International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 6, Nov-Dec-2026, ISSN (Online): 2395-566X

LOCALYN: A Hyperlocal Multi-Source Civic Event Intelligence Platform

Roopa T[1] Ayush Gowda P[2] Kumar K C[3] Tarun C L[4]
Department of Computer Science and Business Systems, B.M.S. College of Engineering, Bengaluru, Karnataka, India

Abstract- Information about road conditions, public infrastructure and local disruptions is recorded by several city-facing organisations. Since each organisation follows its own categories and reporting practices, observations about one incident may remain scattered across separate systems. This paper presents LOCALYN, a hyperlocal event-intelligence platform developed around Bengaluru civic data. The implemented system combines public-source ingestion with a cross-agency association module that links Bruhat Bengaluru Mahanagara Palike (BBMP) grievance categories to Bengaluru Traffic Police (BTP) incident causes. Matching uses a curated administrative crosswalk, ward compatibility, timestamp comparison and text similarity. A four-feature score retains the contribution of each signal, using the all-MiniLM-L6-v2 sentence-embedding model when available and Jaccard similarity as a fallback. The geographic service processes 225 ward polygons through Shapely and an STRtree index, with explicit handling for polygon boundaries and nearby-centroid fallback. Accepted record links are grouped into canonical events, which receive an interpretable Confidence from Heterogeneous Evidence Fusion (CHEF) score. The wider application includes news, BBMP and weather ingestion modules, SQLite storage, FastAPI endpoints and a React–Leaflet interface. In the implemented nine-pair benchmark, LOCALYN obtains an F1-score of 1.00; the date-and-ward baseline obtains 0.833, while the date–ward–category baseline obtains 0.000. These results provide an initial functional comparison of the implemented matching workflow. The paper describes the system architecture, matching procedure and current evaluation, and identifies extensions for live agency integration and larger-scale testing.

Key Words— urban informatics, civic data integration, event resolution, entity matching, geospatial analysis, taxonomy alignment, sentence embeddings, confidence estimation, Bengaluru

I. INTRODUCTION

A single disruption in a city can appear in several information systems. A damaged road may be reported by a resident through a municipal grievance portal, noted by traffic personnel because it slows vehicles, and mentioned in a local news update. Each record describes part of the same situation, but the records are created independently. They may use different incident labels, timestamps, location formats and levels of detail. As a result, a person checking only one source may not see the full picture, while a system that counts every report separately may treat one incident as several.This problem is visible in Bengaluru, where civic services operate across administrative wards and multiple public infor-mation channels. Matching reports is not simply a matter of comparing their text. A municipal category such as a fallen tree may correspond to a traffic cause recorded as an obstruction. Ward names can also change or appear in different forms, and coordinate-based observations need to be interpreted against administrative boundaries.

A useful system must therefore consider the meaning of a report, when it was recorded and where it occurred.LOCALYN was developed to address this problem through a modular civic event-intelligence platform. Its main association workflow links BBMP grievance records with BTP incident observations.

A category crosswalk supplies the administrative relationship between the two agencies; spatial and temporal checks narrow the candidate records; and text similarity contributes an additional comparison. The resulting pair score is retained with its component values so that a match can be inspected rather than presented as an unexplained label.

The platform also includes a ward-resolution service, canonical-event grouping and the CHEF confidence component.

© 2026 IJSRET

Supporting modules collect news, municipal and weather information, store reports in SQLite and expose data through FastAPI. A React interface with Leaflet provides the map-oriented presentation layer. The ingestion and cross-agency association modules are separate parts of the current implemen-tation: public-source collectors populate the civic report store, while the BBMP–BTP matcher provides the taxonomy-aware association workflow. The work makes four practical contributions. First, it provides a structured crosswalk between BBMP grievance subcategories and BTP incident causes, including mapping-strength labels. Second, it implements ward resolution over 225 geographic features, using spatial indexing, polygon tests and fallback handling. Third, it combines time, location, taxonomy and text in a weighted pair-matching procedure with readable decision details. Fourth, it groups accepted links into canonical events and attaches a CHEF score whose factors can be inspected by the application. Together, these components form an end-to-end engineering approach for organising fragmented civic observations around local events.

II. LITERATURE REVIEW

1. Event Detection and Entity Resolution

Event-detection systems identify reports that describe an occurrence and attempt to organise related mentions into a useful representation. Work on social-stream event detection has examined lexical patterns, temporal activity and clustering for identifying emerging events [1]. Entity resolution addresses a closely related task at the record level: deciding whether two records refer to the same real-world object or occurrence, even when their fields differ [2]. These ideas are relevant to

TABLE I

Literature Comparison Relevant To Localyn

Work	Focus	Relevance to LOCALYN
Atefeh and Khreich [1]	Social event detection	Temporal and textual event signals
Imran *et al.* [7]	Crisis-message processing	Heterogeneous public reports
Christen [2]	Record link-age	Pairwise record-resolution principles
Reimers and Gurevych [4]	Sentence em-beddings	Semantic text comparison
Mondal *et al.* [12]	Multimodal event fusion Smart-city	Multi-source event correlation Urban data harmonisation
Orchi *et al.* [13]	in-formation fu-sion	

civic information because reports from different agencies often describe the same incident using different administrative terms. LOCALYN treats cross-agency association as a pairwise record-resolution task. It first uses the BBMP–BTP crosswalk and location compatibility to identify plausible pairs, then evaluates the remaining evidence. Accepted pairs are repre-sented as graph edges and grouped using connected components. This differs from density-based clustering approaches such as DBSCAN, which form groups from density relationships in a feature space [3]. In LOCALYN, the pairwise association deci-sion is made before grouping, so each graph edge corresponds to an accepted record link.

