

Standardising behavioural health data for interoperable human behaviour measurement

Semantic Web
XX(X):1–27
©The Author(s) 0000
Reprints and permission:
sagepub.co.uk/
journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Theophilus, M.J.^{1,2}, Chua, J.Y.³, Kingsnorth, A.P.^{1,2,4}, Williams, A.³, Sherar, L.B.^{1,2,4}, Esliger, D.W.^{1,2,4,5} and Langford, J.^{3,6}

Abstract

Human behavioural health data have increased over the past decade but remain difficult to combine or reuse because sectors describe behavioural events differently. This study aimed to investigate how behavioural health data can be standardised as machine-actionable, interoperable data objects that support exchange across sectors while remaining compatible with routine workflows. Stakeholder consultations and requirements engineering identified 32 requirements, formalised in a traceability matrix and operationalised through eight competency questions. Semantic Web methods were used to develop a standardised behavioural event data object (Atom) as a JSON Schema, classification model ontologies and registries, and persistent identifiers supporting Linked Data projection into knowledge graphs. Evaluation used two weeks of wrist-accelerometer data from eight participants, processed using two classification methods, to demonstrate cross-model data handling. Controlled error seeding and confusion-matrix metrics assessed Atom validation before downstream querying and aggregation. Structural conformance, event-label integrity, duplicate detection, and temporal-gap detection achieved 100% accuracy and F1 scores, with no false positives across both methods. Direct-identifier screening also produced no false positives but showed lower detection performance (F1 scores = 42.42% and 30.85%). Downstream queries successfully retrieved events within specified time windows, summarised behaviour duration, and aggregated classifications into shared behavioural categories through persistent identifiers and knowledge graphs. The resulting data object and supporting resources provide a practical bridge between human behaviour measurement and the Semantic Web, enabling interoperable and auditable behavioural health data across research, clinical, and public health systems. External evaluation is required to assess implementation, scalability, and governance across independent stakeholder systems.

Keywords

Semantic Web; interoperability; movement behaviour; public health; provenance; wearable technology

¹School of Sport, Exercise and Health Sciences, Loughborough University, Epinal Way, Loughborough LE11 3TU, United Kingdom

²National Centre for Sport and Exercise Medicine (East Midlands), Loughborough University, Epinal Way, Loughborough, United Kingdom

³ActivInsights Ltd., 6 Nene Road, Bickton Industrial Park, Kimbolton, Huntingdon PE28 0LF, United Kingdom

⁴NIHR Leicester Biomedical Research Centre, University Hospitals of Leicester NHS Trust, and the University of Leicester, Leicester, United Kingdom

⁵HealthTech Research Centre for Rehabilitation, NIHR, Nottingham University Hospitals NHS Trust, Nottingham, United Kingdom

⁶Faculty of Health and Life Sciences, University of Exeter, Stocker Rd, Exeter EX4 4PY, United Kingdom

Corresponding author:

Millen James Theophilus, School of Sport, Exercise and Health Sciences, Loughborough University, Epinal Way, Loughborough LE11 3TU, United Kingdom
Email: M.J.Theophilus@lboro.ac.uk

1 Introduction

1.1 Background

The measurement of human behaviour is central to healthcare, public health policy, and behaviour change research, yet the resulting data are rarely represented in standardised ways that support data exchange across health, research, and commercial sectors. Human behaviour encompasses all actions a person takes in response to internal or external stimuli and the intersection with health can be viewed at the individual and population level [1]. At the population level, public health policy is the organised effort to prevent disease, prolong life, and enable individuals to realise their birth right of health and longevity [2]. Public health policy organisations are therefore interested in the longitudinal monitoring of behavioural events to associate them with health outcomes and inform guidance [3]. However, despite substantial global investment in digital health, behavioural health data are still rarely managed in ways that support their effective use across research, clinical, and public health contexts.

