Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Pokala HealthIntel: A No-PHI Public-Source Healthcare Evidence Review Workspace

Krishna Sai Pokala

ok50250@umbc.edu

Independent Researcher, Health Informatics, United States https://orcid.org/0009-0001-3012-4018

Method Article

Keywords: health informatics, public-source healthcare data, evidence review, provenance, data lineage, no-PHI, implementation evaluation, public health data systems

Posted Date: July 7th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10245103/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Pokala HealthIntel: A No-PHI Public-Source Healthcare Evidence Review Workspace

Krishna Sai Pokala

Independent Researcher, Health Informatics, United States
Corresponding author: ok50250@umbc.edu

Abstract

Public-source healthcare datasets and APIs can support regulatory, safety, provider, reimbursement, and market-context review. However, review claims derived from those artifacts can become detached from source identity, row-count context, freshness, and limitations. This paper presents Pokala HealthIntel, a no-PHI public-source healthcare evidence review workspace designed to keep selected review claims linked to explicit evidence metadata and use boundaries. The system organizes preprocessed public-source artifacts into claim-linked evidence records and exposes them through a deployed review workspace and route surface. The evaluation assessed dependency installation, frontend build, endpoint availability and latency, evidence-record metadata coverage, canonical lineage verification, public-language boundary scanning, source-manifest status, and validation-tooling status. Dependency installation and frontend build passed. Six checked public routes returned 100% 2xx responses over 20 requests each, with median latencies ranging from 119.87 ms to 136.61 ms. Snapshot metadata reported 172,228 evidence rows and four sources. All four selected top-level evidence records contained claim/title, source, row-count, freshness, and limitation/boundary fields. The canonical lineage verifier passed, the public-language scan found no banned phrases, and source-manifest status was included with supporting materials. One local artifact-validation script remained a tooling-maintenance item after source-layout changes. Pokala HealthIntel demonstrates a modest, auditable pattern for public-source healthcare evidence review. The work is an implementation-level health-informatics systems demonstration, not clinical decision support, a diagnostic model, a treatment recommendation system, patient-level prediction, or a production clinical platform.

Keywords: health informatics; public-source healthcare data; evidence review; provenance; data lineage; no-PHI; implementation evaluation; public health data systems

1 Introduction

Public healthcare data sources can support regulatory review, safety monitoring, provider-context analysis, reimbursement context, and policy or market review. In the United States, public regulatory and health-system data are available through web interfaces, downloads, and APIs such as openFDA, ClinicalTrials.gov, provider registries, and public-use health-system artifacts. openFDA describes itself as an API that serves public FDA data about drugs, devices, and foods; it also states that not all openFDA data have been validated for clinical or production use and that the publicly available data do not contain patient personally identifiable information or other sensitive patient information [1]. ClinicalTrials.gov exposes public study records through a data API and associated data-element definitions [2]. These resources improve access to public artifacts, but access alone does not make downstream review workflows trustworthy.

The narrower problem addressed in this paper is that review claims derived from public artifacts can become separated from context needed for interpretation. A reviewer may see a safety note, provider-density observation, reimbursement-context summary, or market signal without immediately seeing the supporting source, the row scale of the artifact, its freshness, or the limitations that should constrain interpretation. In healthcare settings, this separation matters because public-source review outputs can be misread as clinical evidence, patient-level inference, or validated decision support when the evidence boundary is not visible.

Health informatics has long emphasized that data structure, provenance, documentation, and fitness for use are central to responsible secondary use. Provenance work, including the W3C PROV family, frames provenance as information about entities, activities, and people involved in producing data or things, useful for assessing quality, reliability, or trustworthiness [3]. FAIR principles frame scientific data and related digital objects around findability, accessibility, interoperability, and reusability [4]. Dataset documentation and model reporting traditions further emphasize motivation, composition, intended use, evaluation conditions, and limitations [5, 6]. These traditions do not directly define a public-source healthcare review workspace, but they motivate a design in which evidence summaries remain attached to source context and explicit use boundaries.