2. Semantic Similarity

Text fields are useful when two agencies describe the same event in different words. Sentence-BERT produces sentence-level embeddings that can be compared using cosine similarity [4]. LOCALYN uses the all-MiniLM-L6-v2 model when the SentenceTransformers package and model are available. The text input is assembled from the description, category-related fields and location text supplied by each record. If embedding inference cannot be used, the matcher calculates token-set Jaccard similarity. This keeps the text component available in environments where a transformer model is not installed.

3. Geospatial Matching

Location comparison in civic applications often involves both coordinates and administrative names. A coordinate can be tested against a ward polygon, while a textual ward label must be normalised before it can be compared with another source. LOCALYN uses Shapely for geometry operations and an STRtree to retrieve a smaller set of candidate polygons for each coordinate. The service then performs containment and boundary checks, with a centroid-based fallback for points that do not resolve directly. This arrangement makes ward assignment an explicit stage of the association process rather than relying only on free-text place names.

TABLE II Research Gap Matrix

Capability	Prior work	LOCALYN design	Status
Cross-agency tax-onomy alignment	Domain-specific	Curated BBMP–BTP crosswalk	Implemented
Ward-aware match-ing	Often applicatio n-specific	225-polygon resolver	Implemented
Explainable pair score	Varies by method	Four visible score terms	Implemented
Canonical event grouping	Clustering/linkage	Union–find components	Implemented
Calibrate d confiden ce	Separate evaluatio n task	CHEF heuristic score	Further val-idation
Live multi-agency correlation	Uneven cov-erage	Modular ingestion + benchmark matcher	Partial

© 2026 IJSRET

4. Evidence and Confidence

Multi-source information systems must distinguish the number of reports from the amount of independent support behind an event. Data-fusion research considers how evidence from different sources can be combined, including cases where sources disagree or have different reliability [10]. Dempster–Shafer theory is one established framework for represent-ing uncertain evidence [5], [6]. LOCALYN's current CHEF implementation takes a more direct engineering approach: it combines configured factors for source diversity, agency reliability, agreement, geographic availability and freshness. The resulting score is intended to show how the system assessed the available evidence; it is not presented as a calibrated probability.

5. Urban Information Systems

Urban-computing research describes the integration of in-formation produced by city infrastructure, public services and residents [11]. Crisis-informatics work has also explored how public digital information can supplement conventional reporting during local events. LOCALYN applies these ideas to day-to-day civic observations, with emphasis on adminis-trative categories and ward-level geography. Its design allows additional feeds to be introduced without changing the core pair-scoring interface.

Recent urban-computing research treats city data as a combination of sensing, data management, analytics and service delivery rather than as isolated feeds [8], [11], [18]. Reviews of smart-city fusion emphasise the practical difficulty of harmonising heterogeneous sources with different semantics, coverage and reliability [13]–[15]. Multimodal event-detection research similarly identifies cross-source correlation as a distinct problem from detecting an event within one stream [12].

For civic reporting, volunteered geographic information also introduces questions of positional accuracy, provenance and uneven participation [9], [19], [24]. These observations motivate LOCALYN's explicit treatment of source provenance, ward assignment and cross-agency taxonomy alignment.

Record linkage theory provides the basis for deciding whether two records refer to the same underlying event [2], [16]. Data-fusion research distinguishes the integration

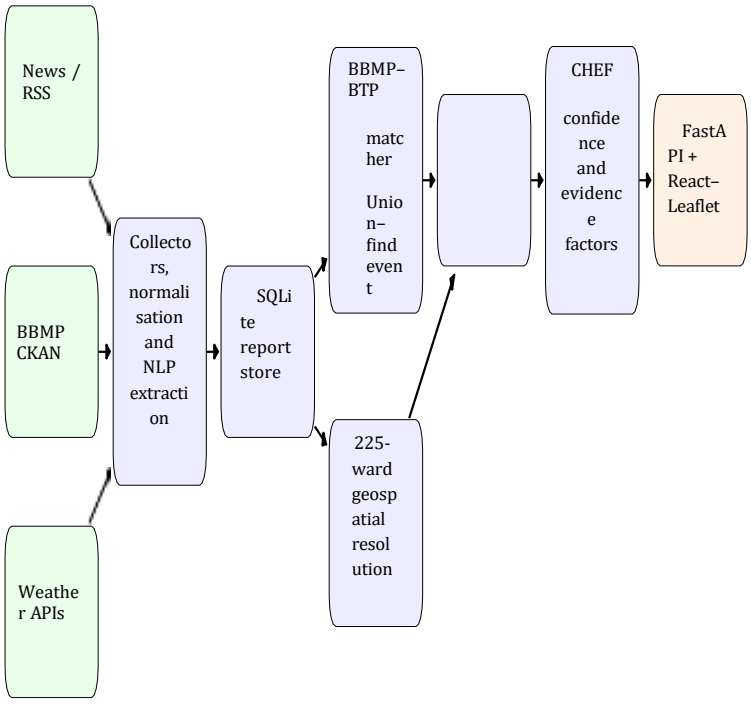


Fig. 1. Overall LOCALYN architecture. Public-source ingestion and the benchmark-driven BBMP–BTP association path are shown as related but distinct processing paths

