

AUDITABLE FORENSIC GEOMETRY

TrajectoryLock

A multi-source system for reconstructing projectile trajectories and testing their geometric compatibility with official accounts

WHITEPAPER + WORKING REFERENCE IMPLEMENTATION

Concept architecture, uncertainty model, chain-of-custody design, source-independence controls, API contract, validation plan, and deployable Python MVP.

Version 0.1 / September 2026

Research prototype. Not a certified forensic instrument and not a substitute for qualified scene reconstruction, medical findings, or laboratory examination.

1. Executive summary

TrajectoryLock is designed to answer one bounded question: **given the available geometry and its uncertainty, how compatible is a reconstructed trajectory with a declared official trajectory?** It accepts multiple videos, photographs, surveyed points, direct trajectory measurements, witness bearings, and an official narrative expressed as a testable 3D line with tolerances.

The system does not assign a single magical confidence number. It separates three quantities that are routinely conflated:

- **Compatibility score:** continuous geometric closeness between reconstructed and declared lines.
- **Threshold match probability:** Monte Carlo probability that angle and offset fall within declared tolerances.
- **Reconstruction confidence:** strength, calibration, and independence of the evidence supporting the estimated line.

A high match probability with low reconstruction confidence is weak evidence. A high-confidence reconstruction with a low match probability is evidence of geometric inconsistency. Intermediate values remain indeterminate.

Core design principle

More files improve precision only when they contribute independent constraints. Re-encodes, reposts, shared camera origins, copied witness accounts, and derivative edits are grouped and down-weighted. This prevents a hundred copies of one clip from masquerading as a hundred independent viewpoints.

2. Problem definition and scope

Public incident review is often trapped between two weak extremes: qualitative visual argument and opaque model output. TrajectoryLock creates a reproducible middle layer. Every conclusion is conditional on declared coordinate systems, calibrations, annotations, source relationships, and uncertainty assumptions.

In scope	Out of scope
Bullets, arrows, fragments, thrown objects, or other approximately linear flight/impact segments	Identifying a shooter, intent, guilt, institutional honesty, or a complete event narrative
3D line reconstruction from calibrated rays, scene survey points, and direct measurements	Unvalidated automatic detection of invisible or motion-blurred projectiles
Geometric agreement testing with transparent uncertainty and tolerances	Replacing firearm/toolmark, autopsy, wound-path, acoustic, or materials experts
Witness direction as labeled corroboration	Treating memory as equal to surveyed physical evidence

Threat model

- Accidental coordinate-frame mismatch, unit mismatch, or camera-calibration error.
- Duplicate/correlated sources that falsely inflate certainty.
- Selection bias: only favorable frames, witnesses, or official parameters are entered.
- Edited media, missing provenance, synthetic content, and timestamp drift.
- Analyst tuning of tolerances after seeing the result.

The production design counters these risks with schema validation, file hashing, immutable revision history, blinded annotations, frozen tolerance declarations, independence groups, sensitivity analysis, and signed result manifests.

3. Evidence and coordinate model

All physical geometry is represented in a right-handed Cartesian world coordinate system measured in metres. Camera calibration converts a marked pixel into a world-space sight ray. Multiple separated sight rays that refer to the same physical feature are intersected by weighted least squares.

Evidence type	Required fields	Role
Visual ray	feature ID, camera origin, direction, angular sigma, range hint	Triangulates a physical point from calibrated video/photo annotations
Survey point	feature ID, XYZ point, positional sigma	Directly constrains entry, exit, strike, or projectile position
Direct line	XYZ point, direction, angular sigma, offset sigma	Represents a validated photogrammetric or scene-reconstruction line
Witness bearing	origin, direction, angular sigma, reliability	Corroboration only; does not silently control physical reconstruction
Official hypothesis	point, direction, uncertainty, frozen angle/offset tolerances	The testable claim being compared

A visual file is not itself a trajectory measurement. It becomes one only after calibration, feature identification, frame selection, annotation, and uncertainty assignment. Each transformation should retain its operator, algorithm version, timestamp, and parent artifact hash.

