Scalable Conflation for Maps Data Replay between Heterogeneous Geospatial Data Sources

Shaishav M. Maisuria
smaisur@amazon.com
Last Mile, Amazon
Bellevue, WA, USA

Vishal Kumar
vishku@amazon.com
Last Mile, Amazon
Bellevue, WA, USA

Ravi Garg
ravigarg@amazon.com
Last Mile, Amazon
Bellevue, WA, USA

Hesham Eraqi
heraqi@amazon.com
Last Mile, Amazon
Bellevue, WA, USA

Abstract

Last-mile logistics efficiency largely depends on map data being accurate, fresh, and having high geospatial coverage. Operators must preserve these qualities when migrating between heterogeneous map data providers. Existing work in the literature addresses road-geometry conflation; however, the conflation of road network attributes and their transition relations remains under-studied, especially at a production-level scale. This leaves significant gaps in the map data when operators switch map providers. In this work, we present a scalable map conflation system to close those gaps and replay accumulated map edits onto a refreshed base map. The proposed system has four components. Change-driven partitioning groups in-scope edits into independently conflatable clusters; connected components of a shared-node graph subdivided by size into a hierarchy of uniform spatial grid cells (S2 cells), so the same pipeline drives both continental-level bootstrapping and small-scale incremental refreshing use cases. A logistic-regression map-matching model (with road-segment normalization) conflates both road segment and intersection-level attributes. A random-forest model formulates turn-restriction conflation as intersection matching followed by road matching and scoring candidate pairs with geometric features. An orchestration layer ties these core components into a scalable, unified production pipeline. On the OpenStreetMap and Overture Maps benchmark, the proposed system successfully conflated ∼94.9% of intersection attributes, ∼94.5% of road-segment attributes, and ∼60% of turn restrictions at 99% precision; the intersection-matching classifier reached ∼92% accuracy.

CCS Concepts

• **Information systems → Geographic information systems**; • **Computing methodologies** → *Spatial data processing*.

Keywords

map conflation, cross-provider migration, road networks, fingerprinting, continental-scale GIS, OpenStreetMap, Overture Maps

SIGSPATIAL '26, Riverside, CA, USA

ACM ISBN 979-8-4007-2950-8/2026/11
https://doi.org/10.1145/3841645.3843272

ACM Reference Format:
Shaishav M. Maisuria, Vishal Kumar, Ravi Garg, and Hesham Eraqi. 2026. Scalable Conflation for Maps Data Replay between Heterogeneous Geospatial Data Sources. In *The 34th ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '26), November 03–06, 2026, Riverside, CA, USA*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3841645.3843272

1 Introduction

Global-scale road-network applications benefit from data assembled from multiple map providers. A mapping deployment may blend open-source OpenStreetMap (OSM) and the Overture Maps Foundation transportation layer [13] in some regions while leveraging region-specific commercial providers in others. Map providers snapshot the world differently when creating maps: the underlying satellite imagery used to digitize roads can originate from varied sources, leading to misaligned geometric representations across maps. Similarly, distinct digitization conventions lead to different road modeling choices, different topology split points at junctions, divergent attribute schemas, and refresh cadences that rarely align.

While map providers continue to improve their data, users deploying routing applications also benefit by investing in their own map quality improvements and enrichments. For example, custom edits such as corrected turn restrictions, refined speed limits, or added road attributes that reflect ground truth not yet captured by any provider significantly improve map quality for downstream applications. At the same time, providers periodically release improved base maps with better coverage, fresher imagery, or higher geometric fidelity, motivating users to switch to a newer base map. When such a switch occurs, the custom enrichments accumulated on the prior base map must also be migrated to the new one in order to preserve maximum quality in the merged map. In this paper, we refer to this problem of migrating desired map data between different map providers as the *Maps Data Replay Conflation* problem.

Due to the differing representations across map providers, maps data replay conflation is not a trivial problem. Geometric misalignment between maps may make it difficult to transfer a custom road attribute to the corresponding feature on the new base map. This difficulty is further compounded by additional complexities such as differences in road-segment length, network digitization density, and modeling conventions. Errors in this reconciliation of enriched data from the old base map to the new one manifest directly as routing inefficiencies, missed turn restrictions, safety hazards, and stale edits that never propagate to the refreshed map. Furthermore, the worldwide scale of map data renders manual transfer of enriched attributes to the new base map both cost-prohibitive and error-prone. It therefore becomes necessary to develop automated techniques that can accurately conflate map attributes for the maps data replay problem at industrial scale.

While geometry alignment in map data reconciliation has been addressed in the literature [14, 15], other complexities remain understudied. For example, due to the interconnected nature of road networks, it is not straightforward to define an atomic unit for re-playability. Similarly, development of a scalable methodology to match road networks that accounts for road-segment length and digitization variability, and transferring heterogeneous attribute types onto the matched road network has not been extensively studied. Intersection point features such as traffic signals and stop signs (nodes), way-attached attributes such as speed limits, access restrictions, and dimensional limits (ways), and relational features such as turn restrictions and lane connectivity (relations) each require specialized post-processing after matching to comply with the distinct payload semantics associated with each attribute type.

In this paper, we address the aforementioned gaps in the existing literature for map data conflation in a maps data replay setting and make the following contributions:

- **Cluster Partitioning**: We propose a clustering methodology that isolates connected components of related changes, each independent of the others. This ensures that each independent change component can be replayed onto the new base map in parallel, addressing a key scaling bottleneck in conflating maps at worldwide scale.
- **Road Linestring and Intersection Map-Matching**:We propose two distinct machine-learned map-matching models. A road map-matching model uses a logistic-regression classifier to compare road-network segments between the change component (the data to be replayed) and the new base map, classifying each candidate pair as matching or non-matching. Matched segments are then assigned their best counterpart using geometric checks such as intersection-over-union. This segment-level correspondence drives the transfer of both way (segment) attributes and node (intersection-level) attributes through a single matched-way frame. A second, random-forest model matches the intersections that anchor turn restrictions, scoring candidate source-to-target pairs over geometric and topological features, so that a turn restriction is recovered only when its anchoring junction is matched correctly. Together the two models achieve strong matching accuracy on the road network, a critical prerequisite for downstream node, way, and relation attribute transfer to the new base map.
- **Attribute Transfer Post-Processing**: The road and intersection matching output is used to replay node attributes (e.g., traffic signals, stop signs), way attributes (e.g., speed limits, access restrictions, lane counts), and relational attributes (e.g., turn restrictions, lane connectivity). This stage also performs important tasks such as unit standardization (to avoid corrupting maps with multiple units for the same attribute), tolerance-bounded conflict detection (to surface conflicting changes to the same road element), conflict volume mitigation using a confidence-based source priority schema, and way-splitting at attribute boundaries when needed.

