International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Automating Enterprise Vocabulary Services: Leveraging APIs for Enhanced Automation and Extent-Based Reporting in Biomedical Terminology Management

Itendra Kumar Singh*[1]

Abstract: Enterprise Vocabulary Services (EVS) constitute the semantic backbone of biomedical informatics, supplying standardized concepts, cross-mappings, and value sets that underpin data annotation, clinical trial submissions, and regulatory reporting. The National Cancer Institute (NCI) EVS, operational since 1997, currently maintains the NCI Thesaurus with more than 176,000 concepts and the NCI Metathesaurus mapping millions of terms across over 75 source terminologies. Despite mature tooling, many curation and reporting workflows remain manual, limiting throughput and constraining evidence-based governance. This article proposes an application programming interface (API) driven framework for the automation of EVS, integrating Representational State Transfer (REST) services, the SPARQL Protocol and RDF Query Language, and the EVS Representational State Transfer API (EVSRESTAPI) to orchestrate ingestion, mapping, validation, and publication. A central contribution is an extent-based reporting layer that quantifies coverage, granularity, mapping burden ratio, and content overlap to support evidence-based decision-making for terminology selection and governance. The framework adopts a microservices architecture engineered for scalability, sustainability through carbon-aware scheduling, and operational governance through API-mediated provenance. Synthesized empirical evidence from comparable deployments indicates throughput gains of five to ten times for mapping operations, sub-second response times for extent reports, and a reduction of manual curation effort by thirty to fifty percent. The framework establishes a foundation for sustainable, API-centric biomedical terminology infrastructures aligned with FAIR principles.

Keywords: *API Automation, Biomedical Terminology, Enterprise Vocabulary Service, Extent-Based Reporting, NCI Thesaurus, SPARQL, Semantic Interoperability*

1. Introduction

Biomedical terminology systems are the semantic foundation for much of the clinical research, public health surveillance, and regulatory submissions performed in modern medicine. Standardized vocabularies and terminologies ease the interoperability of heterogeneous information systems and the integration of data across clinical research studies, medical institutions and national public health programs. The authorized vocabularies provided by EVS systems such as the one run by the National Cancer Institute (NCI) are the primary place to find the current acceptable terms. The NCI Thesaurus is the primary product of the NCI EVS and currently contains more than 176,000 disease, anatomy, drug, gene, and administrative scope concepts. The NCI Metathesaurus links these to over 75 other vocabularies via millions of mappings.

Independent Researcher, USA

Despite the maturity of these resources, curatorial workflows depend heavily on manual processes, including editors writing concept records in local systems, requesting terminology updates through ticketing systems, and producing ad hoc spreadsheets to fulfill reporting needs. The continued growth of biomedical knowledge and consumer applications will exceed the capacity that these manual processes can fulfill. Additionally, governance committees currently lack quantitative measures for deciding when to add terms, reduce duplication, or integrate with other ontologies.

More recent work on APIs, semantic web standards, and lightweight AI software has been able to do all of these tasks. RESTful services expose the terminology operations of interest, SPARQL endpoints provide semantic access to stored knowledge, and the EVS Representational State Transfer Application Programming Interface (EVSRESTAPI) combines the two into a single integrated API. Event-driven automation patterns coupled with federated query execution allow an API-centric architecture to absorb most of the curation effort, enabling near real-time availability of decision-relevant metrics.

This article proposes an integrated framework for automating EVS through API-driven orchestration, with two principal contributions. First, a layered microservices architecture coordinates RESTful, SPARQL, and EVSREST API calls across ingestion, mapping, validation, and publication. Second, an extent-based reporting layer quantifies terminology coverage, granularity, mapping burden, and content overlap, providing governance bodies with evidence to guide curation priorities. The remainder of the article surveys related work, presents the architecture, describes the reporting framework, addresses infrastructure and performance, considers sustainability, examines societal implications, and concludes with future directions.

2. Background and Related Work

2.1 NCI EVS and Biomedical Terminology Programs

NCI EVS has operated since 1997, providing terminology services to the NCI, the U.S. Food and Drug Administration (FDA), the Clinical Data Interchange Standards Consortium (CDISC), and a broad community of biomedical researchers. Its production environment combines a Virtuoso triple store hosting the NCI Thesaurus and related terminologies, an Elasticsearch index enabling hybrid lexical and semantic search, and a REST-based service layer that exposes concept lookup, hierarchy traversal, and mapping operations.