Feature examples

Entry and exit points are common, but the schema also supports named positions such as `projectile_t0`, `projectile_t1`, ricochet points, embedded fragments, arrow shaft endpoints, or line-of-damage marks. At least two reconstructed points or one qualified direct line are required.

4. Reconstruction mathematics

For sight ray i with origin $\mathbf{o}_i$, unit direction $\mathbf{d}_i$, and weight w_i , the matrix $\mathbf{P}_i = \mathbf{I} - \mathbf{d}_i \mathbf{d}_i^T$ projects displacement into the plane perpendicular to the ray. The weighted least-squares point is:

$$\mathbf{A} = \sum_i w_i \mathbf{P}_i \quad \mathbf{b} = \sum_i w_i \mathbf{P}_i \mathbf{o}_i \quad \mathbf{x}_{\text{hat}} = \mathbf{A}^{-1} \mathbf{b}$$

Angular uncertainty is converted to a transverse uncertainty using a declared range hint. Surveyed points add isotropic constraints. The solver rejects ill-conditioned geometry, including parallel views that cannot determine depth.

Line fit

With at least two reconstructed features, weighted principal-component analysis estimates the dominant line direction. Feature uncertainty supplies weights. Direct-line estimates are sign-aligned and blended by inverse angular variance. Because a geometric line has no intrinsic polarity, comparison uses the acute angle between directions.

Line-to-line offset

For lines $(p1, d1)$ and $(p2, d2)$, the shortest skew-line distance is the absolute projection of $(p2 - p1)$ onto the normalized cross product of $d1$ and $d2$. A stable perpendicular-distance form is used when lines are nearly parallel.

Uncertainty propagation

Monte Carlo trials perturb reconstructed and official directions by their angular standard deviations and line origins by positional standard deviations. Every trial records angular difference and line offset. The output includes 95% empirical intervals and the proportion satisfying both frozen thresholds.

$$P(\text{match} \mid \text{declared model}) = \text{count}(\text{angle} \leq T_{\text{angle}} \text{ AND } \text{offset} \leq T_{\text{offset}}) / N$$

This is conditional model agreement, not the probability that the official account is true.

5. Confidence and source independence

TrajectoryLock prevents file count from becoming confidence by using independence groups. Each group contributes its strongest source fully; additional correlated items in the same group add only a small capped increment.

$$\text{effective_group} = \text{best_quality} + 0.2 * (1 - \exp(-\text{sum}(\text{copy_qualities})))$$

The prototype reconstruction-confidence score combines effective source strength (55%), number of reconstructed features (30%), and calibration completeness (15%). It is bounded below 100% to avoid representing model confidence as certainty.

Situation	Expected effect
Second calibrated camera at a separated position	Material gain: adds depth geometry and an independent constraint
A repost or re-encode of camera A	Minimal gain: same independence group
Independent scene survey of entry and exit points	Large gain: different measurement modality
Multiple witnesses who discussed the event together	Grouped or correlation-adjusted; not independent by default
Many frames from one fixed camera	Can reduce annotation noise, but does not create new viewpoint geometry

Calibration is scoreable only when the calibration artifact is preserved and linked. A boolean flag is sufficient for the MVP; production should evaluate reprojection error, control-point coverage, lens model, rolling-shutter behavior, and temporal synchronization.

6. System architecture

The reference implementation is intentionally dependency-light. A NumPy computational core is wrapped by a command-line interface and a local HTTP workbench. The repository includes a browser UI, example case, hashing utilities, Dockerfile, and unit tests.

Layer	Responsibility	Reference module
Integrity	SHA-256 media hashes and canonical JSON fingerprints	integrity.py
Geometry	Vector validation, ray triangulation, PCA line fitting, offsets	geometry.py
Scoring	Monte Carlo comparison, independence score, witness agreement	scoring.py
Pipeline	Schema checks, feature assembly, result manifest	pipeline.py
Interfaces	CLI, JSON HTTP API, local browser workbench	cli.py / server.py
Verification	Synthetic consistent/inconsistent cases and edge tests	tests/

Data flow

```
Original media -> immutable hash -> calibration -> frame/feature annotation -> world-space  
observation -> feature triangulation -> 3D trajectory -> Monte Carlo comparison -> signed result  
manifest
```

Production components

- Encrypted object storage for originals and derivatives; write-once audit event store.
- PostgreSQL case/evidence database with row-level authorization and revision history.
- Worker queue for frame extraction, calibration, photogrammetry, and simulation.
- Role separation for evidence custodian, annotator, reviewer, and report signer.
- Reproducible containers pinned by digest; deterministic seeds and dependency lockfiles.