This paper presents Pokala HealthIntel, a no-PHI public-source healthcare evidence review workspace. The system question is: can selected healthcare review claims remain linked to source, row-count, freshness, and limitation metadata while the deployed workspace passes implementation-level checks for build, endpoint availability, evidence coverage, canonical lineage, public-language boundary, and source-manifest status? The contribution is a systems design and evaluation. It is not a clinical validation study and does not claim diagnostic, treatment, triage, or patient-level predictive capability.

2 Background and related work

2.1 Public-source healthcare data infrastructure

Public healthcare data infrastructure has expanded across regulatory, clinical-trial, provider, payer, and observational-data contexts. openFDA provides public FDA data through endpoint-specific API calls, returning query metadata and endpoint-specific records in JSON form [1]. ClinicalTrials.gov provides public access to study records and data-element definitions through a modernized data API [2]. Public provider and reimbursement resources similarly support review of healthcare-market and health-system context, although their structures, update rhythms, and use constraints vary by source.

Larger health-data infrastructures address related but different problems. The Observational Health Data Sciences and Informatics community and the OMOP Common Data Model support standardized observational health-data research and distributed analytics [7]. The FDA Sentinel Initiative demonstrates national-scale active safety surveillance using structured data networks and repeatable analytic processes [8]. These efforts show the value of common representation, standardization, and transparent analysis. Pokala HealthIntel is narrower. It does not attempt to replace common data models, public registries, or distributed surveillance networks. Instead, it asks how a no-PHI public-source review workspace can keep selected summary claims tied to source, scale, freshness, and limitation metadata.

2.2 Provenance, lineage, and data quality

Data provenance and lineage address the relationship between a data-derived output and the sources, activities, and agents that contributed to it. W3C PROV provides a family of documents for interoperable interchange of provenance information in heterogeneous environments [3]. Database and e-science provenance literature similarly emphasizes why users need to know how outputs were produced, which inputs were used, and how transformations or updates affect downstream products [9, 10].

Health-data quality research adds a complementary perspective: data may be available but still incomplete, inconsistent, stale, or not fit for a particular use. Work on electronic health-record data quality and secondary use emphasizes that completeness, concordance, plausibility, and context influence whether data can support a given analytic purpose [11, 12]. The present system is not an EHR system and does not use PHI, but the same informatics principle applies: a review claim should not be separated from the evidence context needed to judge its fitness for review.

2.3 FAIR principles and documentation traditions

The FAIR Guiding Principles describe findability, accessibility, interoperability, and reusability for scientific data and related digital objects, including algorithms, tools, and workflows [4]. FAIR does not require that every artifact be fully open or unrestricted; instead, it encourages clear metadata, identifiers, access context, and reuse conditions. This is relevant for public-source healthcare review because a derived claim may be visible even when the complete upstream retrieval path, raw public file, or transformation context is not included with the manuscript.

Datasheets for Datasets argues that dataset releases should document motivation, composition, collection process, recommended uses, and other contextual details [5]. Model Cards similarly emphasize intended use, evaluation conditions, ethical considerations, and limitations for released machine-learning models [6]. Pokala HealthIntel is not a predictive model and does not require a model card. However, these documentation traditions motivate the evidence-record requirement that every selected claim include a limitation or boundary field beside source and row-count context.

2.4 Responsible health-informatics boundaries

Clinical AI and medical-device software require evaluation frameworks that are not appropriate for a non-clinical public-source review workspace. Reporting frameworks for prediction models, clinical trials, and AI interventions emphasize intended use, evaluation population, comparator design, outcome definitions, and risk of harm [13, 14, 15]. Those concerns remain important, but this manuscript does not develop a patient-level predictor, diagnostic classifier, or treatment recommendation system. The relevant boundary is therefore different: the system should clearly prevent public-source artifact review from being interpreted as clinical decision support.