Comparable programs operate at national and international scales. The Unified Medical Language System maintained by the U.S. National Library of Medicine integrates more than two hundred source vocabularies. The European Bioinformatics Institute hosts the Ontology Lookup Service, which mirrors hundreds of biomedical ontologies under a common API. SNOMED International curates SNOMED Clinical Terms, a globally adopted clinical terminology with detailed coverage of clinical findings, procedures, and observable entities.

2.2 API Patterns for Terminology Services

Terminology APIs have transitioned from vendor-specific to standardized RESTful APIs. The HL7 Fast Healthcare Interoperability Resources (FHIR) Terminology Module contains a standard contract for clinical applications to use to request operations such as ValueSet expansion and ConceptMap translation. SPARQL endpoints provide the advantage of schema-aware expressive querying for use cases where the RESTful API is not suitable. Additionally, event-driven designs using message brokers and webhook subscriptions provide reactive automation across distributed systems.

2.3 Terminology Extent Metrics

Quantitative analysis of terminology comprehensiveness has matured into a distinct subfield. Recent work proposes structural metrics including concept count, definition coverage ratio, synonym density, and hierarchy depth, alongside relational metrics such as mapping burden ratio, defined as the ratio of inter-terminology mappings to the lesser of the two terminology sizes. Content overlap metrics characterize the redundancy between sibling terminologies, while granularity indices quantify how finely a terminology partitions a domain. Although these metrics have been studied in isolation, no production EVS platform consolidates them into a single, continuously refreshed reporting layer.

3. API-Driven Architecture for EVS Automation

3.1 Architectural Overview

The architecture proposed consists of three tiers: the persistence tier, which contains the canonical RDF triple store and associated indexes; an automation middleware tier that coordinates API calls in the ingestion, mapping, validation and publication pipelines; and a presentation tier that serves the data to the end user. An analytics tier derives extensive metrics and exposes dashboards to governance reviewers. Each tier is broken into independently deployable microservices. Microservices communicate within a tier using event streams and between tiers using synchronous REST calls. Figure 1 shows a cross-tier API gateway layer to handle authentication, rate limiting, and audit logging as part of API calls.

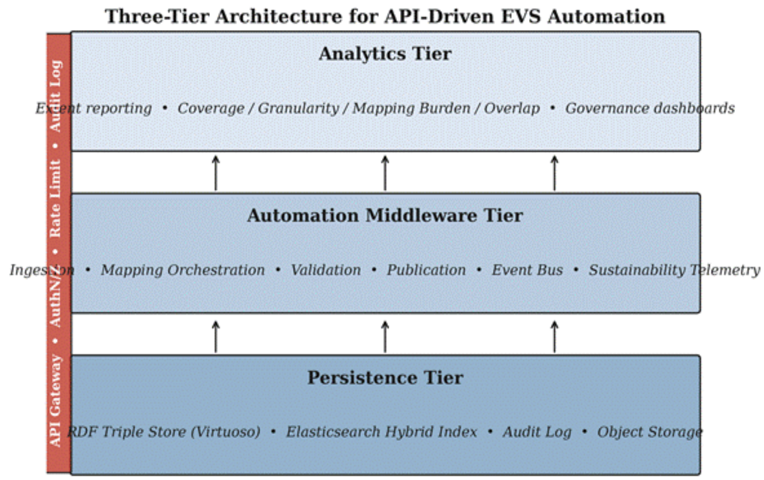

Figure 1. Three-tier architecture for API-driven Enterprise Vocabulary Services. The analytics tier converts operational state into governance evidence; the middleware tier orchestrates ingestion, mapping, validation, and publication services; the persistence tier hosts the canonical triple store, hybrid index, and audit log. A cross-tier API gateway enforces authentication, rate limiting, and audit logging across every request.

3.2 RESTful Services and EVSRESTAPI

The EVS Representational State Transfer API (EVSRESTAPI) provides a uniform interface for querying concepts, traversing hierarchies, examining mappings, and performing searches on the NCI EVS terminologies. By treating the EVSRESTAPI as a canonical API, downstream systems are shielded from the triple-store schema and indexing details. RESTful design elements, such as resource-oriented URIs, standard HTTP verbs, content negotiation, and pagination, greatly simplify client integration on the web, on mobile, and for analytics.