7. API and case contract

The MVP exposes a small JSON API. POST /api/analyze accepts the same object used by the CLI and returns a complete result manifest. GET /api/example supplies a synthetic case. POST /api/hash hashes local paths when the server is explicitly run in a trusted local environment.

Minimal case skeleton

```
{
  "case_id": "CASE-001",
  "sources": [{ "id": "survey-a", "quality": 0.95,
    "calibrated": true, "independence_group": "survey-a" }],
  "observations": [{ "type": "direct_line", "source_id": "survey-a",
    "point": [0, 0, 1.2], "direction": [1, 0.1, 0.02],
    "angular_sigma_deg": 0.5, "offset_sigma_m": 0.02 }],
  "official_hypothesis": { "point": [0, 0, 1.2], "direction": [1, 0.1, 0.02],
    "angular_sigma_deg": 0.7, "offset_sigma_m": 0.04,
    "angle_tolerance_deg": 3.0, "offset_tolerance_m": 0.25 }
}
```

Result contract

- Input and result SHA-256 fingerprints for reproducibility.
- Reconstructed line point, unit direction, angular/offset uncertainty, and residual.
- Per-feature position, uncertainty, ray residual, condition number, and observation count.
- Compatibility, threshold probability, conclusion class, confidence, and witness corroboration.
- Warnings and a mandatory interpretive guardrail.

8. Workflow and human controls

Stage	Required control	Failure response
Acquire	Preserve original bytes, metadata, custody event, and hash	Quarantine unexplained modifications
Synchronize	Estimate clock offsets and rolling-shutter/timing uncertainty	Widen uncertainty or exclude timing claim
Calibrate	Lens model and surveyed control points; retain reprojection residuals	Reject or downgrade bad calibration
Annotate	Blind or dual annotation; label visibility and ambiguity	Adjudicate; keep both originals
Declare	Freeze official line and tolerances before computing comparison	Mark post-hoc changes as new hypothesis revision
Analyze	Pinned code, seed, source groups, and sensitivity runs	Block result if validation fails
Review	Independent technical reviewer and signed interpretation	Return for correction without erasing history

Witness handling

Witness accounts are encoded as distributions, not exact rays. Interview timing, opportunity to observe, contamination, confidence, and directional ambiguity should be retained. Witness agreement is reported separately from physical-line reconstruction so a crowd cannot numerically overrule calibrated geometry.

Official narrative handling

An official account must be decomposed into a falsifiable geometric hypothesis. Vague prose such as 'from the north' should not be converted into a precise line without documenting the conversion. If several official versions exist, they become separately versioned hypotheses and are never merged into a moving target.

9. Validation plan

The included tests verify the mathematics and major score guardrails, but a forensic claim requires empirical validation on controlled physical scenes.

Validation tier	Dataset	Pass criteria
Unit	Synthetic exact and noisy geometry	Known intersections, distances, angles, deterministic results
Bench	Instrumented indoor paths with surveyed ground truth	Bias and 95% coverage within declared tolerances
Field	Multiple cameras, real lenses, lighting, occlusion, rolling shutter	Error curves stratified by range, baseline, frame rate, and visibility
Inter-analyst	Blinded repeated annotations by independent operators	Reported repeatability and adjudication rate
Adversarial	Edited, duplicated, synthetic, mislabeled, and coordinate-swapped inputs	Correct quarantine or uncertainty inflation
External	Independent laboratory executes frozen protocol and code	Reproducible estimates and calibrated confidence intervals

Required performance reporting

- Angular and lateral error distributions, not only mean error.
- Interval coverage: whether nominal 95% intervals contain ground truth about 95% of the time.
- False-consistency and false-inconsistency rates across pre-registered tolerances.
- Sensitivity to camera baseline, calibration residual, range, annotation error, and source correlation.
- Failure rate and explicit abstention rate; a trustworthy system must be able to return indeterminate.

