEEE Intelligent Transportation Systems Magazine*, 13(1):91–106, 2020.

702 Mengwei Xie, Shuang Zeng, Xinyuan Chang, Xinran Liu, Zheng Pan, Mu Xu, and Xing Wei.
 703 Seqgrowgraph: Learning lane topology as a chain of graph expansions. In *Proceedings of the*
 704 *IEEE/CVF International Conference on Computer Vision*, pp. 27166–27175, 2025a.
 705

706 Shaoyuan Xie, Lingdong Kong, Wenwei Zhang, Jiawei Ren, Liang Pan, Kai Chen, and Ziwei Liu.
 707 Benchmarking and improving bird’s eye view perception robustness in autonomous driving. *IEEE*
 708 *Transactions on Pattern Analysis and Machine Intelligence*, 2025b.
 709

710 Sen Yang, Minyue Jiang, Ziwei Fan, Xiaolu Xie, Xiao Tan, Yingying Li, Errui Ding, Liang Wang,
 711 and Jingdong Wang. Toposd: Topology-enhanced lane segment perception with sdmap prior. *arXiv*
 712 preprint *arXiv:2411.14751*, 2024.
 713

714 Hong Zhang, Jiaxu Wan, Ziqi He, Jianbo Song, Yifan Yang, and Ding Yuan. Sparse agent transformer
 715 for unified voxel and image feature extraction and fusion. *Information Fusion*, pp. 102455, 2024a.
 716

717 Hong Zhang, Jiaxu Wan, Jing Zhang, Ding Yuan, XuLiang Li, and Yifan Yang. P2ftrack: Multi-object
 718 tracking with motion prior and feature posterior. *ACM Transactions on Multimedia Computing,
 Communications and Applications*, 2024b.
 719

720 Jiawei Zhang, Lei Huang, Xiao Bai, Jin Zheng, Lin Gu, and Edwin Hancock. Exploring the usage
 721 of pre-trained features for stereo matching. *International Journal of Computer Vision*, 132(10):
 722 4305–4326, 2024c.
 723

724 Jiawei Zhang, Jiahe Li, Lei Huang, Haonan Luo, Xiaohan Yu, Lin Gu, Jin Zheng, and Xiao Bai.
 725 Investigating synthetic-to-real transfer robustness for stereo matching and optical flow estimation.
 726 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025a.
 727

728 Pengcheng Zhang, Xiaohan Yu, Xiao Bai, Jin Zheng, Xin Ning, and Edwin R Hancock. Fully
 729 decoupled end-to-end person search: An approach without conflicting objectives. *International*
 730 *Journal of Computer Vision*, pp. 1–22, 2025b.
 731

732 Shiming Zhang, Zhipeng Luo, Li Yang, Fei Teng, and Tianrui Li. A survey of route recommendations:
 733 Methods, applications, and opportunities. *Information Fusion*, 108:102413, 2024d.
 734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756	CONTENTS OF APPENDIX	
757		
758		
759	A LLM Usage	16
760		
761	B Question and Answer	16
762	B.1 Motivation	16
763	B.2 Distribution	16
764	B.3 Maintenance	16
765	B.4 Composition	17
766	B.5 Collection Process	17
767	B.6 Use	17
768		
769		
770		
771	C Related Work	18
772		
773	C.1 End-to-end Autonomous Driving	18
774	C.2 Mission Comparison	18
775		
776		
777	D Dataset	19
778		
779	D.1 Revised Analysis.	19
780	D.2 Range of Sample	20
781	D.3 Pon split and OP maps	20
782	D.4 Geographical generalization	20
783	D.5 Annotation Workflow	20
784	D.6 Handling Transition Ambiguity	20
785	D.7 Future Expansion Strategy	21
786		
787		
788		
789	E Metric	21
790		
791	E.1 Details of NR-PR	21
792	E.2 Curvature-based Weighted Metric	21
793		
794		
795	F Method	23
796		
797	F.1 Path-aware attention	23
798	F.2 Spatial attention	24
799	F.3 Post Process	24
800		
801	G Ablation Study	25
802		
803	G.1 Hyper Parameter	25
804	G.2 Robustness to Imperfect SD Maps	27
805	G.3 Robustness under Environmental Noise and Occlusion	28
806		
807		
808	H Visualization	29
809		
	H.1 Attention Map	29

810	H.2 Result Compare	30
811	H.3 Failed Cases	32
812	H.4 Generalization on Different OP Generators	33
813		
814		
815	I Implement Details	33
816		
817	I.1 Model Settings	33
818	I.2 Train Configuration	35
819	I.3 Data Augmentations	35
820		
821		
822	A LLM USAGE	
823		
824	The paper has employed LLM tools such as GPT and Writefull for enhancing and refining writing in	
825	the abstract, main text, and appendix.	
826		
827	B QUESTION AND ANSWER	
828		
829	We have included some Q&A to enable readers to grasp the information in our dataset. Every	
830	response follows the double-blind principle.	
831		
832	B.1 MOTIVATION	
833		
834	For what purpose was the dataset created? We introduce this dataset to refine online navigation	
835	by transforming road-level paths on SD maps into lane-level paths on online perception maps. This	
836	approach facilitates cost-effective and highly real-time lane-level navigation. Current lane-level	
837	navigation systems are highly dependent on expensive and globally updated HD maps, which are	
838	costly and lag in real-time efficiency. Online perception maps, a significant focus in autonomous	
839	driving research, create local HD maps near vehicles using real-time onboard sensor data. However,	
840	they lack global topology connections, making them unsuitable for navigation. This dataset envi-	
841	sions aligning SD maps with online perception maps through a map-to-map association paradigm,	
842	translating road-level paths on SD maps into lane-level paths on online perception maps.	
843	What is the relationship between this paper and planning methods? The task and methodology	
844	of this paper focus on delivering road-level navigation on vehicles and drivers' devices. Addresses	
845	the issue of road-level navigation provision in a context where HD maps are no longer used. Recently,	
846	research like NavigScene Peng et al. (2025) has incorporated SD navigation data into autonomous	
847	driving planning modules. We suggest that our approach can act as a preliminary stage for NavigScene,	
848	enhancing the SD navigation data to offer lane-level guidance. This grants NavigScene's planning	
849	module more detailed navigation details, thereby increasing planning precision.	
850	B.2 DISTRIBUTION	
851		
852	Will the dataset be distributed to third parties outside the entity (e.g., company, institution,	
853	organization) on behalf of which the dataset was created? Yes, the dataset can be accessed	
854	publicly on the Internet.	
855		
856	How will the dataset be distributed (for example, tarball on website, API, GitHub)? The dataset	
857	will be released on GitHub and Huggingface.	
858	B.3 MAINTENANCE	
859		
860	Who will be supporting/hosting/maintaining the dataset? The authors will be supporting, hosting,	
861	and maintaining the dataset.	
862		
863	How can the owner / curator / manager of the dataset be contacted (e.g., email address)? You	
	can contact the author by email which will be included in the accepted version of the paper.	

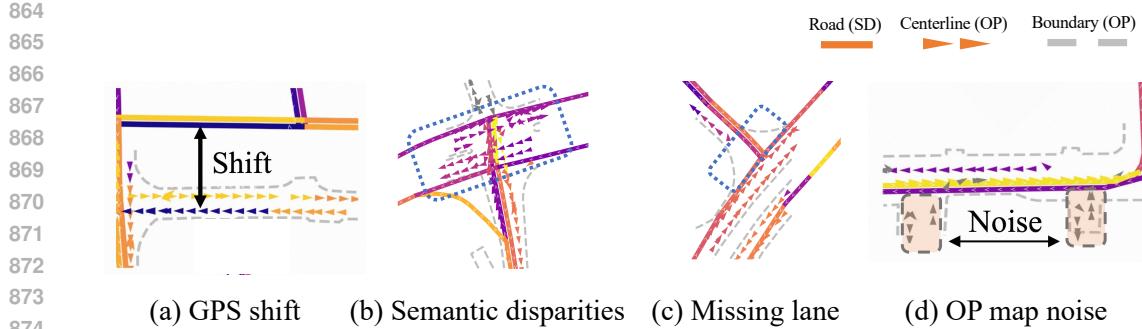


Figure 6: Challenge for map association. (a) GPS shift, (b) Semantic disparities, (c) missing lane, (d) noise in OP map

Is there an erratum? No. We will make a statement if there is any.

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?

Yes. we intend to refresh OMA’s on-line HD map with more robust baselines in the future. The new method will be selected based on comprehensive criteria that include performance measures such as accuracy, efficiency, and generalization.

Will older versions of the dataset continue to be supported/hosted/maintained? Yes.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? Yes. We will make a guide document on the website of OMA.

B.4 COMPOSITION

What do the instances that comprise the dataset represent? In OMA, the basic element is a vector that encompasses the road vectors, the lane vectors, and the boundary vectors. Each vector is defined by two points that indicate its position on the map. We have normalized each vector in OMA to match the ego perspective used in real-world perception settings. Furthermore, each vector is equipped with a set of connectivity relationships that specify the other vectors to which it is connected within the road network. Road vectors include an additional road ID to identify their associated road. Lane vectors have been manually marked to show their association with a specific road.

How many instances are there in total (of each type, if appropriate)? OMA contain over 30K scenarios, 480K road path and 2.6M lane vector with manually annotated associations.

Are relationships between individual instances made explicit? On the SD map, all roads, and on the GT OP map, all lanes within the training and validation sets have thorough descriptions and associated annotations. For the test set, we offer descriptions of roads and lanes, but without associated annotations. Nevertheless, the test set can still be evaluated using annotations from the ground truth OP map in the validation set, thanks to the NR P-R metrics introduced in our paper.

Are there recommended data splits (for example, training, development / validation, testing)? Yes. We have already partitioned our dataset into three distinct splits: training, validation, and testing.