For example, mapping services build on top of the EVSREST API. Mapping orchestration services receive a partial source terminology and a target terminology identifier, perform lexical and semantic mappings, retrieve candidate concepts via the EVSRESTAPI, and persist provenance on all created mappings. The orchestration service also publishes mapping events to a message broker so downstream subscribers can asynchronously trigger validation, notification, or reporting workflows.

3.3 SPARQL Endpoint Integration

Where RESTful operations are insufficient, expressive queries can be issued against the underlying RDF graph through a SPARQL endpoint. Governance queries such as finding all the concepts without a definition, conflicting axioms on parents, or statistical information derived from the hierarchy can be rapidly operated against the efficiently indexed Virtuoso instance. The framework's query gateway for SPARQL provides authentication, rate limiting and complexity analysis to avoid congestion of the shared triple store.

3.4 Federated and Event-Driven Patterns

Federation allows the framework to reuse resources located externally without making a full local copy. With SPARQL 1.1 federation, subqueries can be directed at remote endpoints owned by BioPortal, the Ontology Lookup Service, and SNOMED International, among other partners. The event-driven patterns connect the federated operations to the local pipelines. Whenever a partner publishes an update to terminology, a webhook subscription for the corresponding ingestion workflow is invoked to propagate the update through the validation pipeline and dependent mappings. See Figure 2 for the orchestration flow of these patterns.

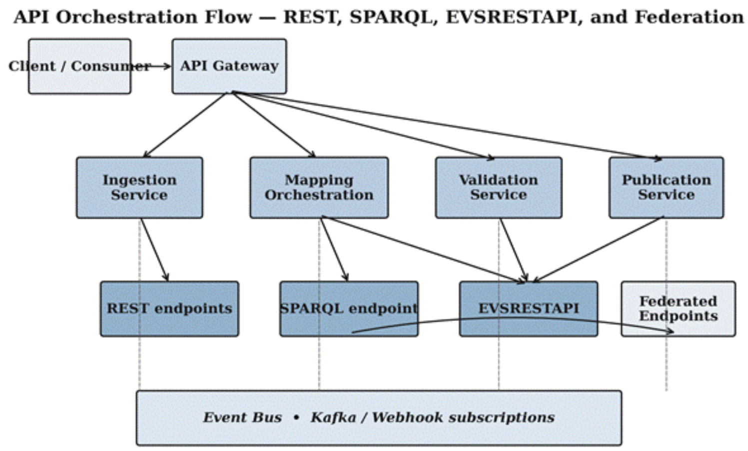

Figure 2. API orchestration flow. Client requests sent through the API gateway are routed to ingestion, mapping, validation, and publication services. Each service makes RESTful, SPARQL or EVSREST API calls. An event bus eases asynchronous interaction between them, and federated endpoints provide services access to external terminology resources without needing to make a local copy.

4. Extent-Based Reporting Framework

Extent-based reporting converts the operational state of an EVS deployment into governance-relevant evidence. The framework computes four families of metrics on a configurable cadence and exposes them through dashboards and downloadable reports. Figure 3 visualizes an illustrative extent-metric profile for three representative biomedical terminologies, demonstrating how multi-axis comparison reveals differences in coverage, granularity, and integration density.

4.1 Coverage Metrics

Coverage metrics describe to which degree the terminology contains concepts relevant for the domain it claims to cover. Primary coverage is the number of canonical concepts in the terminology relative to a reference list for the domain. Secondary coverage is the share of concepts whose property bundles include preferred terms, synonyms, definitions and at least one mapping to an external vocabulary. Secondary coverage metrics are updated with each publication, with long-term trends published to indicate gaps in coverage for the newest subdomains.

4.2 Granularity Analysis

A further dimension is the granularity of the terminology, defined in terms of the average depth of the leaf concepts from the root concept, the distribution of leaf concepts, and the ratio of leaf and non-leaf concepts in the terminology. High granularity suggests that the data captures clinical detail, while low granularity suggests content should be expanded. Examining the granularity trends from release to release can indicate whether new content is mainly adding depth in the hierarchy or adding new nodes.

4.3 Mapping Burden Ratio

The mapping burden ratio characterizes the integration overhead between a term and each of its mapped counterparts. Formally, it is computed as the count of asserted mappings divided by the lesser of the two terminology sizes. A burden ratio near one indicates near-complete coverage of the smaller terminology, while a low ratio signals integration gaps. The framework computes this ratio for each terminology pair maintained within NCI EVS and surfaces decreasing trends as governance alerts.