10. Security, ethics, and evidentiary limits

Trajectory analysis may involve deaths, injuries, criminal investigations, private locations, and identifiable people. Production use requires strict minimization, access control, retention limits, and legal authority. Public transparency is not a license to expose victims, witnesses, or raw location data.

Required safeguards

- Encryption in transit and at rest; case-scoped least privilege and multi-factor authentication.
- Immutable event log for ingest, access, export, annotation, model run, and hypothesis revision.
- Cryptographic hashes for every original and derivative; explicit parent-child provenance.
- No face recognition or identity inference in the core trajectory workflow.
- Human-readable export containing assumptions, exclusions, uncertainty, version, and reviewer sign-off.
- No automated legal conclusion, accusation, or credibility classification.

Admissibility

The software alone does not establish admissibility or scientific acceptance. A competent expert must be able to explain the method, its validation, error rate, calibration, chain of custody, and case-specific limitations. Jurisdiction-specific rules and disclosure duties must be evaluated independently.

Interpretation matrix

Reconstruction confidence	Match probability	Permitted statement
High	High	Geometry is consistent within declared tolerances
High	Low	Geometry is inconsistent within declared tolerances
Low	Any	Evidence is insufficient for a strong geometric conclusion
Any	Middle band	Indeterminate; acquire better constraints or retain ambiguity

11. Development roadmap

Phase	Deliverable	Exit criterion
0 - Reference MVP	Auditable JSON core, web workbench, CLI, hashing, tests, Docker	Delivered with this whitepaper
1 - Evidence workstation	Media ingest, frame extraction, calibration UI, point/ray annotation, project database	End-to-end controlled case without manual JSON
2 - Photogrammetry	Bundle adjustment, surveyed scene model, camera pose solver, synchronization	Validated against instrumented ground truth
3 - Review and reporting	Dual annotation, revision ledger, sensitivity dashboard, signed PDF/JSON report	Independent reviewer can reproduce result
4 - Forensic validation	Pre-registered study, external replication, published error model	Defined intended-use claims are empirically supported
5 - Operational hardening	RBAC, encrypted storage, deployment monitoring, retention/export controls	Security and legal readiness review passed

Near-term engineering priorities

- Add OpenCV-based intrinsic/extrinsic calibration and lens-distortion models as an optional worker.
- Create an annotation UI with epipolar guides and blind duplicate labeling.
- Represent full covariance matrices rather than scalar angular/offset sigmas.
- Add robust estimators (RANSAC/M-estimators), outlier reports, and leave-one-source-out analysis.
- Add coordinate-frame transformations, unit registry, temporal alignment, and projectile physics only when supported by case data.
- Pre-register confidence calibration on controlled datasets before changing score semantics.

12. Reproducibility and operation

The accompanying source bundle runs on Python 3.10+ with NumPy. It can be executed directly, installed as a package, or started in Docker.

Commands

```
python -m trajectorylock.cli demo -o result.json
python -m trajectorylock.cli analyze examples/example_case.json -o result.json
python -m trajectorylock.cli hash-media video.mp4 photo.jpg
python -m trajectorylock.server --port 8080
python -m unittest discover -s tests -v
```

Reproducibility checklist

- Preserve the exact case JSON and every media/calibration hash.
- Record code version, environment, random seed, and Monte Carlo sample count.
- Do not alter official tolerances after reviewing the output without creating a new hypothesis revision.
- Repeat with leave-one-source-out and plausible uncertainty ranges.
- Report warnings, excluded evidence, degenerate geometry, and indeterminate outcomes.

Conclusion

TrajectoryLock converts a potentially rhetorical dispute into an inspectable geometric test. Its value lies less in producing a percentage than in forcing every percentage to reveal what generated it: independent sources, calibrated geometry, declared uncertainty, frozen thresholds, code version, and reproducible calculations.

Reference implementation: TrajectoryLock v0.1. Synthetic example results must never be represented as real-case findings.