B.5 COLLECTION PROCESS

Who was involved in the data collection process (e.g., students, crowd workers, contractors)? The annotations are provided by experienced annotators and multiple validation stages.

B.6 USE

What (other) tasks could the dataset be used for? This dataset is primarily used for map association, a new task for online navigation refinement.

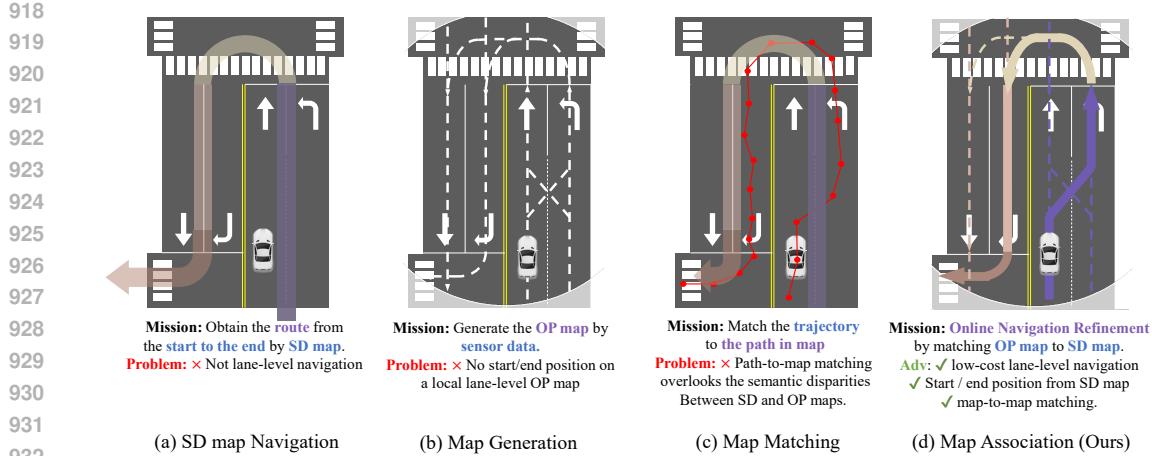


Figure 7: The comparison of (a) SD map navigation, (b) mapping generation, (c) map matching and (d) map association. The blue word is the input of the task and the purple word is the output of the task

C RELATED WORK

C.1 END-TO-END AUTONOMOUS DRIVING

Planning, being an essential component for numerous fields Zhang et al. (2025a; 2024c; 2025b); Bai et al. (2024), is recognized as one of the pivotal research domains in the realm of autonomous driving at present. VAD Jiang et al. (2023) utilizes an ego query mechanism to forecast individual mode trajectories, whereas VAdv2 Chen et al. (2024) enhances this by adopting a probabilistic framework that considers multiple trajectories. SparseDrive Sun et al. (2025) is innovative in its creation of parallel motion and planning modules, effectively decreasing the computational requirements of the BEV features. DiffusionDrive Liao et al. (2025) employs a truncated diffusion policy to improve the probabilistic depiction of trajectories, while MomAD Song et al. (2025) aims to enhance stability and maintain consistency throughout sequential planning decisions. Significantly, unlike other end-to-end planning techniques, NavigScene Peng et al. (2025) uses navigation data based on SD maps, achieving more accurate and reliable autonomous driving, thus underscoring the significance of navigation data. Consequently, we propose that the refinement of lane-level navigation through MAT could further enhance the capabilities of current autonomous driving methodologies.

C.2 MISSION COMPARISON

Fig. 7 compares four tasks: SD map navigation, map generation, map matching, and map association. Initially, navigation based on the standard definition (SD) map achieved a road-level path, which is now widely integrated into navigation apps. However, road-level navigation often does not provide precise directions for specific lanes and directions. As a result, the development of economical and quick lane-level navigation systems has emerged as a significant area of GIS. Next, lane-level online perception (OP) maps are generated using data from vehicle sensors by map generation. Unfortunately, such maps do not have a global topological framework, making it challenging to form a route from the start of the navigation to the endpoint. Subsequently, the task of map matching attempts to align GPS path with global SD or HD maps by following a path-to-map matching model. However, this model falls short when trying to match SD and OP maps due to disregarding the differences between these types. Finally, the map association task seeks to create linkages between SD and OP maps via a map-to-map association model. Using the SD map, the map association identifies a global route from the origin to the destination. Then this global route is meticulously translated into the OP map at the lane level through the association of the SD and OP maps, resulting in a comprehensive lane-level navigation path.

972
973
974 Table 7: Statistics of OMA train, val and test set.
975
976
977
978

Split	Train	Val	Test
HD map Range		$(\pm 15m, \pm 30m)$	
SD map Range		$(\pm 75m, \pm 75m)$	
Scene Segment	26111	5613	5573
Avg. lane per scene	81.3	75.8	310.34
Avg. lane path per scene	7.40	7.97	322.06
Avg. boundary per scene	3.48	3.31	9.65
Avg. length per lane	3.14m	3.19m	2.32m
Avg. length per boundary	44.81m	43.64m	32.17m
Avg. road per scene	15.1	10.8	10.8
Avg. length per road	38.2m	50.3m	50.3m
Avg. Connection per lane	2.0	2.1	2.9
Avg. Connection per road	2.0	1.8	1.8
Avg. Connection per boundary	2.0	2.0	2.0

989
990
991 D DATASET
992993
994 D.1 REVISED ANALYSIS.995
996 As detailed in Section 4.1, the dataset is partitioned into OMA train, val, and test set, with statistics
997 summarized in Tab. 7. The OMA training and val set comprises 26,111 training scenarios and 5,613
998 validation scenarios, totaling 31,724 samples, while the test set contains only 5,573 test scenarios
999 due to the exclusion of low-quality predictions. Both data sets share identical spatial coverage, with
1000 HD maps covering $(\pm 15m, \pm 30m)$ and SD maps extending to $(\pm 75m, \pm 75m)$. Notably, the SD
1001 map’s road density in OMA train and val set decreases from 15.1 roads/scene during training to
1002 10.8 in validation/test splits, suggesting potential domain shifts between training and evaluation
1003 environments.1004 Quantitative discrepancies between train, val and test set reveal systemic geometric and topological
1005 inconsistencies in predicted maps. The test set predicts an average of 310 lanes/scene, more than four
1006 times that of the Train and Val set (73.8), with significantly shorter mean lane lengths (2.32 m vs 3.16
1007 m in trainval), indicating both over-segmentation and false positives. This fragmentation is further
1008 amplified by test’s prediction of 322.06 lane paths/scene (vs. 7.69 in Train), where ground-truth lanes
1009 are frequently split into disconnected fragments. Boundaries exhibit similar degradation: test set
1010 detects 9.65 boundaries/scene (vs 3.40 in trainval) with reduced mean lengths (32.17m vs 44.23m),
1011 reflecting fragmented boundary detection. Meanwhile, validation data show slight degradation
1012 compared to training splits (e.g. 75.8 vs. 81.3 lanes/scene), highlighting inherent variability in
1013 real-world map quality.1014 Connectivity metrics expose deeper structural errors in test set predictions. The average lane con-
1015 nectivity in test set reaches 2.9, substantially higher than trainval’s 2.0/2.1, revealing widespread
1016 mislinking of spatially disjoint lanes. Similarly, the validation data show reduced road connectivity
1017 (1.8 vs. 2.0 in training), suggesting a domain bias toward simpler topologies in training scenarios. Se-
1018 mantic associations between roads and lanes also degrade significantly. Training roads are associated
1019 with 1,547 lanes on average, collapsing to 945 in validation splits, which implies degradation of the
hierarchical structure in complex scenarios.1020 These discrepancies have critical implications for benchmarking perception systems. The severe over-
1021 prediction and fragmentation in test set highlight the need for metrics penalizing false positives and
1022 disconnected paths (e.g., path-length-weighted scores). Furthermore, the mismatch between training
1023 and validation/test distributions (e.g., road count/length differences) necessitates domain adaptation
1024 strategies to ensure generalization. Finally, the collapse of semantic hierarchies in validation data
1025 suggests that end-to-end models may struggle to learn robust associations between roads and their
constituent elements without explicit structural constraints. Together, these findings underscore the

1026
1027
1028 Table 8: Refinement statistics for different scenes in OMA
1029
1030
1031

Scene name	Number of centerlines	Number of refinements	Refinement ratio
Boston	1,205,661	2,341	0.194%
Singapore	910,140	3,682	0.404%

1032
1033 importance of a connectivity-aware association method to avoid overestimating performance on
1034 fragmented or mislinked predictions.
1035

1036 D.2 RANGE OF SAMPLE 1037

1038 The model is trained on nuScenes using synchronized LiDAR and camera inputs with official
1039 configuration. For each sample, an SD map cropping measure $150\text{ m} \times 150\text{ m}$ centered around the ego
1040 vehicle preserves the adjacent topological context followed by the sensor setting of nuScenes Caesar
1041 et al. (2020a). For the OP map, the cropping measure of the OP map $30\text{ m} \times 60\text{ m}$ was centered
1042 around the vehicle of the ego, as referenced in Liao et al. (2023b); Liu et al. (2023a); Li et al. (2022).
1043

1044 D.3 PON SPLIT AND OP MAPS 1045

1046 Both Train, val and test set apply the pon split Roddick & Cipolla (2020) of the nuScenes dataset Cae-
1047 sar et al. (2020b), ensuring that there is no leakage between the training and validation datasets.
1048 For consistent lane prediction segmentation with OMA test set, we re-trained MapTRv2 Liao et al.
1049 (2023b) using the nuScenes dataset with the pon split. Drawing inspiration from the private protocol
1050 in MOT17/MOT20 Dendorfer et al. (2021), we suggest that future research evaluates the test set with
1051 an enhanced centerline prediction network, without relying on MapTRv2 as a baseline. Furthermore,
1052 we will establish an open evaluation protocol allowing submissions to include their own OP maps,
1053 thereby reducing reliance on a single model.
1054