Scientific data management guidelines emphasise that digital research assets should be findable, accessible, interoperable, and reusable (FAIR), particularly when intended for use across contexts and sectors [4]. The Semantic Web is an approach to publishing data with explicit, machine-readable meaning so that people and computer systems can interpret and reuse it consistently. Research in the Semantic Web and Internet of Things (IoT) communities has demonstrated the value of ontologies and structured data objects for unifying diverse health and sensor datasets [5, 6]. Recent initiatives such as Mobilise-D demonstrate that harmonised metadata and standardised processing procedures can improve the comparability of mobility data across devices and reference systems [7]. However, a systematic mapping study of Semantic Web approaches in personal health monitoring identified recurring challenges such as context representation, privacy, and uncertainty in data interpretations [8]. Achieving FAIR-by-design behavioural health data therefore requires not only shared semantics but also privacy-preserving architectures and standardised, machine-readable data objects.

International monitoring data from 2023 indicates that countries use less than 5% of available data in health-related decision making, while only 18% report a functioning unique identifier registry and 27% report no standards for data exchange [9]. Progress remains limited due to heterogeneous measurement approaches, including device-based measurements, patient-reported outcomes, and clinical assessments [10, 11]. Fragmented formats, opaque semantics, privacy concerns, and inconsistent governance collectively impede implementation [9, 12, 13]. In health informatics, this challenge is framed in terms of interoperability, defined as the ability of systems to exchange data and use that data effectively [14]. Semantic interoperability ensures that the meaning of data is preserved when shared across systems and that the data can be correctly interpreted for effective use [15]. Practices to maintain data provenance, defined as the detailed record of the source of data and the methods by which it was produced and transformed, supports reliable assessment and underpins confidence when data are reused for clinical research and regulatory-grade evidence [16]. These indicators reflect broader limitations in data availability, identification, and exchange; they do not measure semantic interoperability directly but illustrate infrastructure gaps that constrain it.

Major health system providers invest heavily in electronic health record upgrades but still struggle to integrate data from connected devices or external providers [17]. These limitations persist despite nearly US \$7 billion in annual digital health invested in the United States in 2018, illustrating that substantial expenditure has not translated into interoperable infrastructures. In parallel, evidence from a systematic review suggests that healthcare data provenance remains under-addressed in the literature, with particularly limited attention to regulatory compliance [18]. Taken together, these findings indicate that clinical and behavioural health data are still treated as local records rather than interoperable health assets. And for human behaviour research, they highlight the need for privacy-preserving representations of data that can connect to existing health data infrastructures.

One notable initiative to increase the interoperability of behavioural data is the Organization for the Advancement of Structured Information Standards (OASIS) Classification of Everyday Living (COEL) standard [19]. The COEL v1.0 standard comprises a specification that defines a hierarchical classification taxonomy of activities of daily living (COEL Model), a JavaScript Object Notation (JSON) artefact for representing behavioural health data (COEL Behavioural Atom), and an operational ecosystem with defined roles and responsibilities. It adopts privacy-by-design principles through pseudonymous identifiers and clear separation of directly identifying information from behavioural health data. The COEL v1.0 standard provides a robust foundation but predates many of the current developments in Semantic Web methods and does not yet fully address the requirements for representing behavioural health data as semantically interoperable resources across clinical, public health, and research sectors.

Building on these challenges, our central proposition is that connecting behavioural health data with Semantic Web methods allows fundamental information to be standardised, validated, linked, and reused across research, clinical care, and public health sectors. Shared, machine-readable data objects support human-centric analysis, privacy-aware data sharing, and provide automated decision support. This study translates these cross-sector needs into explicit requirements and machine-actionable artefacts for COEL v2.0.

1.2 Research objectives

The aim of this study is to establish the specification and implementation basis for COEL v2.0 by investigating how behavioural health data can be standardised and structured to achieve semantic interoperability while supporting routine measurement and analytical workflows.

The specific objectives are:

1. To identify and formalise conceptual, technical, and governance requirements for representing behavioural health data in ways that support use by commercial, research, and data-management stakeholders and enable integration with the Semantic Web.
2. To design and specify a digital system that embodies these requirements, comprising behaviour classification ontologies for semantic interoperability and a data object schema that preserves data provenance and privacy-by-design principles.
3. To evaluate whether the digital system enables semantically consistent querying, summary checks, and basic validation of behavioural health data, and provide recommendations to facilitate wider adoption.