Finally, an orchestration layer ties the aforementioned processing stages into one holistic pipeline. This enables the node, way, and relation attribute replay processes to share intermediate states such as matched road and intersection pairs and way-split decisions. This results in producing a non-conflicting, holistic change set for writing to the new base map.

The remainder of the paper is organized as follows. Section 2 surveys prior work on conflation, attribute transfer, and fingerprint-driven selective computation. Section 3 describes our proposed system, detailing components such as change-driven cluster partitioning; node, way, and relation matching using ML-powered conflation models; a post-processing module to create feature-specific change sets; and cross-stage orchestration. Section 4 reports the evaluation on the OpenStreetMap Full-History Planet source and the Overture Maps target across three regions. Section 5 discusses limitations, future work, and concluding remarks.

2 Related Work

2.1 Geometry conflation

Saalfeld [15] coined *conflation* for rubber-sheeting between spatial datasets that describe overlapping features. Ruiz and colleagues [14] surveyed two decades of follow-on work, organizing approaches by matching unit, matching criterion, and merging strategy. Classical OpenStreetMap-focused work includes sparse matching against georeferenced imagery [21] and polygon-based abstractions for transit-authority alignment [4]. Hacar and Gokgoz [6] observed that different road patterns demand different matching strategies, a finding that informs how we score candidate intersection pairs by local topology in our turn-restriction stage (Section 3.5).

Recent matchers are learned. MAYUR [1] uses early pruning and rank-join to avoid the $O(n^2)$ candidate explosion. Map Stitcher [16] reconciles datasets through graph sampling, while KRAFT [7] builds a knowledge graph over both inputs, aligns features (using both joint graph-alignment and a geospatial feature encoder), and resolves the merge as a mixed-integer program. Both operate on geometry and node-level features; neither carries relational members such as turn restrictions across the merge, which is the gap our relation-conflation stage targets. In autonomous-driving settings, FlexMap Fusion [11] conflates HD maps with OpenStreetMap to reduce map-building labor, and Ali and colleagues [2] align linear-referencing-system maps against OpenStreetMap with open tooling.

We treat centerline geometry conflation as solved infrastructure rather than as our contribution. An upstream geometry matcher aligns source and target centerlines; our pipeline consumes that aligned output. We do not realign centerlines. Instead we transfer the node, way, and relation payload onto centerlines that have already been brought into correspondence, a problem the geometry-conflation literature does not address.

2.2 Attribute conflation

Most attribute-transfer work targets a single attribute class. He and colleagues [8] describe road-attribute inference networks and attribute-learning frameworks that predict missing attributes on a single road graph rather than transferring them across providers. Multi-task attribute prediction shares a backbone across attribute heads for data efficiency.

Fusing attributes from fleet-sourced detections is a related subproblem. Kango and colleagues [10] describe a high-precision pipeline for fleet-sourced traffic signs that resolves duplicate detections and attaches each sign to the correct segment. Our intersection-attribute stage (Section 3.3) takes a complementary view: rather than fusing sensor detections into a single map, we transfer node features already attached on a source provider onto the matched target way, with proximity-bounded duplicate suppression. To our knowledge, prior cross-provider work has not jointly handled the breadth of way-level and node-level attribute classes we target, nor reconciled conflicting jurisdictional attributes such as speed limits using provider-level provenance scoring.

2.3 Relational features on road networks

Turn restrictions are topology-dependent. A restriction binds node and way features by role (`from`, `via`, `to`); if any member is split, renumbered, or absent on the target, the relation is invalidated. Most prior work treats restrictions as provider-internal state: each provider publishes its own restriction schema and expects wholesale ingestion, so a restriction defined against one provider's identifiers does not survive a move to another. We are not aware of prior work that transfers turn restrictions across providers whose topologies disagree.

Our relation matcher reformulates turn-restriction transfer as a learned matching problem rather than a geometry-projection heuristic. It first performs *intersection matching*: it scores candidate (source-intersection, target-intersection) pairs with a random forest over geometric and topological features (Section 3.5), then resolves *road matching* by reassigning the `from` and `to` roles on the target. Casting the via-node correspondence as supervised classification lets the model weigh valence agreement, intersection overlap, and local turning-function shape jointly, instead of relying on a fixed distance-and-overlap rule. Empirically, a heuristic distance-plus-IoU baseline reaches only about 49% recall at the high-precision operating point, whereas the learned intersection-matching classifier reaches roughly 92% test accuracy on held-out intersection pairs and recovers substantially more restrictions at the same precision. This gain motivates the learned reformulation.

Aerial-imagery-driven *attribute inference* operates upstream of any cross-provider transfer. Such systems [8] *infer* road attributes from satellite and GPS data, whereas our pipeline *transfers* attributes already known on the source. The two lines are complementary rather than competing.

2.4 Fingerprinting and selective re-computation

GIS provenance research has tracked feature lineage through service chains [9, 18, 19]. We attach a per-attribute provenance tag to each conflated attribute, recording the originating source and keying the conflict-resolution rule applied to that attribute. The same provenance fingerprint also selects which source edits are in scope for transfer, so the conflation workload is driven by the set of changed edits rather than by the full map. Selective re-computation is classical in incremental systems, but to our knowledge no prior conflation work organizes a continental cross-provider transfer around change-driven connected-component clusters. Vector-map watermarking [5, 17] pursues a different goal, copyright, yet shares the abstraction of per-feature identifiers that support downstream reasoning.