1055 D.4 GEOGRAPHICAL GENERALIZATION 1056

1057 Regarding the geographical generalization of OMA, we elaborate from two aspects: First, OMA’s
1058 data originates from the nuScenes dataset, which is widely used in the autonomous driving field
1059 and contains data from two different countries, Singapore and Boston. Given that nuScenes has
1060 become a core dataset in the autonomous driving field for perception, mapping, and planning since
1061 its introduction, we believe that the nuScenes data itself possesses certain representativeness and
1062 generalizability. Second, we conducted cross-validation ablation experiments in the main text. The
1063 cross-experimental results of MAT on OMA demonstrate that MAT models trained on different
1064 geographical regions possess certain cross-regional generalization capabilities, reflecting the good
1065 geographical generalizability of the OMA dataset itself.
1066

1067 D.5 ANNOTATION WORKFLOW 1068

1069 The OMA data annotation process is divided into two primary stages: First, we performed coarse
1070 alignment based on the GPS coordinates from the SD map of OSM and the GT OP map of NuScenes.
1071 Second, we employed a skilled annotation team to correlate the SD map with the real perception map.
1072 The specialists were then engaged to evaluate and improve the annotations, ensuring their ultimate
1073 quality. Tab. 8 presents the total count and percentage of data changes in various regions in Phase 3.
1074 It is evident that the alteration rate for Singapore and Boston remained less than 1%, suggesting a
1075 generally high standard of data annotation quality.
1076

1077 D.6 HANDLING TRANSITION AMBIGUITY 1078

1079 Addressing the concern regarding transition lines, we acknowledge that assigning roads in such
1080 transition areas presents inherent ambiguity due to the heterogeneity between SD and OP Maps. To
1081 mitigate potential impacts on training stability and evaluation fairness, we implemented specific
1082 strategies: (1) *Standardized Annotation Protocol*: We established a unified rule where the assignment

1080 of a Lane Vector at an SD Link transition boundary is determined by the SD Link closest to the
 1081 Lane Vector’s start point, ensuring topological consistency. (2) *Tolerance in Evaluation Metrics*:
 1082 Our metric assesses the precision of path-level associations across 11 thresholds (50% to 95%).
 1083 Crucially, the 5% tolerance buffer included even at the strictest threshold is specifically designed to
 1084 accommodate inevitable semantic ambiguity at transition boundaries.
 1085

1086 D.7 FUTURE EXPANSION STRATEGY

1088 We acknowledge the importance of scaling the benchmark to encompass more diverse datasets (e.g.,
 1089 Argoverse). To minimize annotation costs during future expansion, we propose a “Human-in-the-
 1090 Loop” iterative annotation strategy. This approach is strongly supported by our data efficiency
 1091 experiments as shown in Tab. 6, which demonstrated that MAT can achieve competitive performance
 1092 with as little as 5% of the training data. The specific pipeline is as follows:

1. **Model-Assisted Pre-annotation:** We utilize the MAT model trained on the existing OMA
 1093 (nuScenes) dataset to perform zero-shot or few-shot inference on unlabeled new data. This
 1094 generates initial “draft” associations (e.g., SD-to-HD correspondences and topology) and
 1095 significantly reduces the cold-start problem.
2. **Lightweight Manual Correction:** Instead of labeling from scratch, annotators focus solely
 1096 on verifying and correcting the model’s high-confidence predictions. Given the model’s
 1097 high data efficiency, the pre-annotation quality improves rapidly even with a small set of
 1098 initial corrections.
3. **Closed-Loop Iteration:** The corrected data is immediately integrated into the training set
 1099 to fine-tune the model. The updated model is then used to pre-annotate the subsequent batch
 1100 of data, creating a positive feedback loop that progressively reduces manual workload.

1104 E METRIC

1107 E.1 DETAILS OF NR-PR

1109 In the main article, we present a narrative explanation of the Navigation Refinement P-R accompanied
 1110 by a schematic diagram. To elucidate the calculation of Navigation Refinement P-R more thoroughly,
 1111 we include the pseudo-code for computing Navigation Refinement P-R, as depicted in Alg. 1.

1112 Furthermore, the formula for $NR\text{-}P^{50:95}$, $NR\text{-}R^{50:95}$ and $NR\text{-}F1^{th}$, $NR\text{-}P^{50:95}$ is as follows:
 1113

$$1115 NR\text{-}P^{50:95} = \sum_{th \in T} NR\text{-}P^{th}, \quad NR\text{-}R^{50:95} = \sum_{th \in T} NR\text{-}R^{th} \quad (7)$$

$$1116 NR\text{-}F1^{th} = \frac{2NR\text{-}P^{th} \cdot NR\text{-}R^{th}}{NR\text{-}P^{th} + NR\text{-}R^{th}}, \quad NR\text{-}F1^{50:95} = \frac{2NR\text{-}P^{50:95} \cdot NR\text{-}R^{50:95}}{NR\text{-}P^{50:95} + NR\text{-}R^{50:95}}$$

1120 where th are the thresholds in association P-R as $[0.5 : 0.05 : 0.95]$ (10 thresholds). Additionally, in
 1121 the validation set, because $NR\text{-}R$ is always 1, we use the $NR\text{-}P$ value directly as $NR\text{-}F1$.
 1122

1123 E.2 CURVATURE-BASED WEIGHTED METRIC

1125 Although the length-interval stratification in the proposed NR-PR metric effectively mitigates bias
 1126 towards short paths, it does not explicitly account for the geometric complexity of roads. In real-world
 1127 datasets, straight roads are statistically more frequent than complex curved roads, potentially masking
 1128 model deficiencies in handling high-curvature topologies. To address this and evaluate the model’s
 1129 fairness across diverse road types, we introduce a **Curvature-based Weighted Metric**.

1130 **Complexity Measure.** Defining complexity via average angular changes can be highly sensitive to
 1131 point sampling rates. Therefore, we employ the **Discrete Fréchet Distance** between the actual path
 1132 and the straight line connecting its endpoints as a robust measure of geometric complexity. A higher
 1133 Fréchet distance indicates a greater deviation from a straight line, representing higher curvature or
 1134 irregularity.

1134 **Algorithm 1** Evaluate Navigation Refinement P-R

1135 1: **function** EVALMETRIC(pred_centerline, gt_centerline, threshold, acc_list)

1136 2: **Input:** pred_centerline, gt_centerline, threshold, acc_list

1137 3: **Step 1: Point Matching**

1138 4: EXTRACTPOINTS(pred_centerline)

1139 5: EXTRACTPOINTS(gt_centerline)

1140 6: POINTMATCH(pred_sample, gt_sample_point, threshold)

1141 7: **Step 2: Path Matching**

1142 8: INITIALIZECOUNTERS(TP, FP, FN, acc_list)

1143 9: **for** all point pairs (i, j) in matched points **do**

1144 10: Find pred_path and gt_path between points i, j

1145 11: PATHMATCH(pred_path, gt_path, threshold)

1146 12: **if** paths match **then**

1147 13: Check sequence consistency and accuracy

1148 14: **for** all $acc \in acc_list$ **do**

1149 15: Update TP/FP based on accuracy vs acc

1150 16: **end for**

1151 17: **end if**

1152 18: **end for**

1153 19: **Step 3: Count Unmatched Paths**

1154 20: **for** all unmatched gt path **do**

1155 21: **for** all $acc \in acc_list$ **do**

1156 22: $FN[acc][k] \leftarrow FN[acc][k] + 1$

1157 23: **end for**

1158 24: **end for**

1159 25: **Step 4: Calculate Precision and Recall**

1160 26: **Initialize:** $Precision \leftarrow \{\}, Recall \leftarrow \{\}$

1161 27: **for** all $acc \in acc_list$ **do**

1162 28: $denominator_p \leftarrow TP[acc] + FP[acc]$

1163 29: $denominator_r \leftarrow TP[acc] + FN[acc]$

1164 30: $Precision[acc] \leftarrow \begin{cases} TP[acc]/denominator_p & \text{if } denominator_p > 0 \\ 0 & \text{otherwise} \end{cases}$

1165 31: $Recall[acc] \leftarrow \begin{cases} TP[acc]/denominator_r & \text{if } denominator_r > 0 \\ 0 & \text{otherwise} \end{cases}$

1166 32: **end for**

1167 33: **Return:** TP, FP, FN, Precision, Recall

1168 34: **end function**

1170

1171

1172 **Stratification Strategy.** As shown in Fig. 8, similar to length-based approach, we adopt a stratified

1173 statistical method to handle the long-tailed distribution of road complexity. We identify the 95th

1174 percentile of complexity scores in the dataset as an upper bound and uniformly divide the range

1175 $[0, 95\text{th percentile}]$ into 10 bins. The final metric is computed by averaging the NR-F1 scores across

1176 these bins, ensuring that complex road geometries contribute equally to the final score, rather than

1177 being overshadowed by the dominant straight roads.

1178 **Quantitative Analysis.** We evaluated the MAT model on the OMA dataset using this weighted

1179 metric. Table 9 illustrates the impact of stratification on the overall score. As the number of

1180 bins increases—forcing the metric to weigh complex roads equally to straight ones—the overall

1181 performance drops significantly (e.g., from 81.6% to 60.3% on the Val set). This confirms that the

1182 dataset is dominated by simple geometries where the model performs exceptionally well, masking

1183 the challenges posed by complex topologies in unweighted metrics.

