# Tyche Atlas

Evidence Carriers, Corpus Scale, and Claim-Support Infrastructure for AI Governance Research

2026-06-07


## Abstract

Tyche Atlas is a thesis-scale evidence-navigation infrastructure for AI governance research. The Atlas organizes heterogeneous evidence carriers into a versioned research object: legal sources, standards and specifications, regulatory sandboxes, public-sector AI registers, public-sector agent materials, trust infrastructure, evidence-package fields, patent and post-quantum cryptography signals, publication-route controls, and claim-support gates.

The public portal exposes the human navigation layer, article PDF, and manifest. Internal support packages remain controlled release artifacts. The public value of the Atlas comes from showing the shape of the research programme: evidence origins, claim strength, repeated-observation control, and context-only signal families.


## Observatory Topics

Tyche Atlas is organized around observatory topics that can support a publication spine, a PhD research map, and later focused papers.

1. AI law and governance duties.
2. Standards and trust infrastructure.
3. Regulatory sandboxes as evidence-producing settings.
4. Public-sector AI agents, registers, and authorities.
5. Evidence packages and action receipts.
6. PKI, identity, credentials, and trust anchors.
7. Patents and post-quantum cryptography implementation signals.
8. Software supply-chain provenance and transparency logs.
9. Machine-readable governance and claim-support status.
10. Publication routing, limitations, and reproducibility boundaries.

## Theoretical Background from Atlas

Atlas uses three theoretical units: evidence carriers, claim gates, and provenance boundaries.

Evidence carriers are records or artifacts that can carry governance-relevant facts across legal, technical, institutional, and publication contexts. Legal sources, standards, registers, sandbox materials, evidence-package fields, patents, and trust-infrastructure records all become evidence carriers when they receive source identity, family, provenance, validation, and claim-link fields.

Claim gates assign manuscript claims to use states. The current Atlas support surfaces use `use-now-cautious`, `candidate-after-verification`, and `appendix-or-separate-paper-only` states. These states connect paper language to support strength, verification needs, and reviewer-facing caution.

Provenance boundaries connect observed source families, canonical records, duplicate clusters, checksums, package manifests, and public release surfaces. This boundary model lets the paper cite corpus structure and support status while raw histories and support packages remain controlled release artifacts.

The theoretical background therefore comes from the Atlas itself: codebooks define record fields, manifests define version boundaries, claim-verification tables define claim status, and evidence-family matrices define the route from source record to manuscript claim.

## Atlas Claim-Support Matrix

The matrix below maps major paper claims to Atlas support surfaces.

| Paper claim | Atlas support surface | Evidence basis | Status |
|---|---|---|---|
| Atlas as research object | Claim-support summary; corpus codebook summary; public manifest v1.9 | Versioning, provenance, deduplication, validation, inclusion rules, checksums, claim review | use-now-cautious |
| Action receipts as research construct | Compact claim-verification table rank 3; evidence-package profile schema | Actor identity, system state, evidence object, trust anchor, verifier-readable event record | use-now-cautious |
| Law-to-evidence crosswalk | Compact claim-verification table rank 2; canonical EU-law source records | AI Act, eIDAS/EUDI, GDPR, NIS2, Cyber Resilience Act, Data Act, Product Liability sources | candidate-after-verification |
| Standards evidence fragments | Compact claim-verification table rank 4; standards hook matrices | SCITT, RATS, C2PA, VC/DID, SBOM, provenance, signing, transparency families | candidate-after-verification |
| Sandboxes as evidence-producing settings | Compact claim-verification table rank 5; sandbox support summaries | AI Act Articles 57-59, AESIA comparator materials, sandbox procedures, logs, documentation surfaces | candidate-after-verification |

| Paper claim | Atlas support surface | Evidence basis | Status |
|---|---|---|---|
| Public registers as schema analogues | Compact claim-verification table rank 6; codebook field families | Public-summary fields, institutional surfaces, authority and validation signals | candidate-after-verification |
| Patent/PQC landscape | Compact claim-verification table rank 10; patent/PQC signal slice | PKI, identity, cryptographic durability, provenance, AI-accountability implementation signals | appendix-or-separate-paper-only |
| Count discipline | Corpus codebook summary; duplicate-control manifest fields; public manifest v1.9 | Raw observations, canonical records, duplicate clusters, provenance retention | use-now-cautious |