Table 1: Coverage of representative prior conflation systems across the dimensions our system targets: node (point-feature) transfer, way (segment-attribute) transfer, relation (turn-restriction) transfer, cross-provider operation (as opposed to single-graph inference), continental scale, and change-driven (fingerprint-gated) re-computation.

System	Node	Way	Rel.	X-prov.	Cont.	Chg.
Saalfeld [15]	✗	✗	✗	✓	✗	✗
MAYUR [1]	✓	✗	✗	✓	✗	✗
Map Stitcher [16]	✓	✗	✗	✓	✗	✗
KRAFT [7]	✓	✗	✗	✓	✗	✗
FlexMap Fusion [11]	✓	✓	✗	✓	✗	✗
Kango et al. [10]	✓	✗	✗	✗	✗	✗
He et al. [8]	✗	✓	✗	✗	✓	✗
This work	✓	✓	✓	✓	✓	✓

2.5 Cross-provider map migration as a problem class

Continental-scale migration from one base map to another, with the non-geometry payload preserved, does not appear to have been treated end-to-end in the open literature. FlexMap Fusion [11] is the closest named system, but it targets AV HD-map construction rather than provider migration at scale, and it does not handle relations. Existing conflation papers concentrate either on the matching model in isolation or on a single attribute class. To our knowledge, no prior open-literature system jointly handles node, way, and relation transfer onto an already-aligned target at continental scale; we are not aware of an end-to-end treatment of this problem class. Table 1 summarizes the axes on which each prior system covers or omits the dimensions we address.

3 Method

3.1 System overview

The system consumes a source map S in the OpenStreetMap node–way–relation (NWR) schema, a target map T already geometry-conflated against S by an upstream road-geometry step, and a fingerprint $F \subseteq S$ marking the source features whose edits are in transfer scope. It produces a changeset C applicable to T: an ordered sequence of `create` and `modify` operations on nodes, ways, and relations, each annotated with a per-attribute provenance tag. The system has four components (Figure 1):

(1) **Partition** (Section 3.2). In-scope source edits, selected by provenance fingerprint and reserved-identifier range, are grouped into independently conflatable clusters, the connected components of a shared-node graph over the affected

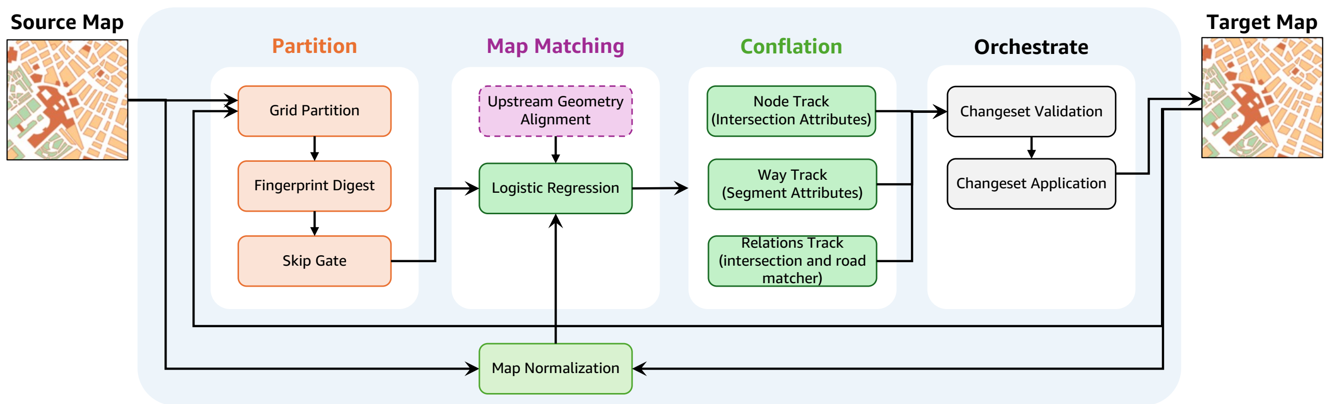

Figure 1: System architecture. In-scope source edits are partitioned into independently conflatable clusters, a logistic-regression model matches and conflates node and way attributes, a random-forest relation matcher recovers turn restrictions by two-stage intersection-and-road matching, and an orchestration layer validates and applies a single changeset to the target. The dashed block marks upstream geometry alignment consumed as infrastructure.

ways. Oversized clusters are subdivided by element count, bounding-box area, and S2 cells.

(2) **Map matching** (Sections 3.3–3.4). A logistic-regression model conflates both intersection-level (node) and road-segment (way) attributes. We specify the map-matching model at interface level, covering its inputs, feature set, and outputs. Its training and internal configuration are out of scope. It consumes normalized source and target segments and emits matched pairs with positional offsets that drive node and way conflation.

(3) **Conflation of relations** (Section 3.5). A random-forest matcher reformulates turn-restriction conflation as intersection matching followed by road matching, scoring candidate intersection pairs with geometric and topological features.

(4) **Orchestrate** (Section 3.6). A coordination layer ties the stages through a shared intermediate state of matched pairs, standardized attribute frames, and way-split decisions, enforcing a stage order in which relation construction consumes post-split way ids.

3.2 Change-driven cluster partitioning

Conflating an entire provider migration as one monolithic job couples cost to map size and makes frequent refresh impractical. We instead partition the workload by the changed edits, grouped into independently conflatable units (Figure 2). The same partitioning serves two modes: a one-time *bootstrap* that migrates a whole provider at continental scale, and an *incremental refresh* that replays ongoing edits, where the change set is small.

Edit selection. From the source snapshot we select the in-scope edits. A way is retained if it carries the in-scope provenance fingerprint or its identifier falls in the reserved organized-edit range. The selection can be constrained further to a region and date window. This step chooses *which* features are candidates for transfer ($F \subseteq S$) by provenance and identity; the new-versus-existing decision is deferred to the per-cluster matching stages.