1184 Table 10 details the performance in specific complexity intervals. The results reveal a strong negative

1185 correlation between road complexity and association precision. On simple straight roads (Bin 1),

1186 the model achieves high precision (87.2% on Val). However, performance degrades drastically on

1187 high-complexity paths (dropping to 43.1% in Bin 10). This trend is even more pronounced in the

1188 Test set, indicating that while MAT is robust, geometric complexity remains a significant challenge

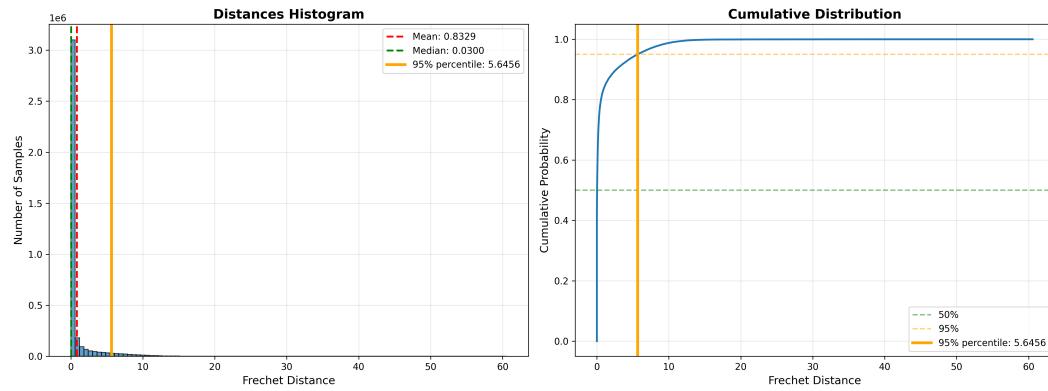


Figure 8: Statistics of Fréchet distances across all paths in the OMA dataset. (1) Histogram of Fréchet distances. A substantial number of paths exhibit Fréchet distances concentrated within the range of $[0, 1]$, revealing a pronounced long-tail distribution in the overall data. (2) Cumulative probability of Fréchet distances. Consistent with (1), the distribution displays a significant long-tail characteristic, underscoring the importance of block-wise statistics for the NR-PR metric.

Table 9: Impact of Complexity Stratification on NR-F1. The metric score decreases as we enforce equal weighting for complex road types (increasing N), revealing the dominance of simple roads in the unweighted score.

Complexity Bins (N)	1	2	5	10
Val Set	81.6	70.2	62.7	60.3
Test Set	51.8	38.7	33.5	32.5

for map association tasks. This analysis serves as a complement to the length-based metric, offering a more comprehensive view of the robustness of the model.

F METHOD

In this section, we delve deeper into the technical aspects of the model, including Path-aware attention, spatial attention, and the model’s post-processing.

F.1 PATH-AWARE ATTENTION

The particular design of path-aware attention (PA) can be seen in Fig 9 (a). The fundamental framework of PA is made up of four components: computing order and its inverse, reorganizing tokens, calculating attention, and inverting tokens.

Path-aware attention uses paths to determine the sequence of tokens. Initially, we define the network of roads or centerlines and then identify all complete paths from a starting point (with no incoming connections) to an endpoint (with no outgoing connections). We concatenate these paths to generate a path-based token sequence. During the reordering of path-aware attention, a token may appear across various paths simultaneously, requiring us to duplicate the token. Then, we compute attention by segregating tokens based on their paths, ensuring that interactions occur only among tokens within the same path. After attention calculation, the tokens are reversed to match the original input sequence. If multiple tokens exist within the path of a single original token, they are averaged.

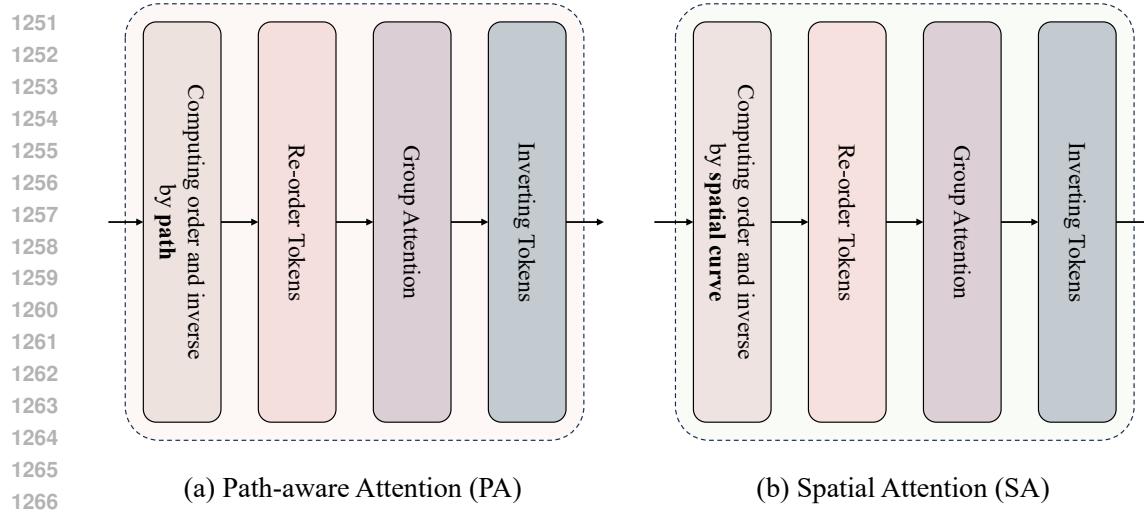
It should be highlighted that the model organizes paths internally to prevent overly lengthy paths. If a path surpasses the patch size, it is divided into several groups, where each group (aside from the final one) matches the patch size. PA’s worst-case time complexity is $O(kn^3)$, where k is the patch size and n is the number of tokens, if and only if $\frac{2}{3}$ centerline or road are either starting points or ending points. Clearly, this is highly improbable in real-world scenarios, and its average time complexity remains $O(k^2n)$, which is linear.

1242 Table 10: Performance per Complexity Bin ($N = 10$). Bin 1 represents straight roads, while Bin
 1243 10 represents high-curvature/irregular roads. The results indicate a performance degradation as
 1244 geometric complexity increases.

1245

Bin Index	1	2	3	4	5	6	7	8	9	10
Val Set	87.2	63.2	56.7	50.7	49.9	49.1	46.5	45.2	42.7	43.1
Test Set	57.9	50.4	32.0	21.2	16.8	14.7	10.7	10.6	10.4	11.6

1246



1265 Figure 9: Overview of Path-aware attention and spatial attention.

1266

F.2 SPATIAL ATTENTION

1271 Fig. 9 (b) provides a detailed description of the architecture of the SA model. Similarly to PA, the
 1272 fundamental structure of SA includes four main components: computing sorting and its reverse,
 1273 rearranging tokens according to the sorted order, calculating attention, and then reverse sorting the
 1274 tokens again.

1275 In order of SA, we apply a space filling curve $\varphi^{-1} : \mathbb{Z}^3 \rightarrow \mathbb{Z}$ to serialize the vectors in a 1D
 1276 sequence, preserving spatial locality. Four curves are used: Z-order, Transposed-Z, Hilbert, and
 1277 Transposed-Hilbert. To avoid bias toward specific curve types, we randomly select one curve per
 1278 training iteration. In addition, similar to PA, if there are multiple tokens that belong to a single token
 1279 after the coordinate calculation, we will average the multiple tokens in the sort and copy that token to
 1280 all the corresponding tokens in the reverse sort.

1281 During the group stage, tokens are partitioned based on the sorted sequence, each partition matching
 1282 the patch size. Following this, the model executes self-attention computations within these partitions.
 1283 Similar to PA, the self-attention (SA) algorithm has a time complexity of $O(kn)$, where k is the patch
 1284 size and n is the number of token, indicating that the complexity of SA remains linear.

1285

F.3 POST PROCESS

1287 In the main manuscript, we offer a narrative explanation of the post-processing. To enhance clarity,
 1288 we also present a mathematical formulation of the post-processing details. Let \mathcal{R} denote the road as
 1289 the vocabulary in the traditional beam search with size $|\mathcal{R}|$, and let $k = 4$ represent the width of the
 1290 beam. At each step t , the algorithm maintains a set B_t of candidates path k , each associated with a
 1291 score $s(h)$ defined as the sum of logarithmic conditional probabilities. The search begins by selecting
 1292 the token w^* with the maximum initial probability $P(w|x)$ given as input x , forming the singleton
 1293 initial set:

1294

1295

$$B_0 = \text{Top}_1(\mathcal{R}, \log P(w|x)), \quad (8)$$

1296 Table 11: Ablation study of patch size of PA and SA. Table 12: Ablation study of group method of
 1297 PA.

Patch Size	64	256	1024	2048	∞
NR-F1 ^{50:95}	77.8	78.4	78.7	78.5	78.4
Latency/ms	77	68	70	72	72

Group Method	N.G	Category	Path
NR-F1 ^{50:95}	78.0	78.2	78.7
Latency/ms	75	73	70

1303 Table 13: Ablation study of Input size of SD Map

Input size of SD Map	NR-F1 ^{50:95}	Latency/ms
60 × 60m	78.6	67
90 × 90m	78.6	69
120 × 120m	78.7	69
150 × 150m	78.7	70

1311
 1312
 1313 which simplifies to:

$$1314 \quad B_0 = \{[w^*]\}, \text{ where } \log P(w^*|x) = \max_{w \in \mathcal{R}} \log P(w|x). \quad (9)$$

1316 This initialization bypasses conventional fixed start tokens and prioritizes high-probability seeds.

1317 In iteration $t \geq 1$, each sequence $h \in B_{t-1}$ generates $2|\mathcal{R}|$ candidates by appending a token $w \in \mathcal{R}$ to the left ($w \cdot h$) or right ($h \cdot w$) of h , forming the expanded candidate set:

$$1320 \quad \mathcal{C}_t = \{w \cdot h \mid h \in B_{t-1}, w \in \mathcal{R}\} \cup \{h \cdot w \mid h \in B_{t-1}, w \in \mathcal{R}\}. \quad (10)$$

1321 h extension updates the sequence score using direction-specific conditional probabilities:

$$1323 \quad s(h') = \begin{cases} s(h) + \log P(w|h, \text{left}, x), & \text{if } h' = w \cdot h \\ s(h) + \log P(w|h, \text{right}, x), & \text{if } h' = h \cdot w \end{cases}. \quad (11)$$

1326 The top- k candidates from \mathcal{C}_t are retained to form B_t :

$$1327 \quad B_t = \text{Top}_k(\mathcal{C}_t, s(\cdot)), \quad (12)$$

1329 i.e.,

$$1330 \quad B_t = \{h'_1, h'_2, \dots, h'_k\}, \text{ with } s(h'_1) \geq s(h'_2) \geq \dots \geq s(h'_k). \quad (13)$$

1332 The process ends at a predefined maximum length T or when all sequences emit an end-of-sequence
 1333 token, with the final output \hat{h} selected as:

$$1334 \quad \hat{h} = \arg \max_{h \in \bigcup_{t=0}^T B_t} s(h). \quad (14)$$

1337 G ABLATION STUDY

1339 For the hyperparameters we proposed, additional ablation studies were executed to exhibit their
 1340 stability.

1342 G.1 HYPER PARAMETER

1344 **Patch Size of PA and SA.** Tab. 11 shows that precision plateaus at patch size 256 (NR-F1^{50:95} =
 1345 78.4%) but increases slightly at 1024 (+0.3%) with a latency trade-off (+2 ms).