TABLE III Data Resources And Their Roles In The Current Prototype

Resource	Access path	Role and boundary
News	NewsData.io / Google News RSS	Public news ingestion; locality and civic-topic filtering
BBMP	OpenCity CKAN	Municipal grievance ingestion and taxonomy source
Weather	OpenWeatherMap / Open-Meteo	Weather observations and contextual reports
Ward map	Bengaluru ward GeoJSON	225 administrative polygon features for spatial resolution
BTP	Curated bench-mark pairs	Used in BBMP–BTP matching benchmark; live connector not present

of evidence from simple concatenation of records [10]. In LOCALYN, sentence-level semantic similarity is an optional text signal rather than a substitute for administrative and geographic constraints [4], [22]. The CHEF component is designed for inspectable evidence aggregation; its current score is deterministic and should not be interpreted as a calibrated probability. Calibration methods such as temperature scaling and proper scoring rules provide a direction for subsequent validation [17], [23]. The display of individual factors is aligned with the broader aim of interpretable model outputs [20], [21].

© 2026 IJSRET

III. RESEARCH GAP AND RESEARCH QUESTIONS

The reviewed literature contains mature methods for entity resolution, semantic similarity, spatial indexing, and evidence aggregation. LOCALYN addresses the practical integration gap between these capabilities by combining an explicit adminis-trative taxonomy crosswalk with ward-boundary compatibility, temporal and textual evidence, and confidence-factor disclosure for the BBMP–BTP civic incident setting. This integration provides a reusable design pattern for other cities where public agencies maintain separate reporting systems and classification schemes.

The study is guided by the following research questions:

- **RQ1**: Does combining taxonomy mapping and ward-level spatial compatibility improve cross-agency incident matching relative to simple date–ward and raw-category baselines?
- **RQ2:** Can the matching decision be exposed through interpretable feature scores and textual explanations?
- **RQ3:** Can matched records be consolidated into canonical events with an explicit, inspectable confidence breakdown?

TABLE IV Data-Quality Dimensions And Current Handling

Dimension	Observed issue	Current handling
Source format	RSS, API JSON, CSV/GeoJSON	Source-specific ingestion adapters
Temporal fields	Missing or inconsistent timestamps	Parsing and fixed missing-time score
Location	Ward aliases, coordinates, boundary and cen-boundaries	Polygon lookup, troid fallback
Taxonomy	Agency-specific labels	159-row curated crosswalk
Text	Short, noisy, differently fall-phrased descriptions	Sentence embedding or Jaccard back
Provenance	Unequal source reliability	Agency reliability constants in CHEF

IV. SYSTEM OVERVIEW AND ARCHITECTURE

1. Architecture

The implemented application comprises a data-ingestion layer, processing and persistence services, a cross-agency fusion module, a geographic ward service, a FastAPI backend, and a React-based interface. The backend persists general reports in SQLite. The frontend uses React and Leaflet for the interactive application and map display.

2. Data Sources and Processing Boundaries

The ingestion code contains connectors for NewsData.io and a Google News RSS query, the OpenCity CKAN API for BBMP grievances, and weather services including OpenWeatherMap/Open-Meteo paths. News items are nor-malised, deduplicated by source and normalised title within a 24-hour interval, and filtered for Bengaluru locality and civic-topic relevance. spaCy is used for location and time extraction, with a custom locality list and regular-expression fallbacks. Text is deterministically classified into TRAFFIC, WEATHER, or OTHER.

The BBMP–BTP matching module operates as a complemen-tary fusion path, using the taxonomy crosswalk and a curated benchmark set to resolve inter-agency record pairs. The modular ingestion architecture is designed to accommodate a live BTP connector as an additional source integration.

3. Database and Service Layer

SQLite stores report identifiers, source metadata, raw text, category, location, coordinates, event time, confidence, confi-dence tier, confidence factors, normalised title, and ingestion time. FastAPI exposes health, report retrieval, ingestion, cross-walk summary and mapping, canonical-event, benchmark, and

TABLE V Bbmp Data Profile Represented In The Implementation

Attribute	Implementation evidence
Source endpoint	OpenCity CKAN
Source type	Municipal grievance records
Taxonomy dimensions	Category and subcategory
Crosswalk registry size	159 rows
Mapped event fami-lies	Road damage, waterlogging, tree fall, construction-related, and others
Raw BBMP row to-tal	Not fixed in repository; depends on source snapshot

TABLE VI PRINCIPAL LOCALYN SOFTWARE MODULES

Layer	Implementation	Responsibility
Ingestion	Python collectors	News, BBMP and weather retrieval
NLP	spaCy + rules	Locality/time extraction and report classification
Storage	SQLite / SQLAlchemy	Report and event persistence
Fusion	matcher, crosswalk, cluster	Candidate scoring and connected components
Geospatial	Shapely / STRtree	Ward polygon resolution and fallback
Confidence	CHEF calculator	Weighted evidence aggregation
API	FastAPI	Ingestion, report, match and benchmark end-points
UI	React, Leaflet	Map-based report and event exploration

© 2026 IJSRET

single-pair matching endpoints. The React frontend includes a Leaflet-based Bengaluru map component. The service layer is designed to support modular API access and map-based exploration; systematic throughput and latency profiling is planned as part of deployment-scale evaluation.