The no-PHI posture is not merely a privacy statement. It shapes what the workspace can and cannot claim. A no-PHI public-source artifact can be useful for exploratory review and context building, but it cannot support patient-specific inference without an entirely different data, governance, and validation design. This paper therefore treats boundary communication as part of the system design and part of the evaluation.

Table 1. Related-work positioning.

Area	Representative focus	Relevance to this work	Boundary of this work
Public healthcare APIs	Access to public records and endpoint-specific data	Provides public artifacts that can support review	Does not by itself solve claim traceability
Common data models	Standardized observational data and analytics	Shows importance of structure and reproducibility	Workspace does not replace OMOP/Sentinel-scale infrastructure
Provenance and lineage	Source-to-output context and derivation	Motivates source, row-count, and freshness fields	Not a complete PROV serialization
FAIR/data stewardship	Metadata, access context, and reuse	Motivates source-manifest and structured artifacts	Not a formal FAIR certification
Dataset/model documentation	Intended use, composition, evaluation, limits	Motivates limitation/boundary fields	Not a predictive model card
Health AI reporting	Clinical model/trial transparency	Clarifies what this system is not	No clinical validation claim

3 System design and evaluation

3.1 Study design

This is a systems demonstration and implementation-level evaluation. The object of study is a deployed public-source evidence review workspace and its selected evidence artifact layer. The evaluation asks whether the system can be installed, built, reached through checked public routes, and inspected for evidence-field coverage, canonical lineage, public-language boundary, and source-manifest status. It does not evaluate clinical safety, patient outcomes, diagnostic accuracy, treatment effect, or clinician decision-making.

3.2 Data-use boundary

The workspace is designed around a no-PHI boundary. It uses preprocessed public-source artifacts and does not redistribute patient-level identifiable data. It does not provide diagnosis, triage, treatment recommendation, risk stratification, or patient-specific guidance. These exclusions are not secondary disclaimers; they define the system scope. The appropriate interpretation is public-source evidence review, not clinical decision support.

3.3 Public-source artifact model

The system uses a snapshot-oriented artifact model. Public-source inputs are represented downstream as preprocessed artifacts and selected evidence records. The evidence layer reports snapshot-level metadata, including evidence-row and source counts, while top-level evidence records expose claim-linked metadata needed for interpretation. Raw public-source files are not redistributed with the manuscript. Source-manifest status is included to document the state of upstream-source disclosure,

Table 2. Claim-linked evidence-record fields.

Field	Purpose	Evaluation status
claim/title	States the review claim or summarized observation	Present in 4/4 selected records
source	Identifies the source category or artifact backing the claim	Present in 4/4 selected records
row count	Provides scale or row-count context	Present in 4/4 selected records
freshness	Provides recency or snapshot context	Present in 4/4 selected records
limitation/boundary	States what the claim should not be used to infer	Present in 4/4 selected records
workspace context	Places the record within review workflow context	Represented in workspace design

and future releases should expand this into retrieval URLs, retrieval dates, checksums, transformation notes, and row-count handoffs.

3.4 Claim-linked evidence schema

The central design choice is to treat a review claim as incomplete unless it remains attached to evidence metadata. The unit of design is not a chart or dashboard card; it is a claim-linked evidence record. Each selected evidence record contains a claim/title, source, row-count context, freshness context, limitation or boundary, and workspace context. This schema is deliberately small so that it can be inspected in a review workflow and evaluated with field-coverage checks.

3.5 Workspace and route-surface implementation

The deployed implementation uses a React/Vite frontend and a Cloudflare Worker route surface. The review workspace exposes summary, trend, evidence-map, evidence-detail, and source/validation contexts. The evaluation checked the public root route and health, intelligence, snapshot, and evidence routes. These route checks were intended to confirm public reachability during the evaluation run, not to provide a production service-level guarantee.