## Corpus Scale

The current Atlas validation matrix records the following scale indicators.

| Layer | Count | Interpretation |
|---|---|---|
| Raw observations | 424,592 | Crawl, package, register, standards, legal, paper, and patent observations retained as provenance. |
| Canonical records | 162,994 | Deduplicated navigation records used for human-facing evidence cards. |
| Duplicate clusters | 16,070 | Repeat observations retained as provenance; evidence weight uses canonical records and claim gates. |
| Curated external records | 245,457 | Coverage-audit subset from curated external source families. |
| Patent BigQuery records | 167,601 | Patent landscape and implementation-signal context. |

These counts are corpus-engineering facts. The Atlas separates raw collection scale from canonical evidence, source authority, claim status, and limitations.

## Evidence Families

The Atlas connects several research families that are often discussed separately.

### AI Law and Governance Duties

The legal layer includes AI Act, eIDAS/EUDI, GDPR, NIS2, Cyber Resilience Act, Data Act, Product Liability, and adjacent official-source surfaces. These records support candidate law-to-evidence crosswalks. Broad legal-duty conclusions require exact clause-level verification.

Safe wording: inspected legal materials can be represented as candidate links to evidence objects; legal-duty wording remains provisional until clause-level verification is complete.

### Standards and Trust Infrastructure

The standards layer includes evidence fragments from SCITT, RATS, COSE/JOSE, verifiable credentials, DID/Data Integrity, C2PA, SPDX, CycloneDX, SLSA, Sigstore, in-toto, SBOM and provenance ecosystems. The Atlas treats these as partial functions for identity, attestation, provenance, signing, supply-chain metadata, and transparency.

Safe wording: the inspected standards and specification families provide partial evidence functions. A stronger common-carrier claim requires exact body, version, and section anchors.

### Sandboxes

The sandbox layer treats regulatory sandboxes as evidence-producing research settings. Sandbox materials can expose procedures, logs, technical documentation, public summaries, and authority interfaces.

Safe wording: AI regulatory sandboxes can be studied as evidence-producing settings for governance research. Claims about legal duties, national implementation, or official designation require clause and official-source verification.

### Public-Sector Registers and Authorities

The public-sector layer includes algorithm registers, supervision bodies, national digital authorities, public datasets, public-sector agent materials, and digital-government evidence surfaces. Public registers can be studied as schema analogues for public accountability and governance documentation.

Safe wording: public algorithm registers are early public-facing evidence-schema analogues. Verifier-complete lifecycle evidence requires additional records.

### Evidence Packages and Action Receipts

The evidence-package layer models action receipts: signed, verifier-readable event records linking actor identity, system state, evidence objects, and trust anchors. This is a research construct for governance design and evidence review.

Safe wording: action receipts are proposed verifier-readable records for research and design. Mandatory-carrier and standards-conformance language requires separate authority anchors.

## Patents, PQC, and Implementation Signals

The Atlas contains a substantial patent and post-quantum cryptography signal layer. This includes post-quantum cryptography, PKI trust structures, evidence durability, decentralized identity, device certificates, software supply-chain provenance, transparency logs, data provenance, AI accountability, and privacy-preserving compute.