V. PROPOSED METHODOLOGY

1. Administrative Taxonomy Crosswalk

The crosswalk registry indexes BBMP category/subcategory combinations and maps them to a LOCALYN event family and a BTP event cause. Each entry carries a mapping strength:

$$w_{map} = \begin{cases} 1.00, & \text{Direct} \\ 0.75, & \text{Probable} \\ 0.50, & \text{Ambiguous} \\ 0, & \text{Do not map.} \end{cases} \quad (1)$$

Only Direct, Probable, and Ambiguous entries are eligible for candidate generation. The CSV supplied in the archive contains 159 rows: 5 Direct, 8 Probable, 2 Ambiguous, and 144 Do not map. This count is taken from the data file rather than comments in the source code that refer to 161 entries. The crosswalk includes event families such as road damage, waterlogging, tree fall, and construction-related incidents.

The registry first attempts an exact category–subcategory lookup, then a subcategory lookup, substring fallback, and finally a category-level fallback. The fallback sequence improves coverage across variations in agency terminology while retaining the selected mapping and its strength for downstream inspection. Its individual contribution will be quantified through planned crosswalk ablation experiments.

TABLE VII BBMP–BTP Crosswalk DistributioN

Mapping strength	Entries	Share
Direct	5	3.14%
Probable	8	5.03%
Ambiguous	2	1.26%
Do not map	144	90.57%
Total	159	100.00%

TABLE VIII Pairwise Matching Feature Design

Feature	Weight	Comparison basis
Temporal	0.25	Difference between report timestamps
Spatial	0.35	Ward resolution and administrative compatibility
Semantic	0.25	Event-family/taxonomy compatibility
Text	0.15	Sentence-BERT cosine or Jaccard fallback
Total	1.00	Weighted pair score

Ward ResolutionThe geographic service loads the 225-feature Bengaluru wardGeoJSON file. The use of administrative geometry as a shared spatial reference is consistent with established geographic-data quality and interoperability practiceciteiso19157,ogcgeosparql2023. Invalid geometries are passed through Shapely's geometry repair function. Geometries are stored in an STRtree. For an input coordinate (ϕ, λ), the service first rejects invalid coordinates and points outside an expanded city bounding envelope. It then queries the spatial index and performs exact polygon containment.

The point-in-polygon decision is:

$$C_{centroid} = \max\left(0.4,\ 1 - \frac{d_H}{5}\right), \quad (3)$$

where Pi is a ward polygon, ci is its centroid, and dH is thehaversine distance. When multiple polygons contain or touch a point, the implementation selects the lowest ward identifier deterministically. Exact containment is assigned confidence 1.00; a boundary match is assigned 0.85. The nearest-centroidfallback reports

$$w(\phi, \lambda) = \begin{cases} w_i, & p \in P_i, \\ \min_{id}\{w_j : p \in \partial P_j\}, & p \text{ touches boundaries}, \\ \arg\min_w d_H(p, c_i), & d_H(p, c_i) \le 3.5 \text{ km}, \\ \emptyset, & \text{otherwise}, \end{cases} \quad (2)$$

rounded to two decimal places, and is permitted only within 3.5 km. The implementation also normalises ward names and maintains a legacy 198-to-225 ward correspondence map.

3. . Candidate Generation

A BBMP grievance b and a BTP observation t are eligible for scoring when the BBMP taxonomy entry exists, is eligible, maps to a compatible BTP event family, and has a ward-compatibility score of at least 0.5. Candidate eligibility is determined by taxonomy and spatial compatibility, while

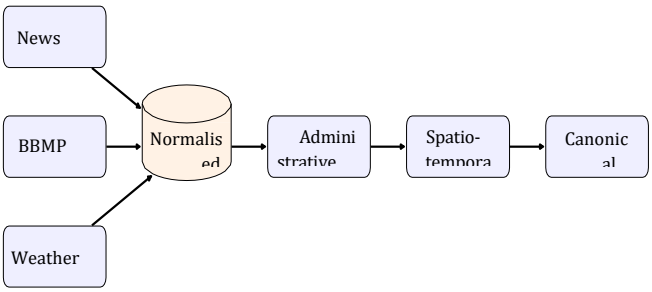


Fig. 2. Data-fusion pipeline. The live ingestion sources populate the general report store; the BBMP–BTP pair-resolution path is benchmarked separately in the current implementation.

© 2026 IJSRET

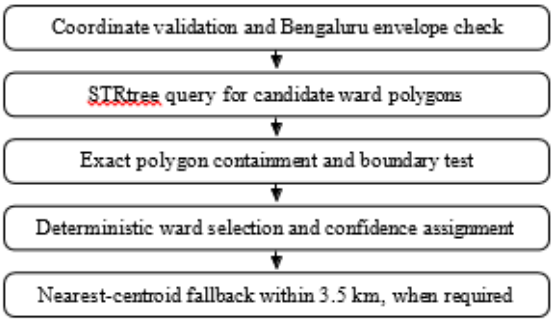


Fig. 3. Ward-resolution sequence implemented by the geographic service.