1346 **Group Method of PA.** For PA grouping (Tab. 12), path-based grouping surpasses non-grouping
 1347 (+0.7%) and category-based baselines (+0.5%), likely due to reduced cross-path interference.

1349 **Input size of SD map.** Table 13 presents the results of the adjustment of the SD map input size. The
 findings reveal that alterations in the SD input size exert minimal influence on NR-F1 (≤ 0.1).

1350

Table 14: Ablation study of δ

1351

1352

δ	NR-F1 ^{50:95}	Latency
1m	74.6	65
0.5m	78.1	68
0.1m	78.7	70
0.01m	78.7	79

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

Table 16: Ablation study of space curve in SA.

Spatial Curve	NR-F1 ^{50:95}	Latency/ms
Z	77.6	70
Z + TZ	78.1	70
H + TH	78.2	70
Z + TZ + H + TH	78.7	70

Sampling size of SA. Table 14 shows the findings for the spatial sample size δ . When δ is set at 1m, there is a drop in model accuracy, which may be attributed to the challenge of differentiating nearby lane lines with such a substantial sampling size. When δ is below 0.5m, the accuracy of the model is decreased (< 0.6). Decreasing the interval further to 0.01 does not improve the accuracy of the model.

Hyperparameter of loss function. As illustrated in Table 15, the ablation studies for β_{ce} and β_{ctc} indicate that high beta values make CTC the main loss, which reduces the precision of the model. In contrast, lower beta values allow CE loss to predominate, leading to stabilized NR-F1 (< 0.3).

Space Curve of SA. Table 16 displays the result of ablation studies focusing on spatial curves in SA. The findings suggest that adding more diverse spatial curves improves the model’s performance. It is crucial to emphasize that including multiple spatial curve types does not affect model latency, given that each layer utilizes only a single spatial curve. Thus, the variety of spatial curve types does not increase the frequency of spatial curve sorting.

Beam width of post-process. Table 17 shows the results of the ablation test regarding the hyperparameter of beam width k during post-processing. The study reveals that setting k to 4 results in the best trade-off between accuracy and processing time.

Length intervals of NR P-R. To further investigate the influence of the path length distribution on the NR-F1 metric, we performed a stratified evaluation based on length intervals. The motivation for introducing length intervals stems from our observation that within the dataset, short-distance paths occur far more frequently than long-distance ones, and their corresponding F1 scores tend to be higher. Without appropriate handling, this imbalance could bias the overall NR-F1 score toward performance on short-distance paths. Inspired by the size-based categorization strategy adopted in the MS COCO benchmark Lin et al. (2014), we categorize paths into discrete length intervals, compute the F1 score within each interval, and subsequently aggregate the results into a unified NR-F1 measure. This procedure ensures that each length range contributes proportionally to the final metric, thus mitigating the dominance of short-distance paths in the evaluation. As shown in Tab. 18, omitting this stratification (that is, computing the metric on all paths without length-based grouping) produces a substantially inflated overall score. In contrast, increasing the number of length intervals causes the aggregated metric to gradually decrease and eventually converge to a stable value. These findings suggest that the length interval-based computation allows NR-F1 to provide a more balanced and fair evaluation on both the short- and long-distance paths, resulting in a more comprehensive evaluation of the model’s ability to capture relational associations.

Distance Threshold of NR P-R. Tab. 19 highlights the effect of various distance thresholds on the final metrics in NR P-R Step 1. These thresholds originate from those used in Reachability P-R. It is

Table 15: Ablation study of β_{ce} and β_{ctc} in loss function

$\beta_{ce} : \beta_{ctc}$	NR-F1 ^{50:95}	Latency
1:1	68.4	70
1:0.1	78.5	70
1:0.01	78.7	70
1:0.001	78.4	70
1:0	78.4	70

Table 17: Ablation study of beam width k in post process.

Beam Width	NR-F1 ^{50:95}	Latency/ms
1	78.2	69
2	78.6	70
4	78.7	70
8	78.7	78

1404
1405
1406 Table 18: Ablation study of length intervals of NR P-R.
1407
1408
1409

Number of length intervals	1	2	3	6	10	15
NR-F1 ^{50:95} / Val	81.8	79.0	78.5	78.2	78.7	78.7
NR-F1 ^{50:95} / Test	51.6	47.6	47.2	45.3	44.9	45.0

1410
1411 Table 19: Ablation study of distance threshold of NR P-R.
1412
1413

Distance threshold	0.5m	1.0m	1.5m	2.0m
NR-F1 ^{50:95} / Val	78.7	78.7	78.7	78.7
NR-F1 ^{50:95} / Test	6.6	45.0	68.2	72.4

1418 important to note that Val uses the ground-truth OP map for predictions, thus all roads are deemed
1419 true roads, while the distance threshold applies solely to the test set. Ablation experiments reveal that
1420 distance thresholds significantly influence the accuracy outcome of NR-F1. With a strict threshold of
1421 0.5m, most of the lane lines on the OP map are regarded false detections, as they do not meet the
1422 requirement. In contrast, a more lenient threshold of 2.0 m allows most lane lines to be perceived
1423 as correct, bringing the NR-F1 to a value comparable to that of the Val set. This also indicates that
1424 there is no domain difference between the Test and Val sets; discrepancies in results are due to the
1425 threshold settings. However, considering that the accuracy of the lane regression is also pivotal, we
1426 opted for a balanced threshold of 1.0m to judge whether the OP map correctly represents a lane.

1427
1428

G.2 ROBUSTNESS TO IMPERFECT SD MAPS

1429 In real-world applications, Standard Definition (SD) maps often suffer from imperfections caused
1430 by GPS localization errors, sensor noise, or map construction delays. To quantitatively evaluate the
1431 robustness of MAT and the proposed metrics against such imperfect inputs, we conducted additional
1432 ablation studies on the OMA Validation set.

1433 We categorized SD map imperfections into three distinct types and injected varying levels of noise
1434 (Low, Medium, High) into the input SD maps. The noise generation protocols are defined as follows:
1435

- 1436 • **Global Shift:** Simulates GPS localization errors by applying a uniform distribution offset
1437 relative to the global map range. We tested noise ratios of 10%, 20%, and 50%. Crucially,
1438 for each SD map sample, the same random offset vector is applied to all road vectors,
1439 preserving the internal shape but shifting the absolute position.
- 1440 • **Element Noise:** Simulates geometric construction errors by applying random perturbations
1441 relative to the global map range. We tested noise ratios of 5%, 10%, and 20%. Unlike
1442 Global Shift, the noise here is generated independently for each vector within the SD map,
1443 resulting in geometric jitter and shape distortion.
- 1444 • **Element Absence:** Simulates map construction incompleteness by randomly masking out
1445 road vectors. We evaluated omission ratios of 10%, 20%, and 30%, where vectors are
1446 randomly discarded from the SD map input.

1448 The results, measured by the NR-F1^{50:95} score, are summarized in Tab. 20. The results demonstrate
1449 the resilience of the proposed method under varying noise conditions:
1450

- 1451 • **Robustness to Shift:** Even under a high global shift of 50%, the model maintains reasonable
1452 performance (74.6 / 41.4). This indicates that our **Path-Aware Attention (PA)** successfully
1453 captures the relative topological structure of the road network, reducing dependency on
1454 absolute coordinate alignment.
- 1455 • **Robustness to Element Noise:** The performance remains highly stable, dropping only
1456 marginally from 78.3 / 44.8 to 77.5 / 44.1 despite a 20% geometric jitter. This suggests that
1457 the **Spatial Attention (SA)** mechanism effectively aggregates global contextual information,
1458 thereby mitigating the impact of local geometric inaccuracies.

1458 Table 20: Ablation study of robustness to imperfect SD Maps on OMA Val / Test set with varying
 1459 noise intensities. The metric reported is NR-F1^{50:95}.

Noise Type	Low	Medium	High
Global Shift (10%, 20%, 50%)	77.5 / 43.3	76.0 / 42.9	74.6 / 41.4
Element Noise (5%, 10%, 20%)	78.3 / 44.8	77.5 / 44.3	77.5 / 44.1
Element Absence (10%, 12%, 30%)	77.5 / 44.0	76.4 / 43.7	75.5 / 43.1

1466 Table 21: Zero-shot robustness evaluation on nuScenes-C under different environmental corruptions.

Condition	Easy	Mid	Hard
Rain	44.7	42.1	40.6
Fog	44.2	43.4	42.1
Dark	44.0	43.8	41.6

- **Robustness to Missing Elements:** The model exhibits strong adaptability to missing data, maintaining a score of 75.5 / 43.1 even when 30% of the SD map vectors are absent. The Transformer architecture effectively infers associations based on the remaining context and the overall graph topology, ensuring that the association process does not fail catastrophically when individual segments are missing.