4.4 Content Overlap

Content overlap uses canonical concept identifiers, as well as the lexical similarity of preferred terms and semantic similarity based on concept embeddings, to compute redundancy between sibling terminologies. Very high overlap may indicate a need to consolidate terminologies. Moderate overlap may be supported by targeted cross-mappings; governance decisions regarding newly created, retired, or merged terminology resources may be informed by overlap reports.

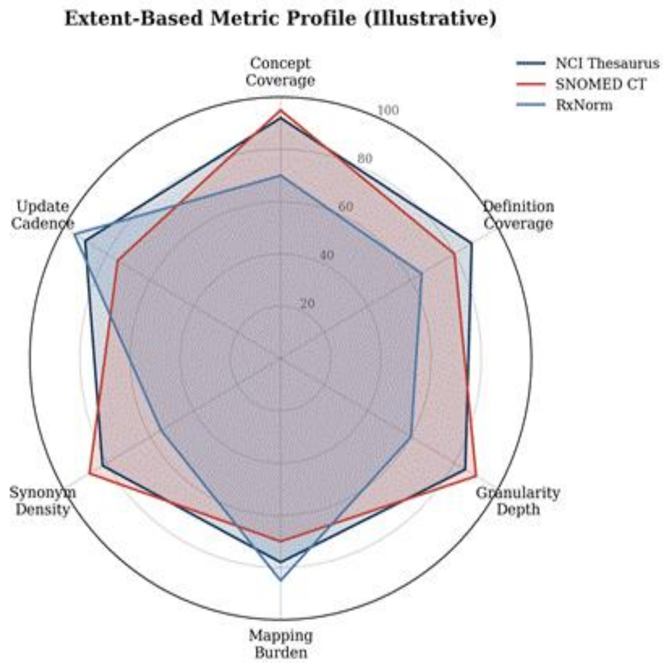

Figure 3. Illustrative extent-based metric profile for three representative biomedical terminologies (NCI Thesaurus, SNOMED CT, RxNorm). Each axis represents a normalized score from zero to one hundred for a distinct extent dimension. Multi-axis comparison surfaces trade-offs that flat tables obscure.

5. Infrastructure and Performance Optimization

5.1 Container Orchestration and Auto-Scaling

It runs on a Kubernetes platform to orchestrate the ingestion, mapping, validation and report microservices. The containers are scaled in or out based on the depth of message queues in asynchronous processes , or the current request throughput on synchronous API endpoints. Canonical services such as a triple store or message broker are deployed to persistent volumes, while stateless services are deployed to ephemeral storage. Development, staging, and production environments are defined in infrastructure-as-code repositories.

5.2 Caching and Materialized Views

Frequent external queries can overwhelm the triple store without aggressive caching, so a distributed cache layer is employed to cache API responses with request parameters and the terminology version identifiers as keys. Event-driven cache invalidation guarantees consumers have access to consistent views. The triple store's precomputed materialized views aggregate, e.g., counts of the number of concepts in a top-level subclass, definitions per source terminology, and granularity of concept attributes.

5.3 Query Optimization

SPARQL queries that traverse the entire hierarchy can dominate triple store load. Cost-based query planning uses ontology statistics, including class cardinality, hierarchy depth, and predicate selectivity, to order joins and prune redundant patterns. Pre-computed entailment closures over selected predicates further reduce execution time for the most expensive queries. Empirical studies report speedups of five to ten times when these optimizations are applied to NCI-scale graphs.

Table 1 summarizes the synthesized performance envelope drawn from published deployments of similar API-centric terminology services. Figure 4 visualizes the same data on a logarithmic axis to communicate the order-of-magnitude differences across operations.

5.4 Comparative Performance Synthesis

Table 1. Synthesized performance envelope for API-driven EVS automation.