Connected-component clustering. We build a graph in which ways are vertices and an edge joins two ways that share a node, and compute clusters as its connected components with a distributed graph engine. Each component is one cluster. Because membership follows shared-node connectivity rather than spatial proximity, edits that topologically interact fall in the same cluster while distinct clusters remain independent by construction. We therefore perform no spatial bounding-box merge. Public ways adjacent to a selected edit are attached as topological context, whereas only the selected edits drive the transfer.

Bounding cluster size. To keep any single unit's fetch-and-apply footprint bounded, oversized clusters are subdivided by an element-count threshold and a bounding-box-area threshold, and large clusters are further divided into geometry-stable S2 cells so that subunits have non-overlapping boundaries and remain independently conflatable. Each cluster is emitted as a self-contained record (cluster id, source NWR to transfer, topological-context NWR), partitioned by country and cluster id.

Apply and version-conflict retry. Because clusters are independent, they distribute across a worker pool and their changesets apply independently. A cluster whose changeset fails to apply because the target element version moved since it was read is re-queued and re-driven. A cluster whose changeset still fails after retries is routed to a dead-letter sink for inspection rather than silently dropped.

3.3 Node conflation

Point features in F (traffic signals, stop signs, combined signs) are properties of intersections, not of a single way. Transferring one requires two decisions. The first is which target way it belongs to, and the second is where along that way it sits. The map-matching model supplies the way correspondence; this stage places the node.

Intersection association. For each source node we take the nearest target intersection within a 50 m buffer, associating it only if the two share a way reference, which prevents attachment to a parallel road

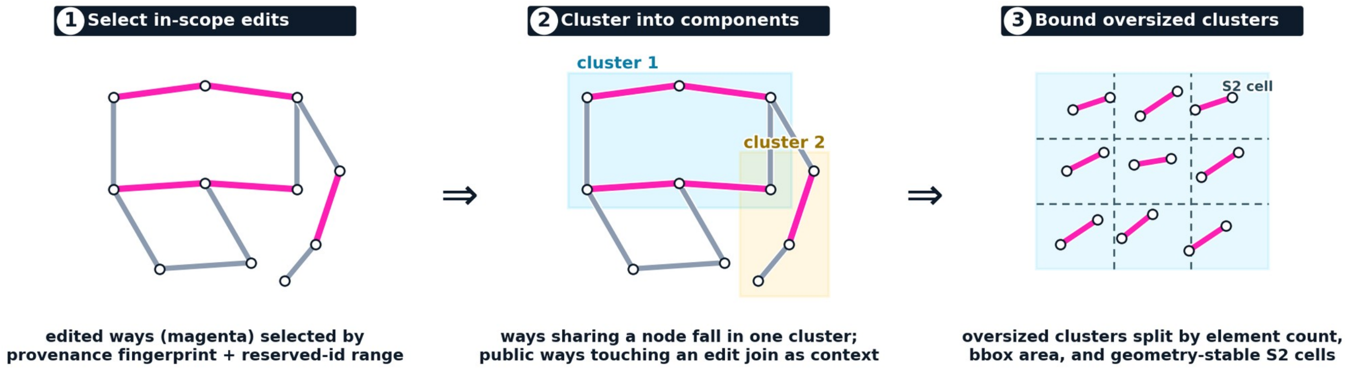

Figure 2: Change-driven cluster partitioning. (1) In-scope edited ways (magenta) are selected from the public network (grey) by provenance fingerprint and reserved-identifier range. (2) Ways are vertices of a shared-node graph (an edge joins two ways sharing a node); its connected components are the clusters, and public ways touching an edit join as topological context. (3) An oversized cluster is subdivided by element count, bounding-box area, and geometry-stable S2 cells, so every unit is bounded and independently conflatable.

that merely passes nearby. We drop the node when its associated intersection sits within 25 m of another intersection on a road of class motorway through tertiary, because such dense, high-class junctions are interchanges where the association is unreliable and a misplaced control device is worse than a missing one.

Placement and deduplication. The on-road position is interpolated about 10 m from the intersection along the associated way, respecting any `direction` tag; service roads are excluded so signals are not synthesized onto driveways. We project the position onto the matched target way, swapping `direction` where the two geometries run opposite. If a target node of the same type already exists within 35 m we skip; otherwise we emit a node-create with a way-node-insert, fingerprinted with the source provider and attribute name.

3.4 Way conflation

The way stage transfers segment-level attributes such as speed limits, lane counts, and access restrictions by vehicle class, among others. It runs in seven steps.

(1) Preprocessing and normalization. Both S and T are filtered to vehicle-access roads. Normalization then splits each multi-intersection way into atomic intersection-to-intersection segments and strips their tags, so a segment carries only its road class and geometry. This common representation lets the matcher compare segments consistently across maps that subdivide the same road differently.

(2) Map matching. The normalized segments are passed to the logistic-regression matcher, which returns matched pairs with their positional offset. Matching runs in both directions within a local map boundary; each source way's target neighbors are scored on start/end distance, Hausdorff and Fréchet distance, intersection-over-union, major-road status, and segment length. A source segment is classified as *existing* if its best-scoring target candidate clears the 0.5 probability threshold, and is matched to that candidate; otherwise it is classified as *new*.

(3) Denormalization and unit standardization. Matched segments map back to their NWR way ids, with 1:N maps resolving over-splits, and values are standardized to canonical units (miles per hour, pounds) while the original tag string is preserved for verbatim re-emission. Sub-attribute values are swapped where source and target geometries run opposite.

(4) Recall accounting. For each attribute category we record, for monitoring, the fraction of in-scope features matched confidently and the fraction rejected by sanity ranges.

(5) Intersection-attribute conflation. Traffic signals and stop signs are conflated against the node stage's matched-way frame, so one map-matching model serves both segment and intersection attributes.

(6) Conflict detection. For each matched pair and attribute we choose *skip*, *add*, or *update*: source missing → skip; target missing → add; both within tolerance → skip; both out of tolerance → update. Tolerances are attribute-specific (for example, 1 mph for speed, 3 inches for dimensions), and an update then passes a sanity-range filter that diverts out-of-range values to a rejection sink. Each attribute carries a split mode: `no_split` for attributes uniform along a way (speed, lane count) and `split_beyond_th` for those that may change along it (height, access restriction).