By delivering these requirements, specifications, and recommendations, this study provides a proof-of-concept pathway to facilitate development and wider adoption of COEL v2.0 in all sectors including the health and the Semantic Web.

2 Related work

Semantic Web methods and data modelling have played a central role in making heterogeneous data interoperable across sectors. Three overlapping approaches are particularly relevant and have emerged as the most notable over the years: schemas and specifications that define data objects, ontologies that provide formal conceptual models, and Linked Data and knowledge graphs that publish and connect these models using shared identifiers [20]. This section provides an overview of these tools, including the languages required to implement them, and situates generic domain-agnostic and domain-specific ontologies within this context.

At the data structure level, many health and sensor systems use JSON together with a schema to define the shape of data objects. JSON schemas provide a way to specify required and optional fields, data types, and simple constraints for JSON objects, and are widely used to validate data objects in application programming interfaces and data pipelines, but they do not themselves provide a formal conceptual model or support logical inference across datasets. Ontologies provide this conceptual layer by defining classes, relations, and constraints in a logic-based language, most commonly the Web Ontology Language (OWL)¹ serialised in the Resource Description Framework (RDF)² format using a language syntax – Turtle, RDF/XML, N-Triples, or JSON-LD³ [21, 22]. In the context of the Semantic Web, ontologies act as reusable schemas for data integration, sharing, and discovery. Knowledge graphs instantiate these schemas as large, directed, labelled diagrams in which entities and their relationships are represented as RDF triples or equivalent structures, often combining multiple ontologies and data sources to support flexible querying and reasoning [23]. Linked Data refers to the practice of publishing such graphs on the web using identifiers, RDF, and typed links so that datasets can be discovered and traversed across organisational boundaries. Together, these tools create the technical foundation on which ontologies are built.

Several domain-agnostic ontologies have been used within sensor-derived information frameworks in the Internet of Things (IoT). The Semantic Sensor Network (SSN) ontology⁴, with its latest update that uses the Sensor, Observation, Sample, and Actuator (SOSA) ontology⁵, provides generic patterns for describing sensors, observations, features of interest, and procedures across a wide range of IoT applications [24]. The Provenance Ontology (PROV-O)⁶ complements such patterns by representing provenance links between derived outputs and the source evidence used to generate them [25], while the Dublin Core Metadata Initiative (DCMI) Metadata Terms⁷ provides lightweight, record-level metadata such as identifiers and conformance references [26]. Further, Simple Knowledge Organization System (SKOS) mapping relations⁸ provide a mechanism for expressing curated correspondence, for example exact or close matches, between concepts without requiring full logical equivalence [27]. These foundational ontologies enable consistent modelling of measurement processes and interoperable descriptions of sensor-derived data, and they are increasingly used in health and environmental domains [5, 28]. However, they do not define domain-specific taxonomies of human behaviours or data objects for representing behavioural events over time.

Domain-specific tools must therefore be layered on top of these infrastructural vocabularies for modelling behaviour. These tools span three recurring emphases: behaviour concept ontologies that define terminology and relationships, activity event representations that encode what happened to whom, and measurement encoding that structure device or clinical outputs. Table 1 describes and positions representative resources by summarising the primary scope and the key strengths and limitations for interoperable, privacy-aware, and provenance-preserving behavioural event records.

¹<https://www.w3.org/TR/owl2-overview/>

²<https://www.w3.org/TR/rdf11-concepts/>

³<https://www.w3.org/TR/json-ld11/>

⁴<https://www.w3.org/TR/vocab-ssn/>

⁵<https://www.w3.org/ns/sosa/>

⁶<https://www.w3.org/TR/prov-o/>

⁷<https://www.dublincore.org/specifications/dublin-core/dcmi-terms/>

⁸<https://www.w3.org/TR/skos-reference/>

Table 1. Summary of related domain-specific ontologies, taxonomies, and specifications for human behaviour modelling.