The deployment architecture is included because a live, reachable system is different from a static manuscript diagram. A public route surface allows reviewers to inspect whether the implementation exists and whether the evidence-review pattern is more than an offline mockup. At the same time, deployment on a public URL does not by itself establish security, production readiness, accessibility compliance, uptime reliability, or clinical suitability. Cloudflare documentation describes Worker code and static assets as deployable together for web applications, and Cloudflare D1 as a managed serverless database with SQLite semantics and Worker/HTTP API access [16, 17]. These references support deployment-context description only; they are not used to claim production reliability or clinical-grade infrastructure.

3.6 Evaluation protocol

The evaluation protocol combined local repository checks, live endpoint checks, evidence coverage inspection, lineage verification, public-language boundary scanning, source-manifest status inspection, and artifact-validation tooling status. Endpoint checks were executed over 20 requests per route. Evidence coverage was evaluated on the selected top-level evidence object and snapshot metadata.

Table 3. Evaluation protocol.

Check family	Procedure	Metric	Interpretation
Repository structure	Verify web and worker directories	Presence/absence	Implementation organization
Dependency installation	Run dependency installation in web app	Pass/fail	Local setup can resolve dependencies
Frontend build	Run frontend build	Pass/fail	Deployable frontend artifact
Endpoint availability	Request each checked route 20 times	2xx success rate	Public route reachability
Endpoint latency	Measure request duration	Median and p95 latency	Responsiveness in one run, not SLA
Evidence coverage	Inspect selected evidence object and metadata	Field completeness counts	Claim-linked evidence completeness
Canonical lineage	Run project lineage verifier	Pass/fail	Evidence/source metadata consistency
Language boundary	Scan public wording for banned phrases	Findings count	Non-clinical/no-overclaim posture
Source-manifest status	Check for status artifact	Present/absent	Reproducibility disclosure
Artifact-validation tooling	Run/check local validator	Pass/fail/maintenance	Validation-tooling condition

The canonical lineage verifier was treated as the authoritative lineage check for reported evidence-row and source-count consistency.

3.7 Evaluation metrics and interpretation rules

Endpoint metrics were 2xx success rate, median latency, p95 latency, and median response size. Evidence metrics were counts and fractions of selected records with claim/title, source, row count, freshness, and limitation/boundary fields. Metadata metrics were reported evidence rows and source count. Boundary metrics were canonical lineage result, public-language scan result, source-manifest status, and artifact-validation tooling status.

The interpretation rules were set before reading the evaluation results. Build, endpoint, evidence-coverage, canonical-lineage, language-boundary, and source-manifest status checks could support implementation-level system claims. The artifact-validation script could identify tooling-maintenance needs, but it was not used to override passed core checks unless it contradicted build, route, evidence, lineage, or language-boundary results. No metric was interpreted as clinical validation, patient-level performance, production uptime, or service-level reliability.

4 Results

4.1 Build checks

Dependency installation passed and the frontend build completed successfully. This supports the limited claim that the current source layout can produce the evaluated web artifact. It does not establish long-term maintainability or production-grade release engineering.

Table 4. Endpoint availability and latency.

Endpoint	2xx success	Median latency (ms)	p95 latency (ms)	Median bytes
/	100.0%	119.87	162.33	616
/api/health	100.0%	120.68	471.40	99
/health	100.0%	136.61	164.61	616
/api/intelligence	100.0%	129.75	150.03	616
/api/snapshot	100.0%	127.11	144.96	616
/api/evidence	100.0%	128.01	542.31	616