1480 These experiments confirm that MAT maintains strong robustness when dealing with imperfect SD
 1481 maps, effectively handling geometric misalignment, shape distortion, and topological incompleteness
 1482 inherent in real-world navigation tasks.

1484 G.3 ROBUSTNESS UNDER ENVIRONMENTAL NOISE AND OCCLUSION

1486 To evaluate the robustness of MAT against sensor failures caused by environmental factors and severe
 1487 occlusions, we conducted extended experiments focusing on two aspects: quantitative evaluation on
 1488 corrupted inputs and qualitative analysis of occlusion scenarios.

1489 **Robustness against Environmental Noise.** We utilized *nuScenes-C* Xie et al. (2025b), a benchmark
 1490 designed to assess robustness against input corruptions, to simulate sensor degradation. Specifically,
 1491 we used MapTRv2 Liao et al. (2023b) to perform inference on the *nuScenes-C* validation set with
 1492 pon split, generating OP maps under three specific weather conditions: **Rain**, **Fog** and **Dark** (Night).
 1493 These conditions were categorized into three difficulty levels: Easy, medium, and Hard followed by
 1494 the nuScenes-C setting.

1495 Crucially, we adopted a **zero-shot inference** setting: the MAT model used for association was trained
 1496 solely on clean data and was **not retrained** on these corrupted samples. This setting serves as a
 1497 rigorous stress test for the generalizability of the model.

1498 Table 21 reports the quantitative results with NR-F1^{50:95} score. As observed, MAT maintains robust
 1499 performance across all conditions. Even in “hard” settings, where visual features are significantly
 1500 compromised, the performance drop is minimal. For example, in Fog scenarios, the metric only
 1501 decreases from 44.2 (Easy) to 42.1 (Hard). This stability demonstrates that, by taking advantage of
 1502 the topological priors of SD maps, MAT effectively mitigates the impact of sensor noise.

1504 **Visualization of Environmental Degradation.** To intuitively display the challenge, Fig. 10 illustrates
 1505 the synthesized degradation samples of the OP maps at different difficulty levels. Fig. 11 and Fig. 12
 1506 further compare the inference results under the “Hard” setting. Despite the severe noise in the input
 1507 OP maps (e.g. missing boundaries or phantom lanes due to fog/rain), MAT successfully recovers the
 1508 correct lane topology by aligning it with the SD map.

1509 **Analysis of Severe Occlusion.** While the explicit quantitative evaluation of occlusion is constrained
 1510 by the lack of ground-truth annotations, we analyze it physically. Vehicle occlusion typically manifests
 1511 itself in OP maps as lane break or missing segments. Since MAT is explicitly designed to model
 global topology, it is inherently robust to such defects.

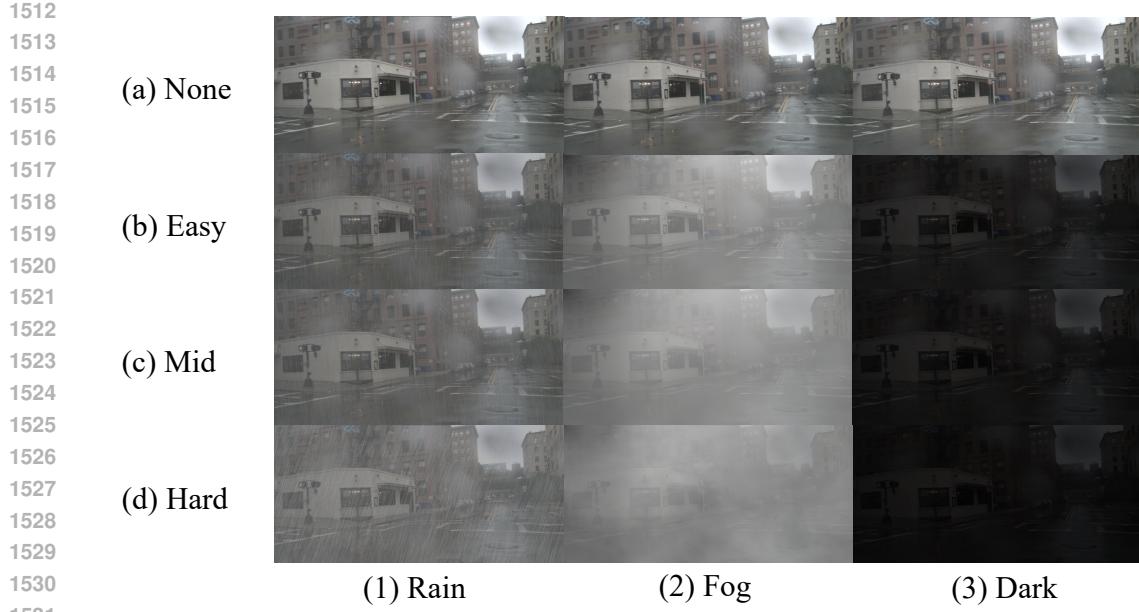


Figure 10: **Visualization of OP map degradation samples from nuScenes-C.** The columns represent different environmental conditions (Rain, Fog, Dark), and the rows represent increasing difficulty levels (Easy, Mid, Hard). These noisy inputs simulate realistic sensor failures.

Fig. 12 presents a case study of **severe vehicle occlusion**. As shown in the visualization, MapTRv2 manages to reconstruct a coherent global road structure despite significant visual blockage by vehicles. We attribute this capability to the temporal modeling design of the upstream perception model (MapTRv2), where single-frame occlusion does not disrupt the map modeling derived from the overall temporal sequence. Taking advantage of the inherent robustness of MapTRv2 against occlusion, our downstream MAT model is consequently largely unaffected by dynamic vehicle occlusion, ensuring stable association performance.

H VISUALIZATION

This section presents the visualization of our model, featuring path-aware attention (PA) and spatial attention (SA) alongside the model’s results. Furthermore, we examine the failed case with an analysis of our model.

H.1 ATTENTION MAP

The upper portion of Fig.13 presents visualizations of Spatial Attention (SA) maps at different stages of the model. As revealed by the analysis, SA provides extensive receptive fields in the early stages, enabling tokens to capture global contextual information. Specifically, during Stage 1 and Stage 2, the SA attention distributions exhibit highly dispersed patterns, allowing each query token to uniformly attend to global regions across the input space. In later stages, the functional role of SA transitions to facilitating cross-category token interactions. For example, in Stages 3-5, distinct attention patterns emerge where tokens primarily interact with their semantically corresponding road elements. Notably, this interaction is not strictly confined to the road tokens directly associated with the centerline token - significant attention weights also develop between the centerline token and adjacent road segments. As exemplified in Stage 4, the tokens establish prominent attention links with multiple road tokens along the same path. We posit that this expanded interaction mechanism constitutes a critical component for precise centerline localization. The propagation of attention observed in later stages effectively enables the maintenance of geometric coherence between spatially distributed road elements while preserving discriminative semantic information through long-range dependencies.

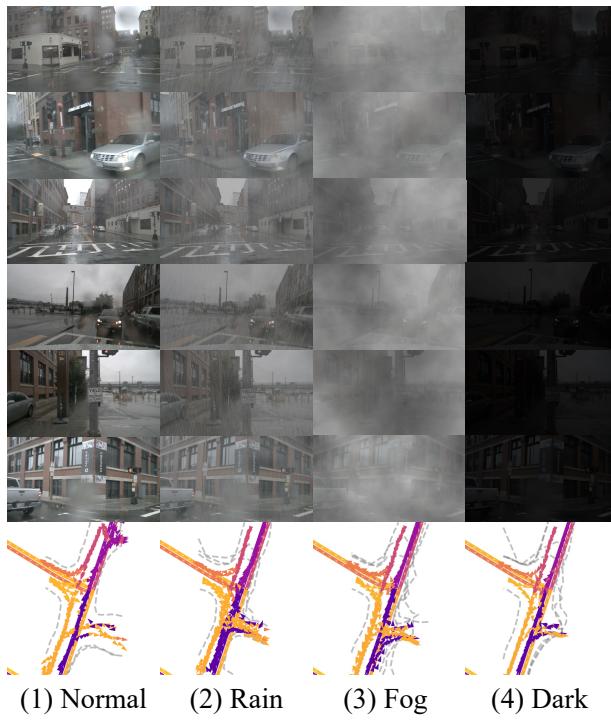


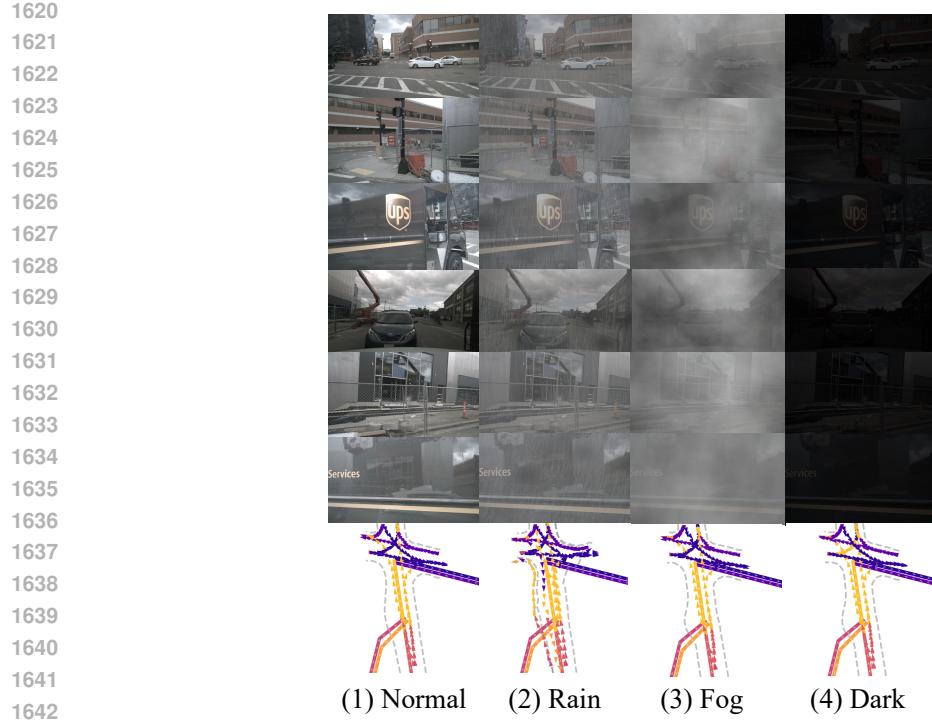
Figure 11: **Qualitative results under “Hard” environmental conditions.** The figure compares the ground truth (GT) with MAT’s predictions under severe Rain, Fog, and Dark conditions. The model demonstrates strong zero-shot robustness, effectively associating noisy OP lanes with SD paths.