(7) Changeset creation and fingerprinting. For each approved target way we walk its geometry and find boundaries where a `split_beyond_th` attribute changes. If all sections agree we emit one way-modify; if they disagree and each fragment exceeds 40 m we split, adding a boundary node, truncating the original way, and creating a new way for the second section. Fragments under 40 m are absorbed into an adjacent section, because routing engines handle tiny fragments poorly and the split would add topology without adding information.

3.5 Relation conflation

Relations are the hardest category because each is defined over several primitives at once, not a single feature. A turn restriction binds a source `from` way, a `via` element (a node, or a sequence of way members), and a `to` way into one prohibition, mandate, or permission. If any member cannot be located on the target topology, the restriction cannot be rebuilt. We recast it as two matching problems in sequence: intersection matching, then road matching (Figure 3).

Intersection matching. A turn restriction is fundamentally an intersection: the `via` element identifies a target junction together with an angularly constrained pair of ways meeting at it. We extract the source intersection either directly from the `via` node when one is given, or by inferring it geometrically as the node shared by the `from`-way and `to`-way. We then enumerate candidate target intersections within roughly a 20 m buffer of the projected source via-coordinate. A candidate is any node with graph valence greater than one, that is, any junction. A binary random-forest classifier labels each (source intersection, candidate target intersection) pair as match or non-match over a set of geometric and topological features: graph valence of the two via-nodes and their difference; intersection-over-union and intersection-over-geometry in both directions; Hausdorff distance between the two intersection footprints; distance statistics over the candidate via-nodes inside the buffer; a dominant-angle turning function, in which each emanating road is treated as a unit vector from the via-node and summed to a resultant angle; and per-quadrant turning-function shape descriptors. We keep the highest-probability candidate that clears a confidence threshold. We evaluated alternative classifiers (gradient-boosted trees and a logistic-regression baseline) and selected the random forest for its behavior at high-precision operating points. Section 4.1 quantifies its recall against a deterministic distance-and-IoU baseline.

Road matching at the matched intersection. Selecting the target intersection leaves the question of which target road plays the `from` role and which plays the `to`. We shift the source intersection onto the matched target intersection and translate the source `from` and `to` ways by the same offset. At the target junction we score every emanating target road by ray-angle difference and assign each role to the road with the minimum angle difference. A 45-degree eligibility gate rejects assignments whose best angle is too large to be credible, and one-way-directionality and road-class consistency checks suppress matches that are angularly plausible but semantically invalid. A final angle-consistency check on the resulting (`from`, `to`) pair confirms the match before the rebuilt restriction is emitted.

Provenance and deduplication. Every conflated relation carries a per-relation provenance tag, so a downstream consumer can pick a precision-recall operating point by threshold filtering rather than re-running the model. Before emission we anti-join against the target's existing relations on the matched (`from`, `to`, `via`, `restriction_value`) key to avoid double-stamping. When the way stage splits a way, any relation that referenced it is repointed to the correct new segment, using the split-decision frame.

3.6 Cross-category coordination and validity

The orchestration layer ties the stages through a single intermediate state: the matched-pair frame with denormalized identifiers, the standardized attribute frame, and the way-split decisions. It enforces a fixed order: node conflation first, way conflation next (which produces the split-decision frame), and relation conflation last (which consumes it). A relation therefore always references post-split target way ids and never a stale identifier. Some relations span a cluster boundary; we detect these at partition time and resolve them in a reconciliation pass after the per-cluster jobs finish, appending the additional relation-create operations the boundary made impossible to emit in a single cluster.

Table 2: Source and target map versions used in experiments and results analysis described in Section 4.

Map	Role	Version
OpenStreetMap (planet snapshot)	Source base-map	2025-11-29
Overture Maps Transportation Theme	Target base-map	2024-07-22

Validity check. Before a per-cluster changeset is written we apply syntactic and semantic checks: every node id referenced by a way-modify must resolve to an existing target node or to a same-changeset create; every relation-create must reference resolvable ways and via-nodes; sanity-range violations appear in the rejection sink rather than the changeset; and every attribute change carries a provenance fingerprint. Changesets that pass are handed to a downstream applicator over a retry-backed queue. By construction, every operation reaching it is topology-safe.

4 Experiments and Results

In this section, we present our experiments and results of the proposed conflation methodology of replaying edited maps data between different base-maps. We evaluate the proposed methodology by replaying edited maps data between two base-maps. A fixed OpenStreetMap version serves as the source base-map, and the edits built on top of it stand in for the custom edits an operator would replay; a snapshot of the Overture Maps Transportation Theme is the target base-map. We deliberately pair a newer OpenStreetMap snapshot with an older Overture release so that the edits OpenStreetMap accumulated over the intervening period, for the regions analyzed, form the body of changes to replay onto the target. Both sources are open and versioned (Table 2), so the benchmark is reproducible from public artifacts. For attribute completeness we evaluate over three anonymized regions that bracket the structural regimes of a continental run (dense-urban Region A, mixed Region B, and rural Region C); the routing-impact and address-proximity analyses are reported for a single region. A held-out edit window $D_1 \rightarrow D_2$ from each region's history supplies the in-scope edit set. Features are conflated independently within clusters, the connected components of the shared-node graph. We report two metric families: conflation accuracy (Section 4.1), and impact on downstream routing and address-proximity coverage measured at scale with a routing engine (Sections 4.2–4.3). We close with an error analysis (Section 4.4).