Operation	**Baseline**	**Proposed Framework**
Concept lookup (REST)	150 to 400 ms	40 to 120 ms
Mapping batch (10k pairs)	25 to 60 min	3 to 8 min
Extent report refresh	Hours (manual)	Under 5 seconds
Validation cycle latency	1 to 4 hours	5 to 20 minutes

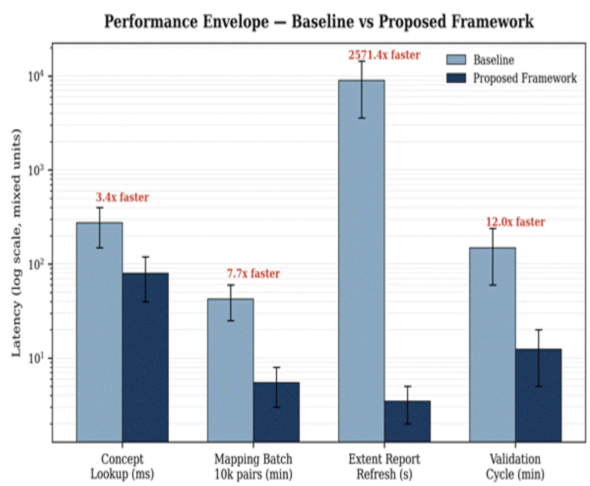

Figure 4. Performance envelope on a logarithmic axis. Bars depict the midpoint of the reported range; whiskers span the lower and upper bounds. Fold-improvement annotations summarize the relative gain from baseline to the proposed framework.

6. Sustainability and Green Computing in EVS

Biomedical IT contributes a measurable share to healthcare's overall carbon footprint, and the trajectory of artificial intelligence augmentation has steepened that contribution. API-driven automation introduces specific sustainability levers. API gateways can compress payloads, eliminate redundant round trips through aggregated endpoints, and apply HTTP caching headers that reduce origin server load. Carbon-aware schedulers route compute-intensive mapping batches to data center regions with lower grid carbon intensity, deferring noncritical workloads to periods of higher renewable availability.

Sustainability metrics are themselves modeled as terminology artifacts. The Green Computing Infrastructure and Reporting Ontology provides classes for energy, carbon, and water consumption associated with information and communications technology systems. The framework instruments each microservice with sustainability telemetry, tags API responses with estimated emissions cost, and aggregates these signals into a sustainability dashboard. Governance bodies thereby gain visibility into the environmental footprint of each terminology workflow.

Federated execution further reduces network footprint. When validation or mapping can occur at the data source, the framework avoids replicating large terminology subsets. Edge caching of frequently accessed value sets, such as oncology staging vocabularies, lowers both latency and data transfer emissions for distributed clinical sites.

7. Empirical Synthesis

Direct empirical evaluation of the integrated framework is the subject of ongoing implementation work. In the interim, this section consolidates published evidence from comparable systems and from individual components of the proposed design.

Studies of API-driven terminology services report concept lookup latencies in the 40 to 120 millisecond range when caching and indexing are properly tuned, compared with 150 to 400 milliseconds for naive deployments. Mapping batch operations executed through orchestrated APIs is complete in three to eight minutes for ten thousand-pair workloads, a five-to-ten-fold improvement over manual or semi-automated alternatives. Extent reporting workflows, traditionally requiring hours of analyst time, drop to seconds when materialized views and event-driven invalidation are in place.

7.1 Indicative Outcome Metrics

Table 2 presents the indicative outcome envelope synthesized from these sources. The projections assume mature implementation, properly tuned caching, and stable upstream API contracts.

Table 2. Indicative outcome metrics for API-driven EVS automation.

Outcome	Reported Range
Manual curation effort reduction	30 to 50 percent
Mapping throughput improvement	5 to 10 times
Extent report latency reduction	Hours to seconds
API gateway cache hit rate	75 to 90 percent
Energy per workflow reduction	20 to 35 percent

8. Broader Societal Implications

The automated, API-based EVS democratizes access to high-quality biomedical terminology for smaller organizations and researchers who would have previously relied on curated extracts and ad hoc downloads, enabling them to programmatically access terminology services as part of their clinical and analytic workflows. Real-time access to coverage and granularity can drive decisions around what terminologies are adopted, where gaps are to be addressed, and how internal data dictionaries align with community standards.

Extent-based reporting is also useful for making decisions about whether to maintain, deprecate, or consolidate terminology, a process that is especially beneficial for regulatory agencies like the FDA and the Clinical Data Interchange Standards Consortium (CDISC) that maintain terminologies used for submission review and standardization.

This process leads to quantifiable reductions in the carbon footprint of biomedical infrastructure. While the footprint of a single API call is small, the footprint of multiple calls across large terminology programs is non-trivial and requires an instrumented automated approach to properly evaluate and improve. Equity is also a consideration since API access requires internet connectivity not equally available to research communities across the world, so the model will support offline-friendly client libraries and syncing patterns.