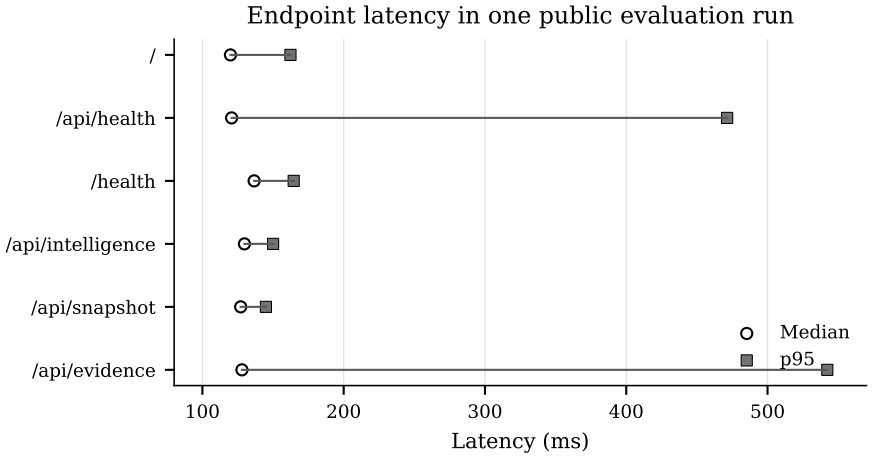


Figure 1. Endpoint latency during the evaluation run. Points show median and p95 latency for each checked endpoint. All endpoints returned 100% 2xx responses over 20 requests each. The values describe one public evaluation run, not an uptime guarantee or production service-level agreement.

4.2 Endpoint availability and latency

All six checked public routes returned 100% 2xx responses over 20 requests each. Median endpoint latency ranged from 119.87 ms to 136.61 ms. p95 values ranged from 144.96 ms to 542.31 ms (Table 4; Figure 1). These measurements describe one evaluation run against the deployed public URL and should not be interpreted as a long-term uptime guarantee, production capacity test, or service-level agreement.

4.3 Evidence metadata coverage

The selected evidence artifact contained four top-level evidence records. Snapshot metadata reported 172,228 evidence rows and four sources. All four selected records contained claim/title, source, row-count, freshness, and limitation/boundary metadata (Table 5). The evaluation checks selected top-level evidence records and snapshot metadata; it does not manually inspect every underlying row.

4.4 Lineage, language boundary, and source-manifest status

The canonical lineage verifier passed. The public-language boundary scan found no banned phrases. Source-manifest status was included with the supporting materials. One local artifact-validation

Table 5. Evidence coverage summary.

Coverage item	Observed value	Interpretation
Selected top-level evidence records	4	Evidence records evaluated
Evidence rows metadata	172,228	Snapshot-level row metadata
Source count metadata	4	Snapshot-level source metadata
Claim/title coverage	4/4	Complete for selected records
Source coverage	4/4	Complete for selected records
Row-count coverage	4/4	Complete for selected records
Freshness coverage	4/4	Complete for selected records
Limitation/boundary coverage	4/4	Complete for selected records

Table 6. Evaluation status and interpretation.

Check	Status	Interpretation
Dependency installation	Passed	Supports local setup of evaluated web app
Frontend build	Passed	Supports deployable web artifact claim
Endpoint availability	Passed	All checked routes returned 100% 2xx responses
Evidence metadata coverage	Passed	All selected records had required fields
Canonical lineage verifier	Passed	Evidence-row and source-count metadata passed canonical verification
Public-language boundary	Passed	No banned public-language phrases found
Source-manifest status	Included	Supports disclosure of source-manifest status
Artifact-validation tooling	Maintenance item	Disclosed tooling limitation, not a core system claim

script remained a tooling-maintenance item after source-layout changes. This remaining item is disclosed as a validation-tooling limitation and future-work item; it is not used as a basis for clinical, production, or raw-source reproducibility claims.

5 Discussion

5.1 Principal finding

The principal finding is that a deployed no-PHI public-source healthcare evidence review workspace can keep selected review claims connected to source, row-count, freshness, and limitation metadata while passing implementation-level checks. The system passed dependency installation, frontend build, endpoint availability, evidence metadata coverage, canonical lineage verification, public-language boundary scanning, and source-manifest status checks. This supports a modest health-informatics systems claim: the workspace demonstrates an auditable evidence-review pattern for public-source healthcare artifacts.

5.2 Why this is more than a dashboard

A dashboard summarizes indicators. Pokala HealthIntel emphasizes the evidence boundary behind a summarized claim. The difference is not only visual. The review unit is structured so the claim,

source, row-count scale, freshness context, and limitation boundary remain close together. This reduces the risk that a reviewer interprets a summary without understanding the artifact context behind it. The system therefore contributes a review pattern rather than a new clinical model.