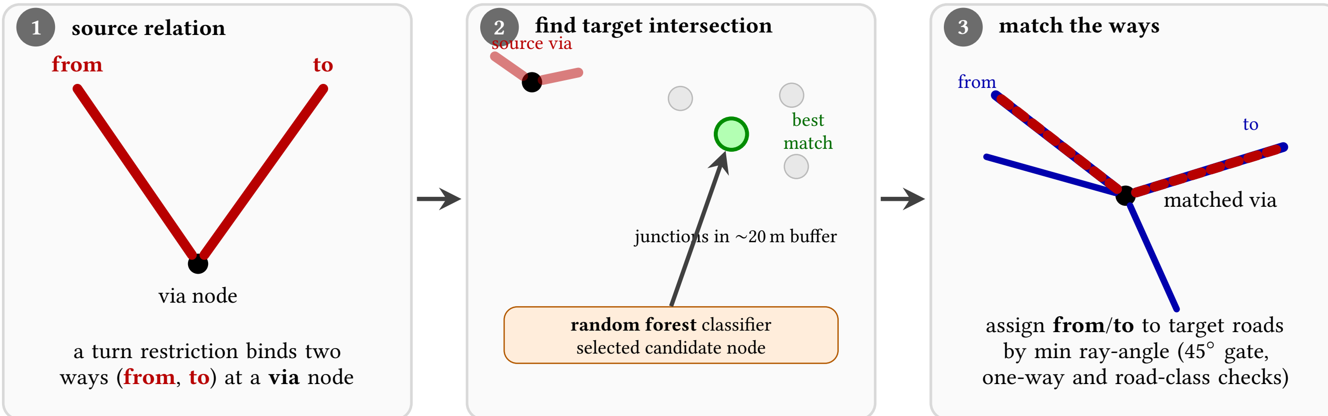

Figure 3: Relation conflation in two steps. *Step 1 (intersection matching):* candidate target junctions within a ~20 m buffer of the projected source intersection are scored by a random forest over ~21 geometric features, and the best-scoring candidate is kept. *Step 2 (road matching):* the source `from`/`to` roads are shifted onto the matched junction and each role is assigned to the emanating target road by minimum ray-angle difference, subject to a 45° eligibility gate and one-way and road-class consistency checks.

4.1 Attribute Replay Experiments

In these experiments, we analyze the accuracy of the attribute replay process, which transfers custom edits to map attributes from the OSM base-map onto Overture maps. The attributes span node features such as stop signs and traffic signals, way features such as road directionality, speed limits, and road names, and relation features such as turn restrictions. Our evaluation uses human auditors, who compare the pre-conflated attributes on the OSM map against the post-conflated Overture map. For each transfer, auditors evaluate whether the road was matched correctly, whether the attribute was transferred to the correct extent of the road, and whether way splitting was required. Because operating these maps at world scale demands high accuracy, we run at a high precision of 99%. Table 3 reports recall at 99% precision by entity class against a hand-curated ground-truth set.

Relation rows give turn-restriction recall at two precision points, reflecting fully automated (~99%P) versus human-review (~98%P) consumption. Across the three regions the system transfers ~95% of intersection and road-segment attributes and recovers 55–62% of turn restrictions at ~99%P, rising to 68–73% at ~98%P. Attribute transfer is stable across regimes, while relation recall varies more and is lowest in the rural regime, where sparse, irregular intersections make restriction matching harder. Relation recall trails attribute transfer throughout because a restriction is recovered only if both bounding intersections match *and* the `from`/`to` roles are assigned correctly, so per-stage errors compound.

4.2 Routability Metrics Evaluations

Analysis in Section 4.1 measures the conflation quality of map elements, but it does not capture downstream impact on routing and navigation. Since routing runs on a road network graph *derived from* a base map, we re-route a large trip sample on a graph from the OpenStreetMap snapshot (control) and from the conflated map (experiment), comparing routed distance per trip [20]. The instrument localizes *where* routing shifts, a release-gating signal whose correctness interpretation we return to in Section 4.4. Figure 4a illustrates a single diverging trip.

Table 3: Per-category transfer completeness on the open OSM→Overture benchmark, over the three structural regimes. Node and way rows are attribute-transfer recall; relation rows are recall at a fixed precision point.

Category	Region A	Region B	Region C
Node attrs. (recall)	~94.9%	~95.0%	~96.0%
Way attrs. (recall)	~94.5%	~95.0%	~94.5%
Relation (R@~99%P)	~60%	~62%	~55%
Relation (R@~98%P)	~73%	~70%	~68%

Table 4 reports a localized urban area, scoring the trips routable on both graphs against the control. The change is *bidirectional*: routes shorten and lengthen in differing proportions on each graph, and only a small fraction are unchanged. For the conflated map the shorter and longer routes nearly balance, so the net barely moves, whereas the unconflated target is more imbalanced; in both cases each affected route shifts substantially more than the net, which understates per-route churn. The headline is the parity gap. Without replay the vanilla Overture target routes +4.8% longer in aggregate than the enriched OpenStreetMap control; replaying the custom edits narrows this to +0.94%, bringing the conflated target to near parity with the enriched control. The residual gap of under 1% is attributable to recall loss in the replay process (Section 4.4). We repeated the same comparison across many region cells of a country-scale map. A trip-weighted aggregate over the full sample, which folds in trips routable on only one graph and is thus a broader and noisier denominator than the comparable subset, rises modestly in distance and by less in time. The change concentrates in denser regions, where the highest trip-density tertile shifts most and the

Table 4: Routing impact for a localized urban area. Directional split of trips re-routed on the conflated-map graph (E) and the unconflated target graph (U), each against the control (C). The control is the OpenStreetMap planet snapshot and U the unconflated Overture target, both from Table 2. Mean $|\Delta|$ is the mean absolute change in routed distance over the affected routes. The same comparison was repeated across many region cells of a country-scale map. Each net aggregate understates the per-route churn, and the unconflated target diverges from the control more than the conflated map does.

	Conflated (E)		Unconflated (U)	
Outcome vs. C	Routes	$\|\Delta\|$ (m)	Routes	$\|\Delta\|$ (m)
Shorter	49.9%	36.6	65.5%	48.5
Longer	44.2%	59.4	29.0%	20.5
Unchanged	5.9%	–	6.0%	–
Net distance	+0.94%		+4.8%	

lowest barely changes, and is consistent with changes in geometry and topology. Clustering divergence-onset points within a fixed radius [3, 12] localizes recurring changes to re-segmented corridors.