Resource - Abbreviation (Type) [Host]	Primary scope	Strengths and limitations
Human Behaviour Change Ontology - HBO [29] (OWL ontology) [BCIO, GitHub]	Describing full range of potentially observable human behaviours for behaviour change intervention research.	Strengths: Broad coverage and formal identifiers; developed and refined through stakeholder input; peer-reviewed process; versioned and openly maintained. Limitations: No event data object or exchange format; no privacy constructs; not designed around sensor evidence systems or device processing workflows.
Physical Activity Concept Ontology - PACO [30] (OWL ontology) [BioPortal]	Standardising heterogeneous physical activity concepts, primarily from questionnaires or scales.	Strengths: Formal structured activity concepts; evaluated for logical and structural consistency; demonstrated automated classification using an OWL reasoner. Limitations: No event data object or time-stamped exchange format; scope is driven by questionnaire-derived descriptions; no privacy or data provenance constructs.
Ontology for assessing Physical Activity and sedentary behaviour - OPA [31] (OWL ontology) [BioPortal]	Characterising physical activity, sedentary behaviour, and the context in which they occur (time, space, social, weather).	Strengths: Explicit context dimensions; designed for standardised description and extensibility; reuses established ontologies where relevant. Limitations: No event data object or exchange format; no privacy or data provenance constructs.
Semantic Mining of Activity, Social, and Health data - SMASH [32] (OWL ontology) [BioPortal]	Integrating data about physical activity, social interaction, and biomedical measures for weight loss intervention research.	Strengths: Integrates social, activity, and biomarker concepts in one modular OWL resource; open and publicly hosted. Limitations: Alpha status; intervention-specific focus; no event data object or exchange format; no privacy or data provenance constructs.
2024 Compendium of Physical Activities [33] (Tabular compendium)	Cataloguing activities by energy cost to support consistent coding and estimation of energy cost across studies in adults.	Strengths: Widely used; large catalogue; measured MET values with references and supporting material. Limitations: Not a formal ontology; no event data object; intended for adults; no privacy-by-design constructs; no data provenance constructs.
Activity Streams 2.0 [34] (JSON/JSON-LD) [W3C]	Representing activities on the web, intended to be used with vocabularies that define activity and object types.	Strengths: W3C Recommendation; extensible vocabulary and supports JSON-LD contexts for Linked Data use. Limitations: No event data object; no behaviour taxonomy; no privacy-by-design constructs; no data provenance constructs.
Open mHealth schemas [35] (JSON) [GitHub]	Standardising mobile and sensor measurement encoding, intended to improve interoperability and support optional alignment to FHIR.	Strengths: Publicly documented schema library; schema-based validation; optional FHIR mappings and profiles to support clinical integration workflows. Limitations: Measurement-centric rather than a general event record; no privacy or data provenance constructs (beyond what may be carried externally or via mappings).
Classification of Everyday Living v1.0 Standard – COEL v1.0 [19] (Taxonomy + JSON schema) [OASIS]	Representing behavioural health data as discrete events using the COEL Model taxonomy and the COEL Behavioural Atom data object.	Strengths: Event data object; behaviour taxonomy for labelling; privacy-by-design constructs (pseudonymisation and role separation). Limitations: No ontology representation of the taxonomy; limited extension mechanism; limited data provenance constructs; operational ecosystem assumptions hinder adoption; taxonomy bound.

BCIO – Behaviour Change Intervention Ontologies; W3C – World Wide Web Consortium; OWL – Web Ontology Language; JSON – JavaScript Object Notation; JSON-LD – JSON for Linked Data; OASIS – Organization for the Advancement of Structured Information Standards; FHIR – Fast Healthcare Interoperability Resources; MET – Metabolic equivalent of task

The analysis of features in Table 1 shows that most domain-specific resources emphasise either behavioural concepts or measurement encoding, but do not provide an event-level behavioural record with privacy and data provenance constructs. In healthcare deployments, data representations must also integrate with broader infrastructure standards that define what is captured, how it is structured, and how it is exchanged. Examples include the NIH Common Data Elements Repository [36], the European Union Rare Disease Platform set of common data elements [37], the Observational Medical Outcomes Partnership common data model [38], and exchange standards such as FHIR by HL7 [39].