5.3 Relationship to provenance and FAIR principles

The workspace is provenance-inspired and FAIR-aligned, but it is not a formal provenance or FAIR certification. The evidence-record design reflects provenance concerns by attaching source and context to claims. The source-manifest status and structured evidence fields reflect FAIR concerns for findability, accessibility, interoperability, and reusability. Future work should formalize these relationships more completely, including upstream retrieval metadata, stable identifiers, checksums, and machine-readable provenance statements.

5.4 Health-informatics relevance

The no-PHI and non-clinical boundary is central to the design. Public-source healthcare review tools can be useful without becoming clinical decision support. The manuscript therefore avoids claims about diagnosis, treatment, patient-specific prediction, or clinical outcome impact. This boundary is important for scientific honesty, public deployment, and scholarly credibility. It also reflects a practical health-informatics skill: translating public healthcare artifacts into responsible review workflows without overstating the evidentiary status of those artifacts.

5.5 What remains before stronger research claims

A stronger research release should expand source-manifest status into a complete upstream-source manifest with retrieval URLs, retrieval dates, checksums, transformation notes, and row-count handoffs. Local validation scripts should be consolidated under the current source layout. A reviewer-task evaluation could test whether users can find source, row-count, freshness, and limitation context quickly and accurately. External comparison against other public-healthcare review tools could clarify practical value. A production-readiness study would require uptime monitoring, security review, accessibility review, load testing, release governance, and operational incident handling. A clinical study would require an entirely different protocol and is outside the scope of this work.

6 Limitations

This work has important limitations. First, the system does not use PHI and does not perform patient-level prediction. It does not diagnose, triage, recommend treatment, or provide clinical decision support. Second, the evaluation is implementation-level; it does not evaluate clinical validity, patient outcomes, clinician behavior, safety impact, or prospective deployment. Third, endpoint latency reflects one evaluation run against the public URL and should not be interpreted as an uptime guarantee, capacity benchmark, or service-level agreement.

Fourth, evidence coverage was evaluated for four selected top-level evidence records and snapshot metadata. The evaluation did not manually inspect every underlying row represented by the reported evidence-row metadata. Fifth, the source-manifest status file was included, but future work should expand upstream retrieval metadata, retrieval dates, checksums, and transformation notes. Sixth, one local artifact-validation script remained a tooling-maintenance item after source-layout changes. Seventh, no formal user study, accessibility audit, security review, load test, or comparison

Table 7. Manuscript claim boundaries.

Supported claim	Excluded claim	Reason
Deployed no-PHI evidence review workspace	Clinical decision support	No diagnosis, triage, treatment recommendation, or patient-specific guidance is provided
Claim-linked evidence metadata for selected records	Clinical validation	No clinical outcome, clinician-behavior, or prospective study was performed
Implementation-level route availability during one evaluation run	Production service-level guarantee	Endpoint measurements reflect one run, not longitudinal uptime or load testing
Canonical lineage verification of reported metadata	Complete raw-source reproducibility	Upstream retrieval URLs, dates, and checksums remain future work
Public-language boundary scan	Regulatory clearance or medical-device status	No medical-device or clinical workflow evaluation was performed

against institutional evidence-review tools was performed. These limitations restrict the claims to a deployed systems demonstration and implementation-level evaluation.

7 Conclusion

Pokala HealthIntel demonstrates that a no-PHI public-source healthcare evidence review workspace can keep selected review claims linked to source, row-count, freshness, and limitation metadata while passing implementation-level checks for build, endpoint availability, evidence metadata coverage, canonical lineage, public-language boundary, and source-manifest status. The contribution is an evidence-review pattern for public-source healthcare artifacts, not a clinical model. The work should be interpreted as a health-informatics systems demonstration rather than clinical validation, patient-level prediction, treatment guidance, or production clinical decision support.