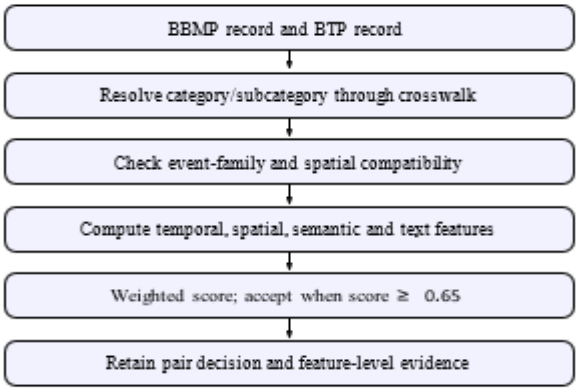


Fig. 4. Cross-agency event-association workflow.

temporal distance is represented as a graded feature in the subsequent composite score. This separation allows the matcher to retain potentially relevant records with wider temporal separation while reducing their final score.
Temporal Compatibility

Let Δt denote the absolute difference between parsed BBMP and BTP timestamps, in hours. The temporal feature is discretised as

$$T(\Delta t) = \begin{cases} 1.00, & 0 \le \Delta t \le 2, \\ 0.95, & 2 < \Delta t \le 6, \\ 0.90, & 6 < \Delta t \le 12, \\ 0.80, & 12 < \Delta t \le 24, \\ 0.50, & 24 < \Delta t \le 48, \\ 0.30, & 48 < \Delta t \le 72, \\ 0, & \Delta t > 72. \end{cases} \quad (4)$$

If timestamps cannot be parsed, the implementation assigns 0.40. Date-only strings on the same calendar date are treated as zero temporal distance.

5. Spatial Compatibility

The spatial feature S is calculated by the ward-compatibility routine. It uses coordinate-to-polygon resolution when BTP coordinates are available, then checks normalised ward-name equality, legacy ward correspondence, and assembly-constituency compatibility. The implementation returns fixed compatibility levels (including 1.00 for strict polygon contain-ment, 0.90 for boundary or legacy correspondence, 0.85 for a proximity/adjacency case, and 0 for mismatch), with textual reasons. This is an administrative compatibility score, not a calibrated probability of event identity.

6. Taxonomy-Semantic and Text Similarity

Let M denote the semantic feature. When a crosswalk entry is compatible with the BTP event cause, M equals its mapping-strength weight. When the entry is eligible but does not satisfy the family-compatibility test, the score is discounted to 0.6wmap; otherwise it is zero.

For text, LOCALYN concatenates BBMP description, subcat-egory, and staff remarks, and BTP description, title, and location name. When the sentence-transformer model is available, cosine similarity is rescaled to [0, 1]:

$$x_{\text{SBERT}} = \frac{1 + \cos(\mathbf{e}_b, \mathbf{e}_t)}{2}. \quad (5)$$

If the model is unavailable or inference fails, the system uses Jaccard similarity over filtered token sets:

$$x_{\text{Jac}} = \frac{|A \cap B|}{|A \cup B|}. \quad (6)$$

Algorithm 1 LOCALYN BBMP–BTP pair matching

Require: BBMP record b, BTP record t, crosswalk registry C
Ensure: Match decision and feature explanation

```
m ← C.lookup(b.category, b.subcategory)
if m is missing or ineligible then
    return non-match: taxonomy unavailable
end if
if m is incompatible with t.cause then
    return non-match: event-family mismatch
end if
(S, r_s) ← wardCompatibility(b, t)
if S < 0.5 then
    return non-match: spatial incompatibility
end if
T ← temporalCompatibility(b.time, t.time)
M ← crosswalkStrength(m)
X ← textSimilarity(b, t)
R ← 0.25T + 0.35S + 0.25M + 0.15X
ŷ ← (R ≥ 0.65)
return ŷ, R, (T, S, M, X), r_s
```

© 2026 IJSRET

The selected text feature is denoted by X.

7. Composite Match Score

For each eligible candidate pair, the composite score is:

$$R(b, t) = 0.25T + 0.35S + 0.25M + 0.15X. \quad (7)$$

A pair is accepted when:

$$(b, t) = I[R(b, t) \geq 0.65]. \quad (8)$$

The weights and threshold are fixed configuration values in the submitted implementation; no training or data-driven optimisation procedure is present. The system stores the individual feature values and human-readable explanations for each decision.

8. . Canonical Event Construction

Accepted pairs are represented as edges in an undirected graph. Record identifiers form vertices; an accepted BBMP–BTP pair creates an edge. A disjoint-set/union–find procedure groups connected records into components. Each component becomes a canonical event with a representative category, ward, title, earliest and latest timestamps, participating agencies, observation count, and averaged coordinates where available. The centroid of the observation coordinates is passed through ward resolution to assign a ward identifier when possible. This stage is connected-component consolidation, not density-based clustering.

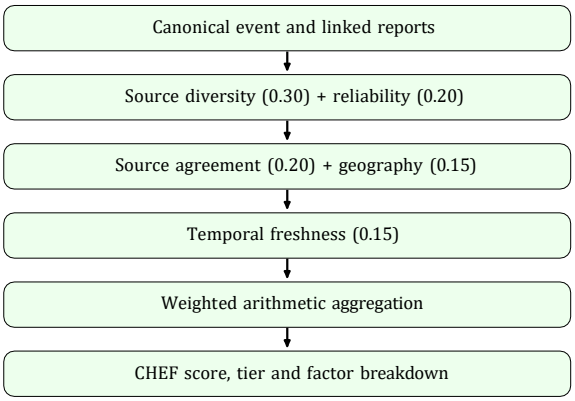


Fig. 5. CHEF calculation workflow. The implemented principal score is a deterministic weighted heuristic rather than a calibrated probability.