In contrast, the lower part of Fig.13 visualizes the Path-aware Attention (PA) maps at different stages of the model. The visualization reveals that PA primarily focuses on neighboring tokens adjacent to the target path tokens, effectively serving as a local information extractor. Experimental results demonstrate that this localized information extraction capability plays a pivotal role in the model performance, exhibiting a marked contrast with the global perception mechanism of SA. We posit that SA specializes in capturing global contextual patterns while PA emphasizes localized feature extraction. This dual-attention paradigm establishes a synergistic interplay between global and local perception, achieving an optimal balance between comprehensive understanding and fine-grained detail processing, thereby substantially enhancing the model’s overall effectiveness.

H.2 RESULT COMPARE

Fig.14 illustrates comparisons of model ground truth on KNN, HMM, MAT and MAT w/ postprocess. The top rows of (d) and (e) exhibit our post-processing module’s effectiveness. MAT predictions initially display incorrect topological connections where roads are mistakenly linked (marked with red circles). Our postprocessing, which utilizes topology-aware beam searching, rectifies this by eliminating non-sequential transitions and reconstructing precise topological paths. The second row demonstrates MAT’s superior handling of complex topologies. Although HMM targets sequential path associations, its single path paradigm often underperforms in complex topologies with intersections. In contrast, our model uses spatial attention to grasp global information and cross-path associations, facilitating adaptive learning of complex topological patterns for accurate connectivity inference. The third row showcases our model’s improved ability to localize associations. Using path-aware attention, the model emphasizes detailed extraction of local features along paths. This targeted local perception ensures precise associations at challenging points, such as junctions, where HMM is typically short due to limited contextual understanding.

Fig. 15 illustrates a comparison of ground truth results for KNN, HMM, MAT, and MAT with post-processing. Significantly, the visualization demonstrates that our model excels in map association in



1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

Figure 12: **Qualitative results under “Hard” environmental conditions and severe occlusion.** Similar to the setup in Fig. 10, this figure displays the inference results under Rain, Fog, and Dark scenarios, but specifically highlights a case of severe vehicle occlusion. As observed, the model maintains robust lane reconstruction capabilities even when the visual information is corrupted by extreme weather or significantly obstructed by other vehicles.

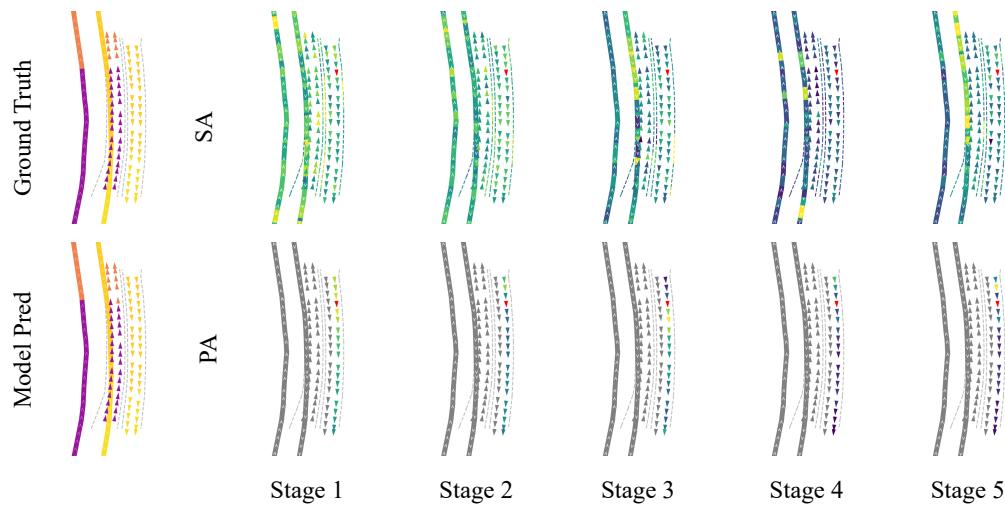


Figure 13: Visualization of attention map of Path-aware attention and spatial attention. SA means Spatial Attention. PA mean Path-aware Attention. The red triangle represents the token corresponding to the current attention map.

noisy scenarios with inaccurate centerline predictions, surpassing KNN and HMM by integrating the complementary benefits of global association (SA) and local detail refinement (PA).

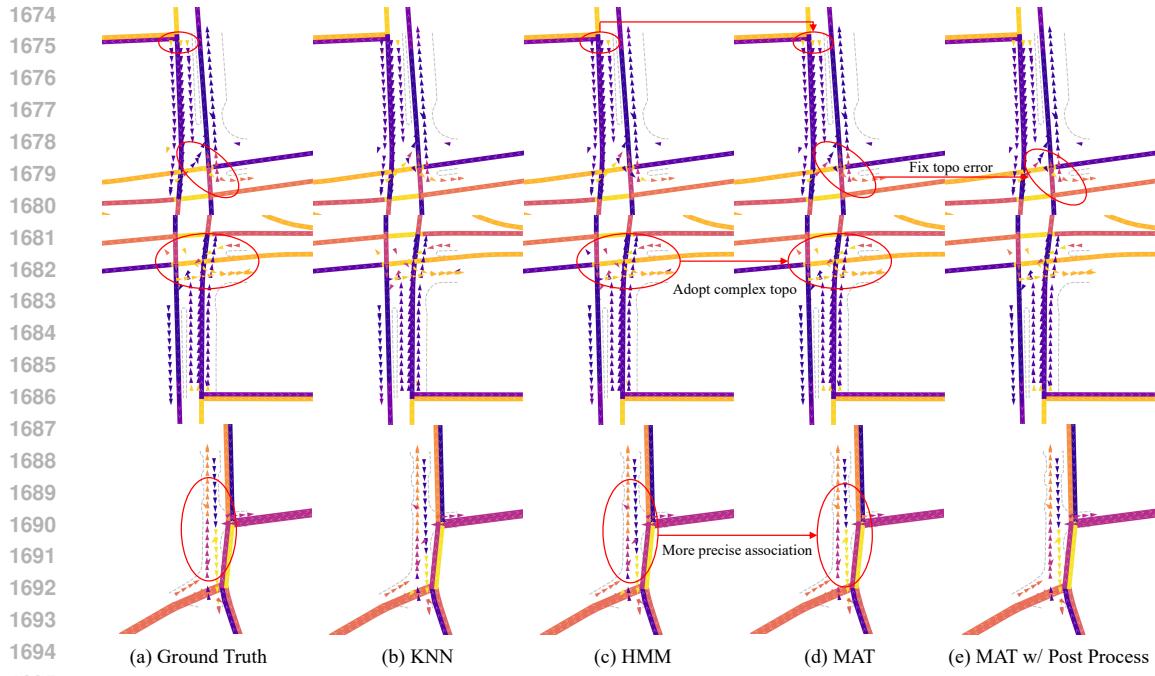


Figure 14: Visualization of result in OMA Val set.

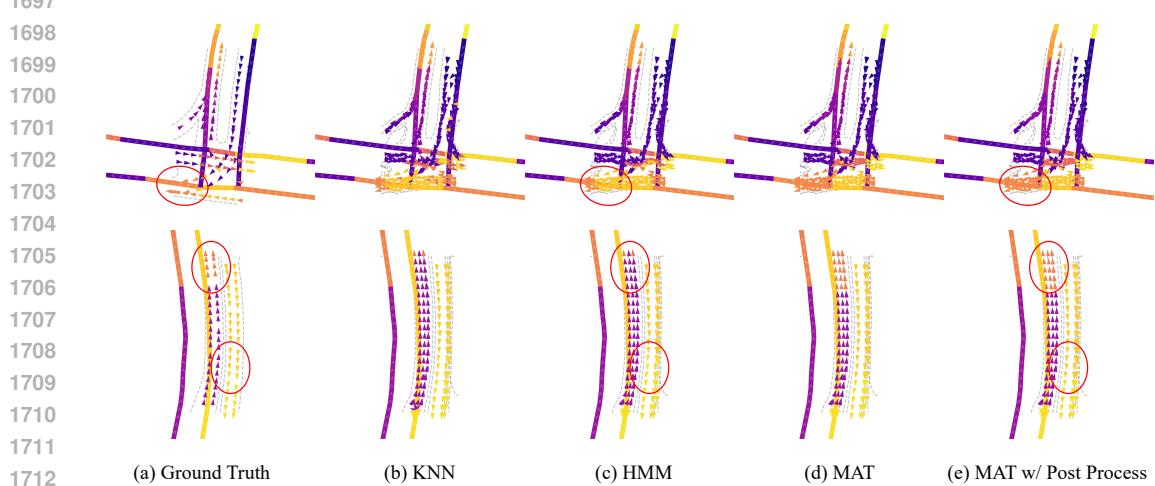


Figure 15: Visualization of result in OMA Test set.