Among the domain-specific resources reviewed, the COEL v1.0 standard occupies a distinctive position by combining a hierarchical taxonomy of everyday activities (COEL Model) with a schema (COEL Behavioural Atom) that encodes behavioural events as pseudonymised packets of data [19]. The COEL Model is a four-level hierarchical taxonomy of decreasing granularity: Cluster (highest class) contains Classes, which contain SubClasses, which contains Elements. The COEL Behavioural Atom (also referred to as “Atom” in short) is a single event-level record that encodes a time-bounded behavioural event for a pseudonymous actor. Therefore, the COEL v1.0 standard already addresses several gaps identified in previous work, including a human-centric behaviour model and a privacy-by-design event-level data object. At the same time, the current data object and model remain tied to a JSON artefact without an accompanying ontology and is scarcely adopted in the public health, clinical, and academic sectors.

3 Methods

This study specified and designed an extended COEL digital system⁹ building on the COEL v1.0 standard and evaluated it using a wrist-worn accelerometer case study. Three primary deliverables are provided alongside this manuscript:

- (i) a requirements traceability matrix linking stakeholder needs, competency questions, and verification targets (Supplementary A),
- (ii) the COEL Behavioural Atom v2.0 specification and normative JSON Schema¹⁰ (Supplementary B),
- (iii) a versioned public repository containing the complete set of machine-actionable artefacts and evaluation resources, including registries, ontology serialisations, mapping tables, a JSON-LD context, and R code.

The methods comprise of five stages: (i) stakeholder engagement, (ii) competency questions creation, (iii) requirements engineering and traceability, (iv) artefact and system development, and (v) system evaluation.

3.1 Stakeholder engagement

Design needs were elicited through consultations with three stakeholder groups representing complementary ecosystem roles: a personal data platform provider (Mydex CIC), a digital health technology company (ActivInsights Ltd., UK), and academic researchers (Loughborough University, UK). COEL v1.0 was decomposed into its core components: the interface ecosystem, the hierarchical COEL Model (behaviour taxonomy), and the COEL Behavioural Atom (event record). A structured review of COEL v1.0 identified strengths and limitations relevant to universality, implementation, and adoption. Outcomes from the review and stakeholder discussions were captured in the COEL Behavioural Atom v2.0 specification and consolidated into single testable requirement statements to guide subsequent design decisions.

3.2 Competency questions

Competency questions (CQs) articulate the requirements that a system is expected to satisfy and provide a practical means of testing whether those design requirements are fulfilled [40]. In this study, eight CQs were drafted from stakeholder-informed design needs to define a testable evaluation protocol for the extended COEL digital system. The CQs were designed to be executable over multiple Atoms grouped together (Atom payloads) and assess two classes of capability: validation (structural correctness, semantic usability, governance constraints, and temporal integrity) and scoping (retrieval and summarisation over participants and time windows) [41]. Table 2 lists the CQs and specifies the implementation level. The CQs were additionally used within requirements engineering as explicit sources in the traceability process and as verification targets during evaluation.

Table 2. List of competency questions derived from identified design needs for the extended COEL digital system.

CQ	Category	Competency question	Implementation
CQ1	Structural, validation	Can the system detect and report structural errors in Atom payloads, including missing and incorrectly formatted data for required parts?	JSON
CQ2	Interoperability, validation	Can the system verify that event labels used in Atoms exist in the published model registries, and flag Atoms with unknown or inconsistent labels?	JSON
CQ3	Governance, validation	Can the system check that Atoms contain no directly identifiable personal information and flag those that have?	JSON
CQ4	Governance, validation	Can the system structurally identify how many unique Atoms are in the payload and flag duplicate Atoms?	JSON
CQ5	Temporal, validation	Can the system identify temporal gaps in the payload longer than a specified threshold and identify the adjacent Atoms?	JSON
CQ6	Temporal, scoping	For a single participant and defined time window, can the system retrieve all Atoms that are partially or fully contained within that window?	JSON
CQ7	Summary, scoping	For a single participant, can the system summarise Atom durations by label, aggregated across the full observation period or per day, with stratification by evidence type and classification model?	JSON-LD
CQ8	Extended interoperability, scoping	Can the system aggregate durations across different classification models using declared equivalence mappings into shared higher-level categories?	JSON-LD