Patent records show implementation patterns around trust, provenance, cryptography, identity, and AI accountability. Deployment, legal duty, market adoption, compliance, certification, and field effectiveness require separate validation.

| Signal family | Examples | Safe interpretation |
|---|---|---|
| PQC and quantum-resilient trust | post-quantum cryptography, quantum-key communication, quantum threat detection, smartcard PQC | landscape signal for trust-infrastructure research |
| PKI, identity, credentials | eIDAS/EUDI, verifiable credentials, decentralized identity, device certificates | partial evidence-carrier families |
| AI accountability | AI Act evidence, model documentation, algorithmic accountability, public registers | candidate governance evidence map |
| Software supply chain | SBOM/provenance, transparency logs, certificates of authenticity, data provenance | reusable evidence fragments |

## Claim-Support Gate

The controlling claim-support posture is conservative. Atlas claims have different publication states.

| Claim state | Examples | Public posture |
|---|---|---|
| Use now, cautiously | Atlas as citable research object; action receipts as a proposed research construct | Supported with boundaries. |
| Candidate after verification | Law-to-evidence crosswalks; standards fragments; sandbox evidence infrastructure; public-register schema analogies | Useful, with exact verification required for stronger claims. |
| Appendix or separate paper | Pallas-style country readiness; patents as implementation signals | Valuable as adjacent or separate-paper material. |

The claim gate protects the publication spine by separating legal compliance, certification, official audit, completeness, and field-effect claims from corpus-method claims.

## Duplicate Resolution

The Atlas raw corpus contains repeated crawl and export observations. Public counts foreground raw scale, canonical records, and duplicate clusters. Detailed repeat counts remain in audit records.

The largest duplicate clusters include major legal, standards, and public authority sources. Stable official sources are repeatedly observed across crawl, validation, and package cycles. Repeated sightings are treated as provenance.

## Publication Spine

The Atlas article route is a data/method and evidence-infrastructure route. The paper explains the research programme and the evidence carriers that connect its parts. The publication object is the infrastructure: corpus design, codebook, validation, deduplication, limitations, source authority, claim routing, and reproducibility boundaries.

Data and Information Management is modeled as a route for this paper because the fit is data/method infrastructure. Submission, repository deposit, licensing, disclosure, and venue decisions remain author-owned steps outside the portal.

## Limitations

Atlas has real limitations:

1. Legal and standards mappings remain candidate until clause, article, standard, and section anchors are verified.
2. Patent records are implementation signals.
3. Source-family volume and source authority are separate properties.
4. Raw observation counts and independent evidence weight are separate properties.
5. Submission status, repository deposit, license choice, disclosure filing, DOI reservation, and journal decision status require human-owned workflow records.
6. The public portal is a navigation surface; raw corpus histories and internal package audits remain controlled release artifacts.

## Public Portal

The public Atlas portal is available at:

https://atlas.eatf.eu/

The public manifest is available at:

https://atlas.eatf.eu/site-manifest.json

This PDF is a public article artifact. Source files and internal support packages remain controlled release artifacts.

## Selected Canonical Evidence Clusters

The following examples show canonical evidence carriers inside the larger navigation layer. Repeat observations are retained in audit records. The public manuscript reports source family and interpretation.

| Evidence carrier | Family | Public interpretation |
|---|---|---|
| Regulation (EU) 2016/679, GDPR | EU law | canonical legal-source card |
| Data Governance Act | EU law | official-source evidence card |
| Data Act | EU law | adjacent law-to-evidence source |
| NIS2 Directive | EU law | cybersecurity governance source |
| Cyber Resilience Act | EU law | product and cybersecurity evidence source |
| Product Liability Directive | EU law | adjacent liability source |
| Estonian Data Protection Inspectorate | Public authority | authority surface |
| Estonian Ministry of Economic Affairs and Communications | Public authority | national digital governance surface |
| Spanish AESIA | Public authority | AI supervision comparator |
| Estonian Information System Authority | Public authority | digital trust and infrastructure surface |
| CEN-CENELEC Artificial Intelligence | Standards | AI standardisation evidence card |
| IETF SCITT | Trust infrastructure | supply-chain integrity and transparency card |
| IETF RATS | Trust infrastructure | remote attestation card |