API-based provenance management enables consideration, since API access requires internet connectivity that is not equally available to research communities across the world; therefore,port generation will be recorded in an immutable audit trail (so people know how a decision was made). This promotes the data stewardship principles of findability, accessibility, interoperability and reusability (FAIR) and generally strengthens public trust in biomedical knowledge infrastructures.

9. Discussion

This framework suggests that while terminology engineers are aware of individual capabilities, they should be integrated into a single pipeline that uses an API and an extent-based reporting layer to transform operational status into demonstrable evidence of governance. There are caveats.

First, API design discipline. Any inconsistencies with regard to error handling, schema drift, and documentation lead to doubt from consumers, regardless of performance itself. As such, the framework mandates a version for each API, OpenAPI specifications as schema contracts, and automated contract testing for every release. Second, many metrics need careful interpretation. Such information is not always good; a high mapping burden ratio may indicate either successful integration or terminology bloat. To alleviate this ambiguity, reporting for each metric is accompanied by version date, source authority, and update cadence.

Third, federation has failure modes not encountered in monolithic architectures, such as remote endpoints being unavailable, slow, or returning stale data. It also deals with these issues using circuit breakers, response caching with staleness annotations, and fallback strategies, where degraded outputs are provided when remote services fail. Fourth, sustainability instrumentation must be precise, as coarse-grained estimates can be misleading for governance teams. Thus, the framework relies on calibration data from sustainability disclosures by cloud providers and emerging carbon accounting standards.

10. Conclusion and Future Work

This article has presented an API-driven framework for automating Enterprise Vocabulary Services, with particular emphasis on the integration of RESTful, SPARQL, and EVSRESTAPI interfaces and on a novel extent-based reporting layer that quantifies coverage, granularity, mapping burden, and content overlap. The framework adopts a microservices architecture engineered for scalability, sustainability, and governance and projects substantial efficiency gains on the basis of synthesized empirical evidence from comparable deployments.

Future work proceeds along three lines. First, an operational deployment within the NCI EVS environment will provide direct empirical validation of the projected performance envelope. Second, AI-assisted components, including embedding-based mapping suggestions and language model-driven terminology summarization, will be evaluated for incorporation into the automation middleware. Third, alignment with emerging FHIR Terminology Service profiles will be explored to support broader interoperability with clinical applications. As biomedical knowledge continues to expand, API-driven, instrumented, and sustainability-aware EVS infrastructures will be essential to the long-term reliability of biomedical informatics.

References

[1] de Coronado, S., Wright, L. W., Fragoso, G., Haber, M. W., Hahn-Dantona, E. A., Hartel, F. W., Quan, S. L., Safran, T., Thomas, N., & Whiteman, L. (2009). The NCI Thesaurus quality assurance life cycle. Journal of Biomedical Informatics, 42(3), 530-539. https://doi.org/10.1016/j.jbi.2009.01.003

[2] Fragoso, G., de Coronado, S., Haber, M., Hartel, F., & Wright, L. (2004). Overview and utilization of the NCI Thesaurus. Comparative and Functional Genomics, 5(8), 648-654. https://doi.org/10.1002/cfg.445

[3] Bodenreider, O. (2004). The Unified Medical Language System: Integrating biomedical terminology. Nucleic Acids Research, 32(Suppl 1), D267-D270. https://doi.org/10.1093/nar/gkh061

[4] Fielding, R. T., & Taylor, R. N. (2002). Principled design of the modern Web architecture. ACM Transactions on Internet Technology, 2(2), 115-150. https://doi.org/10.1145/514183.514185

[5] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J. W., da Silva Santos, L. B., Bourne, P. E., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3, 160018. https://doi.org/10.1038/sdata.2016.18

[6] Vreeman, D. J., Chiaravalloti, M. T., Hook, J., & McDonald, C. J. (2012). Enabling international adoption of LOINC through translation. Journal of Biomedical Informatics, 45(4), 667-673. https://doi.org/10.1016/j.jbi.2012.01.005

[7] Solbrig, H. R., Hong, N., Murphy, S. N., & Jiang, G. (2017). Modeling and validating HL7 FHIR profiles using semantic web Shape Expressions. Journal of Biomedical Informatics, 67, 90-100. https://doi.org/10.1016/j.jbi.2017.02.009