TABLE IX Chef Confidence Components

Component	Weight	Current operationalisation
Source diversity	0.30	Number of distinct agencies
Source reliability	0.20	Fixed agency reliability values
Source agreement	0.20	Agreement proxy based on agency count
Geographic consistency	0.15	Ward presence/consistency indicator
Temporal freshness	0.15	Recency factor
Total	1.00	Deterministic weighted heuristic

VI. CONFIDENCE ESTIMATION

LOCALYN contains two related deterministic confidence paths. The single-report path uses source reliability, temporal freshness, corroboration, source agreement, and geographic consistency with weights 0.30, 0.25, 0.25, 0.10, and 0.10, respectively. Source reliability is rule-based: weather APIs receive 0.95, listed established news outlets 0.85, other named sources 0.70, and unknown sources 0.50. Corroboration is assigned 0.50 for no other reports, 0.75 for one, and 0.95 for two or more. In this path, source agreement and geographic consistency are currently set to 1.0 rather than inferred from independent measurements.

The full canonical-event path uses:

$$C_{full} = 0.30D + 0.20Q + 0.20A + 0.15G + 0.15F, \quad (9)$$

where D is source diversity, Q is average agency reliability, A is source agreement, G is geographic consistency, and F is temporal freshness. The code assigns diversity values of 0.30, 0.50, 0.85, and 0.98 for zero, one, two, and at least three distinct agencies, respectively. Agency reliability constants include 0.90 for BTP, 0.85 for BBMP, 0.95 for weather observations, 0.80 for news, and 0.65 for citizen reports. Agreement is 0.95 when at least two distinct agencies participate and 0.70 otherwise. Geographic consistency is 1.00 when a ward name is present and 0.70 otherwise.

Both paths return a score clipped to [0, 1] and a tier: HIGH for scores at least 0.80, MODERATE for scores at least 0.60, and LOW otherwise. The score is an interpretable heuristic and has not been calibrated against event truth labels. Although

TABLE X Implemented Benchmark Results On Nine Definitive Pairs

Method	TP	FP	FN	F1
Date + ward	5	2	0	0.833
Date + ward + naive category	0	0	5	0.000
LOCALYN	5	0	0	1.000

the code contains a Dempster–Shafer combination helper, the reported confidence in the principal CHEF routines is based on weighted arithmetic aggregation. Accordingly, no claim of validated Dempster–Shafer fusion or source-ancestry/Sybil resistance is made.

© 2026 IJSRET

VII. EXPERIMENTAL SETUP

1. Implementation Environment

The backend uses Python, FastAPI and SQLite. spaCy supports extraction of locality and time expressions, Shapely handles ward geometries, and SentenceTransformers supplies the optional sentence-embedding model. The frontend is built with React and Vite, with Leaflet used for map display. The repository also includes tests for API routes, pair matching, ward resolution, weather processing, scheduling and confidence calculations. Runtime and resource measurements are planned for a larger deployment study.

2. Benchmark Dataset

The benchmark implementation contains ten BBMP–BTP record pairs: five positive, four negative, and one uncertain. The uncertain pair is excluded from definitive metric calculations. The examples cover direct pothole, waterlogging, and tree-fall correspondences as well as negative and ambiguous cases. The repository encodes benchmark labels and explanatory notes alongside the test cases. This benchmark provides a reproducible functional check of the matching logic; broader event coverage, independent annotation, and held-out evaluation are planned for subsequent study.

3. Baselines and Metrics

Three approaches are compared:

- **Baseline 1:** same calendar date and compatible ward.
- **Baseline 2**: same date, compatible ward, and naive raw category/subcategory string compatibility.
- **LOCALYN:** crosswalk eligibility and event-family com-patibility, ward point-in-polygon/administrative compati-bility, and the four-feature weighted score.

Precision, recall, and F1-score are calculated over positive and negative labels; uncertain pairs are skipped:

$$P = \frac{TP}{TP + FP}, \qquad R = \frac{TP}{TP + FN}, \tag{10}$$

$$F_1 = \frac{2PR}{P + R}. \tag{11}$$

TABLE XI Precision, Recall And F1 Comparison

Method	Precision	Recall	F1
Date + ward	0.714	1.000	0.833
Date + ward + naive category	0.000	0.000	0.000
LOCALYN	1.000	1.000	1.000

TABLE XII Recorded Benchmark Performance On Nine Definitive Pairs

Approach	TP	FP	FN	P	R	F1
Date + ward	5	2	0	0.714	1.000	0.833
Date + ward + naive category	0	0	5	0.000	0.000	0.000
LOCALYN	5	0	0	1.000	1.000	1.000

VIII. RESULTS AND DISCUSSION

Table X reports the benchmark values stored in the repository. LOCALYN identifies all five positive examples and produces no false positives among the four negative examples, yielding precision, recall, and F1-score of 1.000 on the definitive benchmark pairs. The date–ward baseline identifies all five positives but produces two false positives, resulting in an F1-score of 0.833. The naive category baseline misses all five positive examples, illustrating the value of an explicit semantic bridge when agencies use different vocabularies.

The comparison highlights the role of the crosswalk in this set of examples. The naive category rule relies on similar category strings, whereas the LOCALYN workflow compares mapped event families and retains the mapping strength as part of the score. The date–ward baseline uses two useful signals but does not distinguish between different types of incidents occurring in the same area. In the curated examples, combining taxonomy and spatial evidence gives the matcher a more specific basis for association. A broader dataset will be used to examine how this behaviour carries over to other wards and incident types.

Feature-level ablation is a planned next step to quantify the individual contributions of taxonomy mapping, ward compatibil-ity, temporal scoring, and text similarity. The optional Sentence-BERT/Jaccard design also enables a controlled comparison of semantic and lexical matching under a fixed runtime configuration.