4.3 Network coverage: address-to-road proximity

A complementary instrument measures address to road proximity. For each address point we compute the geodesic distance to the nearest *vehicular* segment in each graph (k-nearest candidates refined to the true perpendicular distance), excluding classes addresses do not front (motorway, trunk, emergency, pedestrian, construction, ferry). Greater near-range mass means more reliable attachment [2], though proximity to *a* segment does not guarantee it is the correct one. Conflation moves the distribution toward the control, recovering near-range coverage that the unconflated target leaves in the mid and far buckets: the unconflated target reaches only ∼96% within 50 m, whereas the conflated and control graphs sit at near parity, with medians within 0.3 m and over 98% within 50 m (Tables 5–6). The conflated and control graphs differ only in the upper tail, where the conflated graph cuts points beyond 100 m from 1,966 to 345. Figure 4b shows the per-address mechanism.

4.4 Error Analysis

Relation recall loss concentrates on three patterns. The first, *carriageway modeling differences*, arises when a dual carriageway is represented as two one-way roads in one map and a single bidirectional centerline in the other, so a restriction's bounding intersections do not match cleanly. The second, *very-short via-roads at sharp angles*, occurs where a small coordinate shift swings the dominant-angle turning feature past threshold and pushes the true candidate below the confidence cutoff. The third is *training-label noise*, where ground-truth restriction labels are themselves ambiguous, capping achievable recall. Way-conflation residuals are instead dominated by attribute schema drift, in which directional sub-attributes are encoded under different provider keys, so the attribute is present

Table 5: Address-to-road distance distribution (% of points per bucket) for a region, on the control (C), conflated-map (E), and unconflated target (U) graphs. C is the OpenStreetMap planet snapshot and U the unconflated Overture target (Table 2). The conflated graph reduces the share of addresses left far (>100 m) from any road. Bucket rows are independently rounded from higher-precision values, so a column need not sum to exactly 100.

Distance (m)	C	E	U
0–25	85.81	84.83	81.20
25–50	13.00	13.85	14.90
50–100	1.05	1.29	3.77
100–200	0.10	0.02	0.09
>200	0.03	0.00	0.04
within 50 m	**98.81**	**98.68**	**96.10**

Table 6: Address-to-road proximity percentiles (m); lower is better, for the control (C), conflated-map (E), and unconflated target (U) graphs. C is the OpenStreetMap planet snapshot and U the unconflated Overture target (Table 2). The unconflated target's percentile distances run slightly above the conflated and control graphs, with the gap widening toward the upper percentiles, and its within-50 m coverage is correspondingly lower. The final row is not a distance percentile but the coverage percentage within 50 m (%), repeated from Table 5 for reference.

Statistic	C	E	U
Median (m)	12.5	12.8	13.2
p95 (m)	34.6	35.3	36.3
p99 (m)	52.0	53.7	55.5
% within 50 m	98.81	98.68	96.10

in the target but not recognized as the same field. Each residual reflects a deliberate precision-favoring choice rather than a matching error, and each is a target for the relaxed-precision operating point of Section 4.1.

The operational instruments carry a different limitation. Both measure *change*, not *correctness*. A shorter route may signal a corrected geometry or a spurious connection, and a closer snap may attach an address to the wrong segment. They are thus release-gating and regression-detection tools. Certifying a flagged shift as an improvement needs a ground-truth audit against observed trajectories, which we leave to future work.

5 Conclusion

We presented a scalable conflation system that lets operators replay accumulated map edits onto a refreshed base map, or swap base-map providers, without losing the road-network entities and transition relations that ordinary geometry conflation leaves behind. The system rests on four components. Change-driven partitioning selects in-scope edits and groups them into independently conflatable clusters, the connected components of a shared-node graph, subdivided by size and geometry-stable S2 cells, so that one pipeline

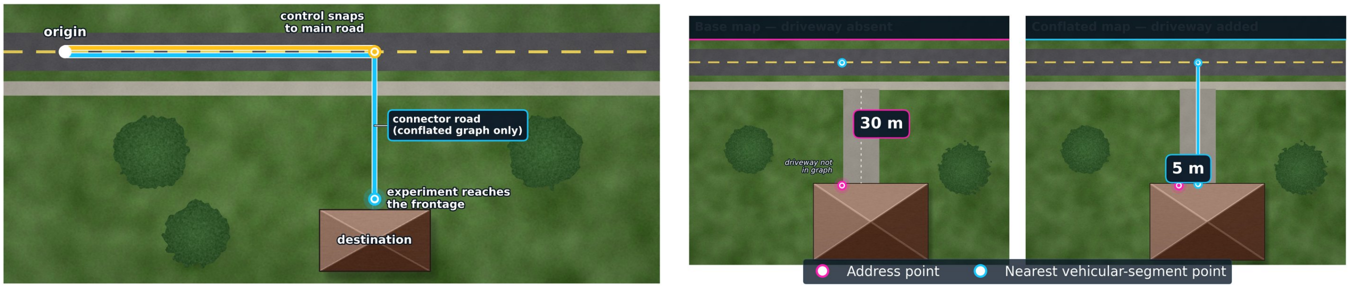

(a) Route divergence under recosting for one trip. A connector road runs from the main road toward the destination; it exists in the conflated-map graph but is absent from the control graph. The amber route (control, the OpenStreetMap snapshot) lacks the connector and ends at its nearest snap point on the main road; the cyan route (conflated map, experiment) descends the connector and reaches the building frontage.

(b) Proximity change from a missing driveway. Left: the driveway is absent from the graph, so the address snaps to the distant street. Right: conflation adds the driveway as a vehicular segment, so the address snaps nearby.

Figure 4: Operational instruments illustrated on single cases: recosting (a) and address-to-road proximity (b).

serves both a continental bootstrap migration and a small-scale incremental refresh. A logistic-regression map-matching model with road-segment normalization conflates both road-segment (way) and intersection-level (node) attributes through a single matching interface. A random-forest relation matcher recasts turn-restriction conflation as intersection matching followed by road matching, scoring candidate intersection pairs with geometric and topological features and then transferring the restriction onto the matched roads. An orchestration layer ties these stages into one production pipeline, sharing matched pairs, standardized attribute frames, and way-split decisions across the work.