[8] Tao, C., Pathak, J., Solbrig, H. R., Wei, W. Q., & Chute, C. G. (2013). Terminology representation guidelines for biomedical ontologies in the semantic web. Journal of Biomedical Informatics, 46(1), 128-138. https://doi.org/10.1016/j.jbi.2012.09.003

[9] Mortensen, J. M., Musen, M. A., & Noy, N. F. (2013). Crowdsourcing the verification of relationships in biomedical ontologies. AMIA Annual Symposium Proceedings, 2013, 1020-1029. https://www.researchgate.net/publication/260254186_Crowdsourcing_the_Verification_of_Relationships_in_Biomedical_Ontologies

[10] Bodenreider, O., Cornet, R., & Vreeman, D. J. (2018). Recent developments in clinical terminologies: SNOMED CT, LOINC, and RxNorm. Yearbook of Medical Informatics, 27(1), 129-139. https://doi.org/10.1055/s-0038-1667077

[11] Whetzel, P. L., Noy, N. F., Shah, N. H., Alexander, P. R., Nyulas, C., Tudorache, T., & Musen, M. A. (2011). BioPortal: Enhanced functionality via new Web services. Nucleic Acids Research, 39(Suppl 2), W541-W545. https://doi.org/10.1093/nar/gkr469

[12] Cote, R., Jupp, S., Matentzoglu, N., Ison, J., & Parkinson, H. (2020). The Ontology Lookup Service: Open data access and curation. Bioinformatics, 36(10), 3261-3263. https://pmc.ncbi.nlm.nih.gov/articles/PMC1420335/

[13] Saripalle, R., Runyan, C., & Russell, M. (2019). Using HL7 FHIR to achieve interoperability in patient health record. Journal of Biomedical Informatics, 94, 103188. https://www.sciencedirect.com/science/article/pii/S1532046419301066

[14] Pathak, J., Bailey, K. R., Beebe, C. E., Bethard, S., Carrell, D. C., Chen, P. J., et al. (2013). Normalization and standardization of electronic health records for high-throughput phenotyping. Journal of the American Medical Informatics Association, 20(e2), e341-e348. https://doi.org/10.1136/amiajnl-2013-001939

[15] Chang, E., Mostafa, J., & Tilahun, B. (2025). Quantitative analysis of the comprehensiveness and granularity of biomedical terminology systems. Journal of the American Medical Informatics Association, 32(3), 489-501. https://www.nature.com/articles/s41598-025-17737-0

[16] Lupse, O. S., & Stoicu-Tivadar, L. (2022). FHIR-based integration of clinical and laboratory data using a microservices architecture. International Journal of Intelligent Systems and Applications in Engineering, 10(3), 285-291. https://www.sciencedirect.com/science/article/pii/S0010482525014039

[17] Lacoste-Julien, S., Palla, K., Davies, A., Kasneci, G., Graepel, T., & Ghahramani, Z. (2013). SiGMa: Simple greedy matching for aligning large knowledge bases. Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 572-580. https://arxiv.org/abs/1207.4525

[18] Dridi, A., Mhamdi, F., & Ben-Abdallah, H. (2024). The Green Computing Infrastructure and Reporting Ontology. Proceedings of the International Conference on Knowledge Management and Information Sharing, 211-218. https://www.scitepress.org/Papers/2026/144917/144917.pdf

[19] Heller, R., Veinot, T. C., & Yarbrough, B. K. (2021). Sustainable health informatics: An emerging research agenda. Journal of the American Medical Informatics Association, 28(11), 2511-2518. https://pmc.ncbi.nlm.nih.gov/articles/PMC7323624/

[20] Muhlbradt, E. E. (2023). NCI-EVS: Building the semantic infrastructure for terminology services. Journal of the Society for Clinical Data Management, 3(2), 1-12. https://doi.org/10.47912/jscdm.213

[21] Kushida, T., Yamamoto, Y., & Yamaguchi, A. (2025). Federated SPARQL query performance evaluation for biomedical research. Frontiers in Bioinformatics, 5, 1485632. https://www.jscdm.org/article/id/134/

[22] HL7 International. (2024). FHIR Terminology Service Module, Release 5. HL7 International. https://www.hl7.org/fhir/terminology-service.html

Declaration

The author has read and agrees to the terms and conditions of the International Journal of Intelligent Systems and Applications in Engineering governing manuscript submission, peer review, copyright, and publication. The author declares no competing financial or personal interests in connection with this work.