## Selected Patent and PQC Signals

Patent records reveal implementation imagination around trust infrastructure. The following examples are landscape signals.

| Signal family | Example patent signal | Jurisdiction / assignee | Safe use |
|---|---|---|---|
| PQC and PKI | Secure communication using post-quantum cryptography | United States / John A. Nix | post-quantum trust landscape |
| PQC and smartcards | Post-quantum cryptography on a smartcard | United States / Wells Fargo Bank | financial-sector PQC signal |
| PQC optimization | Post-quantum cryptography optimization | United States / Wells Fargo Bank | implementation-pattern signal |
| Quantum threat detection | Disparate quantum computing threat detection | United States / Wells Fargo Bank | risk and threat-model signal |

| Signal family | Example patent signal | Jurisdiction / assignee | Safe use |
|---|---|---|---|
| Quantum-resistant communication | Anti-quantum HTTPS communication based on CA and key pool | China / Nanjing Ruban Quantum Technology | PQC/CA architecture signal |
| Quantum key distribution | Network having quantum key distribution | European Patent Office / QinetiQ | quantum trust-infrastructure signal |
| Decentralized identity | Decentralized identity authentication framework for distributed data | United States / Ledgerdomain | identity and evidence-governance signal |
| Verifiable credentials | Digitally signed contracts with verifiable credentials | United States / Portable Data Corporation | credential carrier signal |
| Data provenance | System to manage data provenance | United States / Anahit Tarkhanyan | provenance infrastructure signal |
| Data supply chain | Provenance for data supply chains | United States / Intel | supply-chain evidence signal |
| Certificates of authenticity | Systems implementing certificates of authenticity | United States / Salesforce | software/product provenance signal |
| Zero trust access | Auditable and tamper-resistant remote zero trust access | United States / Raytheon | auditability and trust signal |
| AI accountability | AI agent decision platform with deontic reasoning | United States / Qomplx | AI accountability signal |
| Privacy-enabled AI networks | Privacy-enabled collaborative AI network | United States / Qomplx | privacy-preserving compute signal |
| Geospatial information separation | Functionally separating geospatial information | United States / Anonos Innovations | public-sector/accountability signal |

The patent layer helps identify where governance claims might touch technical implementation. Legality, adoption, field effectiveness, market relevance, and standards conformance require separate evidence.

## Reviewer Risks, Plainly Stated

Atlas records are canonicalized, deduplicated, grouped by source family, connected to manifests, and routed through claim gates.

Large counts are corpus-engineering facts. Raw observations, canonical records, and duplicate clusters are separated so repeated sightings of the same source remain provenance.

Legal claims remain candidate until exact clause-level verification is complete. Standards claims remain partial until exact body, version, and section anchors support a stronger statement.

Patent and PQC records are implementation signals. Public registers are public-facing schema analogues. Verifier-complete lifecycle evidence requires additional records.

Journal submission, repository deposit, license selection, disclosures, DOI reservation, and venue acceptance remain separate author-owned workflow decisions.

## Practical Use

Atlas can be used by a researcher to ask:

1. Which source families support a claim?
2. Which records are canonical?
3. Which claims are use-now, candidate, or appendix-only?
4. Which legal or standards claims need exact clause or section pinning?
5. Which public evidence carriers resemble verifier-readable governance objects?
6. Which patent/PQC signals show technical implementation imagination?
7. Which route controls remain before journal submission or repository deposit?

The portal exposes the research programme's topology before raw corpus publication.

## Copyright, Use, and Scope



Tyche Atlas is a research object and evidence-navigation infrastructure. The scope covers corpus method, evidence-family mapping, public navigation, claim-support status, and publication-readiness documentation. Legal advice, compliance assessment, standards conformance, certification, official audit, journal acceptance, repository deposit, DOI reservation, and third-party endorsement require separate competent-party records.