1. Comparative Analysis

The naive category baseline performs poorly because agency-specific category vocabularies are not necessarily lexically aligned. The crosswalk provides a domain-specific bridge between BBMP grievance descriptions and BTP event causes. The spatial subsystem provides an additional administrative constraint: a BTP coordinate can be resolved to a ward polygon rather than relying only on a locality string. Finally, the composite score retains temporal and textual signals as graded

evidence instead of treating any single feature as decisive.

© 2026 IJSRET

2. Operational and Scientific Implications

The system keeps the reasons for a pair decision alongside the score. An operator can inspect whether the records were close in time, compatible at ward level, linked by the taxonomy and similar in their text. The same principle is used by CHEF: its factors are visible rather than hidden inside a single output value. This is useful when civic information is reviewed by people who need to understand why reports have been grouped together.

TABLE XIII Planned Ablation Matrix

Configuration	Status	Purpose
Full LOCALYN matcher	Evaluated	Reference configuration
Without taxonomy gate	Planned	Measure crosswalk contribution
Without spatial feature	Planned	Measure ward/spatial contribution
Without temporal feature	Planned	Measure time contribution
Without text feature	Planned	Measure semantic-text contribution
SBERT vs. Jaccard only	Planned	Compare text backends

1. ADVANTAGES

The current implementation offers several practical benefits:

- Administrative interoperability: explicit mapping-strength labels preserve uncertainty in cross-agency category correspondences.
- Ward-aware geospatial resolution: the 225-feature ward dataset, STRtree candidate lookup, deterministic bound-ary policy, and centroid fallback support locality-level resolution.
- Interpretable matching: the four component scores and explanatory messages make individual match decisions inspectable.
- Graceful text-similarity fallback: the matcher can use Jaccard similarity if the optional sentence-transformer model is unavailable.
- Canonical event representation: accepted cross-agency links can be consolidated into an event record with temporal bounds, participating agencies, and observation provenance.

2. Implementation Scope And Future Development

LOCALYN currently combines a live public-source ingestion and civic-report management path with a taxonomy-aware BBMP–BTP matching and benchmark path. The present implementation includes the BBMP–BTP crosswalk, candidate generator, pair scorer, ward service, connected-component grouping and CHEF calculation. Its benchmark contains ten curated pairs, with nine definitive labels used for the reported comparison.

CHEF is a deterministic evidence score based on configured factors. Candidate eligibility is determined by taxonomy and spatial compatibility, while temporal proximity contributes to the final pair score. These modules provide a working base for adding further agency data and extending the evaluation.

- Run ablations: remove taxonomy mapping, spatial compatibility, temporal compatibility, SBERT, and Jaccard features one at a time. Compare full and reduced models.
- Tune and validate thresholds: perform grid search only on training/validation data; report precision–recall curves, ROC-AUC where appropriate, and confidence intervals from bootstrap resampling.
- Compare text models: evaluate SBERT and Jaccard on the same fixed environment and annotated set; report model version, inference latency, and fallback frequency.
- Calibrate CHEF: compare confidence against observed match correctness using reliability diagrams, Brier score, expected calibration error, and calibration methods such as isotonic regression.
- Test source dependence: identify syndicated or copied reports and compare naive report counts with unique-source or shared-ancestry aggregation. This is not currently implemented.
- Measure geographic accuracy: validate ward assignment against authoritative labelled coordinates, especially near boundaries and in centroid-fallback cases. Report accuracy, ambiguity rate, and distance-to-boundary sen-sitivity.
- Evaluate operational performance: measure candidate-generation reduction, throughput, p50/p95 latency, mem-ory use, and behaviour as the number of reports grows.
- Integrate live BTP data: connect a documented, legally accessible BTP data source to the same ingestion and per-sistence pipeline, with source provenance and timestamp semantics recorded.

IX. CONCLUSION

LOCALYN brings together several steps that are usually handled separately when civic reports are collected from different organisations. The implemented platform includes public-source ingestion, a BBMP–BTP taxonomy crosswalk, ward-level geographic resolution, temporal and text comparison, canonical-event grouping and the CHEF evidence score. Its application layer stores reports in SQLite, exposes operations through FastAPI and presents geographic information through a React–Leaflet interface.

© 2026 IJSRET

The matching workflow combines four signals rather than depending on a single field. The crosswalk relates agency-specific incident labels, the ward service provides a common

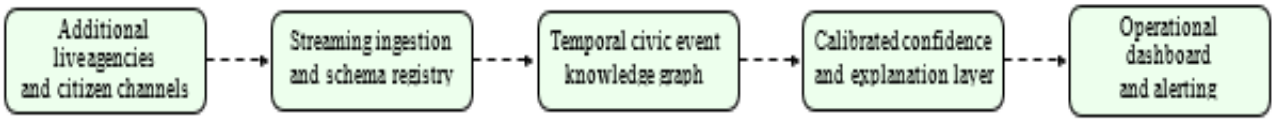


Fig. 6. Proposed extension architecture. Dashed links indicate planned components, not features claimed as implemented