H.3 FAILED CASES

1718 Fig. 16 illustrates the failed cases of the MAT. Our study reveals that the key challenge is the
1719 localization errors associated with spatial misalignment between the predicted paths and the actual
1720 labels. This discrepancy significantly affects the accuracy of the association, particularly at critical
1721 junctures where complex path interactions create ambiguous topological patterns. Although all
1722 baseline methods exhibit substantial association errors under these difficult conditions, our model
1723 achieves notable error reduction due to its dual-attention framework. However, discrepancies between
1724 our predictions and the ground truth remain, indicating potential for further enhancement. We
1725 propose that improving the path-aware attention (PA) mechanism by incorporating local operators
1726 such as convolutional kernels could be advantageous. This hybrid approach would preserve model
1727 efficiency while allowing for more precise spatial-temporal feature extraction at path intersections,
thus improving local association accuracy without compromising inference speed.

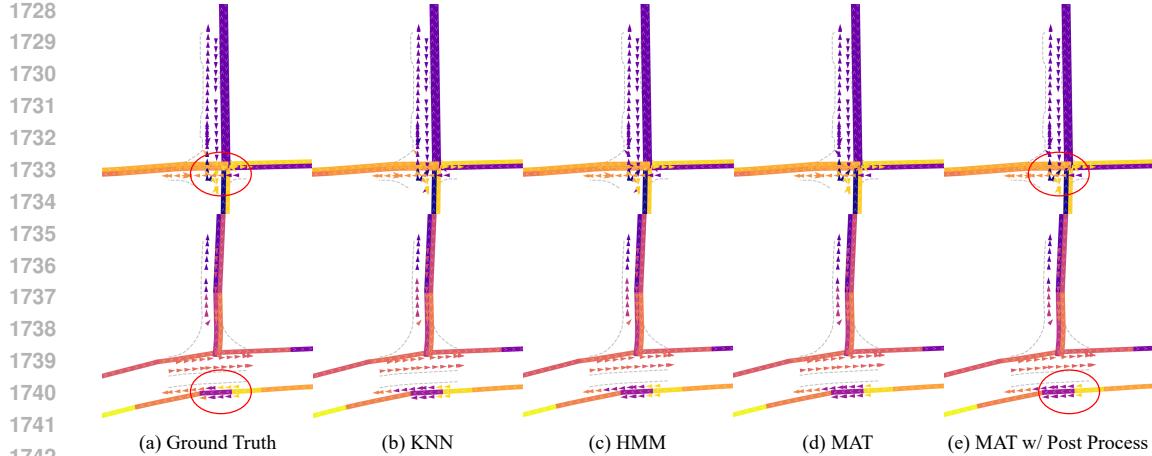


Figure 16: Visualization of failed cases.

H.4 GENERALIZATION ON DIFFERENT OP GENERATORS

To further validate the robustness of MAT against varying noise patterns inherent to different map generation paradigms, we visualize the association results on the OMA Test set using three distinct baselines: MapTR (Liao et al., 2023a), MapTRv2 (Liao et al., 2023b), and SeqGrowGraph (SSG) (Xie et al., 2025a).

As illustrated in Figure 17, each generator exhibits unique geometric and topological characteristics:

- **Row 1 (MapTR):** Often produces fragmented centerlines with localized geometric jitter, creating a challenge to ensure topological continuity.
- **Row 2 (MapTRv2):** Offers improved geometric stability, but still suffers from occasional detection gaps and instance discontinuities.
- **Row 3 (SeqGrowGraph):** Generates highly connected graphs via expansion, which reduces fragmentation, but may introduce erroneous topological links (over-connection) in complex intersections.

Despite these substantial domain gaps, MAT consistently establishes accurate associations (indicated by the consistent coloring of lane instances that match the SD map topology) across all three inputs. This qualitative evidence reinforces the quantitative results in Tab. 1, demonstrating that our proposed Path-Aware and Spatial Attention mechanisms effectively generalize across disparate upstream perception generators without requiring generator-specific fine-tuning.

I IMPLEMENT DETAILS

I.1 MODEL SETTINGS

Tab. 22 summarizes the architectural configurations of our proposed MAT variants (MAT-T and MAT-L). All variants adopt identical channel dimensions, attention head counts, spatial curve orders and hybrid attention mechanisms combining spatial attention (SA) and path-aware attention (PA). In particular, parameters such as patch sizes, MLP ratios (matching spatial curve orders), and stochastic depth rates are uniformly inherited across architectures, reflecting ablation study results that optimized these values for balanced accuracy-latency trade-offs. A distinctive design choice lies in the shuffling strategy, where MAT-T/L progressively refine the shuffle order to enhance token mixing in spatial attention, aligning with their increasing computational budgets. This structured configuration hierarchy enables systematic evaluation of model capacity versus efficiency, as validated by the ascending La / ms metrics (34 → 70) corresponding to deeper transformer layers.

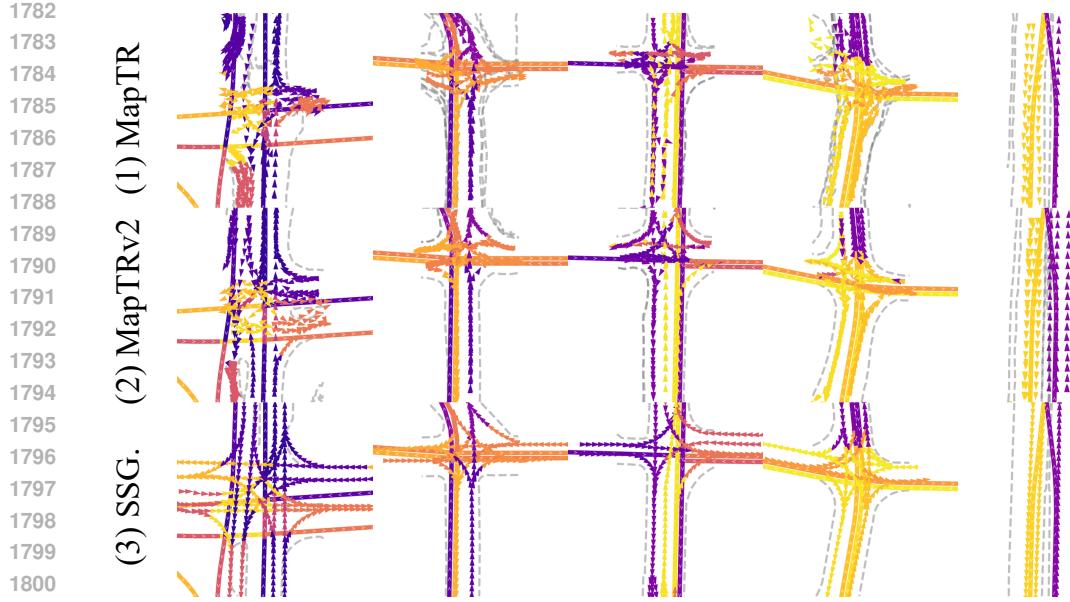


Figure 17: Visualization of MAT association results across different Online Perception (OP) map generators on the OMA Test set. Rows correspond to (1) MapTR, (2) MapTRv2, and (3) Seq-GrowGraph (SSG). The columns display diverse scenarios including intersections and straight roads. Consistent coloring between lane segments indicates that MAT successfully associates fragmented or noisy OP elements with the correct semantic road instances from the SD map, demonstrating strong cross-generator robustness.

Table 22: Model settings.

Parameter	MAT-T	MAT-L
Blocks	[2, 2, 2, 2, 2]	[4, 4, 4, 12, 4]
Attention Head		[4, 4, 8, 8, 8]
MLP Ratio		[4, 4, 4, 4, 4]
Drop Path		[0.3, 0.3, 0.3, 0.3, 0.3]
Channels		[96, 192, 384, 768, 1536]
Path Size		[1024, 1024, 1024, 1024, 1024]
Attention Order		["Spatial Attention", "Path-aware Attention"]
Spatial Curve		["z", "z-trans", "hilbert", "hilbert-trans"]
Shuffle		[Shuffle Order, Shuffle Order, Shuffle Order, Shuffle Order, Shuffle Order]
Latency/ms	34	70

Table 23: Train Configuration and Data augmentations.

Training Configuration			
optimizer	AdamW	batch size	128
scheduler	Cosine	weight decay	5e-3
learning rate	1e-4	epochs	50
criteria	CrossEntropy, CTC Loss	warmup epochs	2
Data Augmentation			
random rotate	axis: z, angle: [-1, 1], p: 0.5	random scale	scale: [0.9, 1.1]
random flip	p: 0.5	random jitter	sigma: 0.005, clip: 0.02
grid sampling	grid size: [0.1, 0.1, π/16]		

1836
1837

I.2 TRAIN CONFIGURATION

1838
1839
1840
1841
1842

The details of the implementation are summarized in Tab. 23. Our training protocol uses the AdamW optimizer with a base learning rate of 1×10^{-4} and a cosine learning rate decay, operating in batch size of 128. The weight decay regularization is set to 5×10^{-3} . The training process spans 50 epochs with a 2-epoch warm-up phase for learning rate initialization. Model optimization combines CrossEntropy loss for classification tasks and CTC loss for sequence alignment objectives.

1843
1844

I.3 DATA AUGMENTATIONS

1845
1846
1847
1848
1849
1850
1851

For data augmentation as shown in Tab. 23, we implement a series of randomized transformations that include axis-aligned rotation around the z-axis with $\pm 1^\circ$ angular variation at a 50% application probability, isotropic scaling within the range $[0.9, 1.1]$, random flipping with equal probability 50%, point cloud jittering characterized by $\sigma = 0.005$ and a clip limit of 0.02, and grid sampling with spatial discretization parameters set to $[0.1, 0.1, \pi/16]$. These enhancement strategies were systematically validated through ablation studies to optimize the balance between model accuracy and computational efficiency while ensuring robustness to input variations.

1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889