On the OpenStreetMap and Overture Maps benchmark, the system conflates about 94.9% of intersection attributes and 94.5% of road-segment attributes, and recovers roughly 60% of turn restrictions at the high-precision operating point, well above the ~49% a distance-and-IoU heuristic reaches at the same precision. The intersection-matching classifier reaches about 92% test accuracy. Two directions remain open. The hand-tuned geometric scorer rejects a long tail of correspondences. Recovering these with learned geometry models is one avenue we intend to explore. A second is reconciling relations whose members straddle cell boundaries, since the current partitioning processes each cell in isolation.

References

[1] Gorisha Agarwal, Laks V. S. Lakshmanan, Xiaoming Gao, Kevin Ventullo, Saurav Mohapatra, and Saikat Basu. 2021. MAYUR: Map conflAtion using earlY prUning and Rank join. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '21)*. ACM, 550–553.

[2] Gibran Ali, Neal Feierabend, Prarthana Doshi, Whoibin Chung, Simona Babiceanu, and Michael Fontaine. 2025. Automated Route-based Conflation between Linear Referencing System Maps and OpenStreetMap Using Open-source Tools. In *2025 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 4443–4450. arXiv:2507.13939.

[3] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD '96)*. AAAI Press, 226–231.

[4] Hongchao Fan, Bisheng Yang, Alexander Zipf, and Adam Rousell. 2016. A polygon-based approach for matching OpenStreetMap road networks with regional transit authority data. *International Journal of Geographical Information Science* 30, 4 (2016). doi:10.1080/13658816.2015.1100732

[5] Danila Glazkov, Nikolay Chupshev, and Victor Fedoseev. 2023. Vector Tile Geospatial Data Protection Using Quantization-Based Watermarking. In *Proceedings of the 9th International Conference on Geographical Information Systems Theory, Applications and Management (GISTAM)*. SciTePress, 251–258. doi:10.5220/0012044500003473

[6] Muslum Hacar and Turkay Gokgoz. 2019. A New, Score-Based Multi-Stage Matching Approach for Road Network Conflation in Different Road Patterns. *ISPRS International Journal of Geo-Information* 8, 2 (2019). doi:10.3390/ijgi8020081

[7] Farnoosh Hashemi and Laks V. S. Lakshmanan. 2025. KRAFT: A Knowledge Graph-Based Framework for Automated Map Conflation. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM)*. 802–812. arXiv:2509.04684. doi:10.1145/3746252.3761291

[8] Yang He, Emre Eftelioglu, Mohamed Moustafa, and Amber Roy Chowdhury. 2023. A Highly Efficient and Effective Attribute Learning Framework for Road Graph from Aerial Imagery and GPS. In *BigSpatial '23: Proceedings of the 11th ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data*. doi:10.1145/3615833.3628594

[9] Liangcun Jiang, Peng Yue, Werner Kuhn, Chenxiao Zhang, Changhui Yu, and Xia Guo. 2018. Advancing interoperability of geospatial data provenance on the web: Gap analysis and strategies. *Computers & Geosciences* (2018). doi:10.1016/j.cageo.2018.05.001

[10] Vaibhav Kango, Hesham Eraqi, and Mohamed Moustafa. 2024. High Precision Map Conflation of Fleet Sourced Traffic Signs. In *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE. doi:10.1109/IGARSS53475.2024.10642165

[11] Maximilian Leitenstern, Florian Sauerbeck, Dominik Kulmer, and Johannes Betz. 2024. FlexMap Fusion: Georeferencing and Automated Conflation of HD Maps with OpenStreetMap. In *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 272–278. arXiv:2404.10879. doi:10.1109/IV55156.2024.10588414

[12] Paul Newson and John Krumm. 2009. Hidden Markov Map Matching Through Noise and Sparseness. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '09)*. ACM, 336–343. doi:10.1145/1653771.1653818

[13] Overture Maps Foundation. 2026. Overture Maps Foundation Transportation Theme. https://overturemaps.org/. ODbL transportation theme; accessed 2026.

[14] Juan J. Ruiz, Francisco J. Ariza, Manuel A. Urena, and Emilio B. Blazquez. 2011. Digital map conflation: A review of the process and a proposal for classification. *International Journal of Geographical Information Science* 25, 9 (2011), 1439–1466. doi:10.1080/13658816.2010.519707

[15] Alan Saalfeld. 1988. Conflation Automated Map Compilation. *International Journal of Geographical Information Systems* 2, 3 (1988), 217–228. doi:10.1080/02693798808927897

[16] Erfan Hosseini Sereshgi and Carola Wenk. 2024. Map Stitcher: Graph Sampling-based Map Conflation. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Spatial Big Data and AI for Industrial Applications* (Atlanta, GA, USA) *(GeoIndustry '24)*. Association for Computing Machinery, New York, NY, USA, 5–15. doi:10.1145/3681766.3699604

[17] Heyan Wang, Nannan Tang, Changqing Zhu, Na Ren, and Changhong Wang. 2025. A Blockchain Copyright Protection Model Based on Vector Map Unique Identification. *ISPRS International Journal of Geo-Information* 14, 2 (2025), 53.

doi:10.3390/ijgi14020053

[18] Jie Yuan, Peng Yue, Jianya Gong, and Mingda Zhang. 2013. A Linked Data Approach for Geospatial Data Provenance. *IEEE Transactions on Geoscience and Remote Sensing* (2013). doi:10.1109/TGRS.2013.2249523

[19] Peng Yue, Jianya Gong, and Liping Di. 2010. Augmenting geospatial data provenance through metadata tracking in geospatial service chaining. *Computers & Geosciences* 36, 3 (2010). doi:10.1016/j.cageo.2009.09.002

[20] Chiqun Zhang, Dragomir Yankov, Antonios Karatzoglou, Michael R. Evans, Florin Sabau, and Oussama Dhifallah. 2023. A Post-routing ETA Model Providing Confidence Feedback. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23)*. ACM. doi:10.1145/3589132.3625606

[21] Jiantong Zhang, Yueqin Zhu, and Liqiu Meng. 2011. Conflation of road network and geo-referenced image using sparse matching. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. doi:10.1145/2093973.2094011