TABLE XIV Implementation Status Of Major Components

Component	Status	Evidence / next step
News/BBMP/weather ingestion	Implemented	Connectors and persistence present
BBMP–BTP taxonomy cross-walk	Implemented	159-row CSV registry
Ward resolver	Implemented	225 ward polygons and spatial index
Pairwise matcher	Implemented	Weighted features and threshold 0.65
Canonical event grouping	Implemented	Union–find connected compo-nents
CHEF score	Implemented	Weighted deterministic heuristic
Live BTP ingestion	Future inte-gration	Benchmark path currently sepa-rate
Large-scale validation	Planned	Expand and annotate gold set
Confidence calibration	Planned	Reliability and calibration study

administrative location reference, and the text and time features

add detail to each pair decision. In the current benchmark, LOCALYN records an F1-score of 1.00 across nine definitive pairs, compared with 0.833 for date–ward matching and 0.000 for the date–ward–category baseline. This initial comparison shows how the implemented components work together on the selected examples. The next stage is to connect further agency feeds, extend the annotated examples and study system behaviour with a larger volume of reports. The modular design provides a basis for those extensions while retaining inspectable matching and confidence outputs.

REFERENCES

1. F. Atefeh and W. Khreich, “A survey of techniques for event detection in Twitter,” Computational Intelligence, vol. 31, no. 1, pp. 132–164, 2015, doi: 10.1111/coin.12017.
2. P. Christen, “A survey of indexing techniques for scalable record linkage and deduplication,” IEEE Trans. Knowl. Data Eng., vol. 24, no. 9, pp. 1537–1555, 2012, doi: 10.1109/TKDE.2011.127.
3. M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in Proc. KDD, 1996, pp. 226–231.
4. N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in Proc. EMNLP-IJCNLP, 2019, pp. 3982–3992, doi: 10.18653/v1/D19-1410.
5. A. P. Dempster, “Upper and lower probabilities induced by a multivalued mapping,” Ann. Math. Statist., vol. 38, no. 2, pp. 325–339, 1967, doi: 10.1214/aoms/1177698950. cities: Emerging trends and persistent challenges,” Information Fusion, vol. 114, Art. no. 102667, 2025, doi: 10.1016/j.inffus.2024.102667.
6. G. Shafer, A Mathematical Theory of Evidence. Princeton, NJ, USA: Princeton Univ. Press, 1976.
7. M. Imran, C. Castillo, F. Diaz, and S. Vieweg, “Processing social mediamessages in mass emergency: A survey,” ACM Comput. Surv., vol. 47, no. 4, Art. no. 67, 2015, doi: 10.1145/2700390..
8. M. Batty et al., “Smart cities of the future,” Eur. Phys. J. Special Topics, vol. 214, pp. 481–518, 2012, doi: 10.1140/epjst/e2012-01703-3.
9. M. F. Goodchild, “Citizens as sensors: The world of volunteered geography,” GeoJournal, vol. 69, pp. 211–221, 2007, doi: 10.1007/s10708-007-9111-y.
10. J. Bleiholder and F. Naumann, “Data fusion,” ACM Comput. Surv., vol. 41, no. 1, Art. no. 1, 2009, doi: 10.1145/1456650.1456651.
11. Y. Zheng, L. Capra, O. Wolfson, and H. Yang, “Urban computing:
Concepts, methodologies, and applications,” ACM Trans. Intell. Syst. Technol., vol. 5, no. 3, Art. no. 38, 2014, doi: 10.1145/2629592.
12. M. Mondal, M. Khayati, H.-A. Sandlin, et al., “A survey of multimodalevent detection based on data fusion,” VLDB J., vol. 34, Art. no. 9, 2025, doi: 10.1007/s00778-024-00878-5.
13. H. Orchi, A. B. Diallo, H. Elbiaze, E. Sabir, and M. Sadik, “A contem-porary survey on multisource information fusion for smart sustainable

© 2026 IJSRET

cities: Emerging trends and persistent challenges," Information Fusion, vol. 114, Art. no. 102667, 2025, doi: 10.1016/j.inffus.2024.102667
14. X. Zou et al., "Deep learning for cross-domain data fusion in urban computing: Taxonomy, advances, and outlook," Information Fusion, vol. 113, Art. no. 102606, 2025, doi: 10.1016/j.inffus.2024.102606.
15. I. Jeddouba, G.-A. Nysa, R. Hajji, and R. Billen, "Data integration across urban digital twin lifecycle: A comprehensive review of current initiatives," Annals of GIS, vol. 31, no. 3, pp. 367–386, 2025, doi: 10.1080/19475683.2024.2416135.
16. I. P. Fellegi and A. B. Sunter, "A theory for record linkage," J. Amer. Stat. Assoc., vol. 64, no. 328, pp. 1183–1210, 1969, doi: 10.1080/01621459.1969.10501049.
17. C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in Proc. ICML, vol. 70, 2017, pp. 1321–1330.
18. R. Kitchin, "The real-time city? Big data and smart urbanism," GeoJour-nal, vol. 79, pp. 1–14, 2014, doi: 10.1007/s10708-013-9516-8.
19. International Organization for Standardization, Geographic Information—Data Quality—Part 1: General Requirements, ISO 19157-1:2023, 2023.
20. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in Proc. ACM SIGKDD, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
21. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, 2017.
22. W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, "MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers," in Advances in Neural Information Processing Systems, vol. 33, 2020.
23. G. W. Brier, "Verification of forecasts expressed in terms of probability," Monthly Weather Review, vol. 78, no. 1, pp. 1–3, 1950, doi: 10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2.
24. Open Geospatial Consortium, "OGC GeoSPARQL 1.1," OGC Standard 22-047r1, 2023.

© 2026 IJSRET