Ethics statement

This work uses no PHI and no individual-level identifiable data. The manuscript describes a public-source evidence review workspace and implementation-level evaluation. No human-subjects intervention, clinical trial, patient-level prediction, diagnosis, triage, or treatment recommendation was performed.

Data availability

The deployed workspace uses preprocessed public-source artifacts and selected evidence records. Raw public-source files are not redistributed with the manuscript. Supporting materials include source-manifest status and evaluation summaries. Future releases should expand the upstream-source manifest with retrieval URLs, retrieval dates, checksums, and transformation notes.

Code availability

The project is intended to be made available as a public repository accompanying the preprint. The manuscript's core claims are based on the evaluated deployed system, selected evidence artifacts, and implementation-level checks described above.

Funding

No external funding was received.

Competing interests

The author declares no competing interests.

Author contributions

Krishna Sai Pokala conceived the system, developed the software, prepared the evidence-review design, ran the evaluation workflow, interpreted the results, and prepared the manuscript.

References

[1] U.S. Food and Drug Administration. openFDA documentation. `https://open.fda.gov/apis/`.

[2] National Library of Medicine. ClinicalTrials.gov API and data-element documentation. `https://clinicaltrials.gov/data-api/about-api`.

[3] W3C. PROV-Overview: An overview of the PROV family of documents. `https://www.w3.org/TR/prov-overview/`.

[4] Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*. 2016;3:160018.

[5] Gebru T, Morgenstern J, Vecchione B, et al. Datasheets for Datasets. *Communications of the ACM*. 2021;64(12):86-92.

[6] Mitchell M, Wu S, Zaldivar A, et al. Model Cards for Model Reporting. Proceedings of FAT*. 2019.

[7] Hripcsak G, Duke JD, Shah NH, et al. Observational Health Data Sciences and Informatics (OHDSI): Opportunities for observational researchers. *Studies in Health Technology and Informatics*. 2015;216:574-578.

[8] Platt R, Brown JS, Robb M, et al. The FDA Sentinel Initiative - An evolving national resource. *New England Journal of Medicine*. 2018;379:2091-2093.

[9] Simmhan YL, Plale B, Gannon D. A survey of data provenance in e-science. *SIGMOD Record*. 2005;34(3):31-36.

[10] Cheney J, Chiticariu L, Tan WC. Provenance in databases: why, how, and where. *Foundations and Trends in Databases*. 2009;1:379-474.

[11] Kahn MG, Callahan TJ, Barnard J, et al. A harmonized data quality assessment terminology and framework for the secondary use of electronic health record data. *eGEMs*. 2016;4(1):1244.

[12] Weiskopf NG, Weng C. Methods and dimensions of electronic health record data quality assessment. *Journal of the American Medical Informatics Association*. 2013;20(1):144-151.

[13] Collins GS, Moons KGM, Dhiman P, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ*. 2024.

[14] Liu X, Rivera SC, Moher D, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nature Medicine*. 2020;26:1364-1374.

[15] Vasey B, Nagendran M, Campbell B, et al. DECIDE-AI: new reporting guidelines to bridge development and evaluation of clinical artificial intelligence systems. *Nature Medicine*. 2022;28:190-191.

[16] Cloudflare. Cloudflare Workers documentation. `https://developers.cloudflare.com/workers/`.

[17] Cloudflare. D1 documentation. `https://developers.cloudflare.com/d1/`.

[18] Centers for Medicare & Medicaid Services. Public data and research files. `https://www.cms.gov/data-research`.

[19] National Plan and Provider Enumeration System. NPI Registry public search and data dissemination resources. `https://npiregistry.cms.hhs.gov/`.

[20] Office of the National Coordinator for Health Information Technology. Health IT data and standards resources. `https://www.healthit.gov/`.

[21] Greenes RA. *Clinical Decision Support: The Road to Broad Adoption*. Academic Press; 2014.

[22] Hersh WR. *Health Informatics: Practical Guide for Healthcare and Information Technology Professionals*. Lulu; 2022.