CQ – competency question, JSON – JavaScript Object Notation; JSON-LD – JSON for Linked Data; COEL – Classification of Everyday Living

3.3 Requirements engineering and traceability

Traceability was used to maintain an auditable link between stakeholder needs, the resulting design decisions, and the artefacts produced in this study. Following IEEE 830 principles, traceability was treated as a property of the requirements set that permits each requirement to be traced to its origin and to the artefacts created during development [42]. A requirements

⁹<https://w3id.org/coel/>

¹⁰<https://w3id.org/coel/atom/2.0/>

traceability matrix (RTM) was maintained as the primary instrument (Supplementary A). Each RTM entry comprised a unique identifier, a single, testable requirement statement, its source and rationale, the artefacts impacted or produced, and an explicit verification method [43]. Requirement sources were derived from stakeholder engagement, the COEL Behavioural Atom v2.0 specification (Supplementary B), and CQs to operationalise system capabilities as testable queries. Forward traceability was established by mapping each requirement to specific implementation targets. Backward traceability linked each design decision and each evaluation outcome to its originating requirement. This RTM-driven workflow ensured that changes to the Atom specification, registries, and Linked Data projection were justified, reviewable, and verifiable.

3.4 Artefact and system development

Artefact and system development followed a staged workflow that preserves the COEL Model concept hierarchy while making it machine-actionable and aligned with Semantic Web practice. First, the COEL Model¹¹ was normalised into a versioned registry containing stable Internationalized Resource Identifiers (IRIs), human-readable labels, and codes that preserve hierarchy for each possible event classification. The registry was serialised as an OWL ontology (Turtle) with events represented as OWL classes with `rdfs:subClassOf`¹² relations arranged to mirror the original four-level structure and annotated with codes and human-readable labels to support mapping and consistency checks [44]. In parallel, two minimal classification taxonomies used in system evaluation (ActivInsights Behavioural Bout Model and ActivInsights Rest Activity Model)¹³ were curated and published using the same conventions, together with mapping tables linking their classification labels to COEL Model codes using SKOS mapping concepts to support cross-model aggregation [27]. Persistent IRIs were minted using the w3id identifiers¹⁴ [45] for each classification model and behaviour label to support Linked Data querying. Protégé Desktop (Version 5.6.7) was used to visualise the class hierarchy and generate ontology summary statistics [46].

Second, the COEL Behavioural Atom was redesigned based on the requirements captured in the RTM (Supplementary A) and specified in the COEL Behavioural Atom v2.0 specification (Supplementary B). In broad terms, an Atom represents a time-bounded behavioural event linked to a pseudonymous participant and is accompanied by classification and provenance information. The schema requires four top-level *elements* (*Header*, *When*, *Who*, *What*) with several *sub-elements* (*elements.subelements*). *Header* identifies the Atom and related resources. *When* defines the time interval of the event. *Who* records the pseudonymous actor associated with the event. *What* records the behavioural label(s) assigned to the event and their identifier(s). An optional *How* element records how the Atom was derived, including evidence and classification model metadata. Additional optional elements support *Where* (location), *Context* (additional contextual descriptors), *Consent* (governance metadata), and *Extension* (implementation-specific fields such as derived measures). JSON was retained as the normative exchange representation to preserve adoption across non-semantic infrastructures.

Third, a JSON-LD context file for the normative JSON data object was created to support Linked Data style querying [47]. The export workflow performs a lightweight JSON-LD projection in which the original Atom JSON structure is preserved, each Atom is assigned a locally minted URN identifier (@id) derived from *Header.AtomID*, and is wrapped in a JSON-LD document comprising a user-supplied or default @context and an @graph of Atoms. Semantic interpretation is therefore delegated to the context file. In the study context, linkage was achieved by aligning Atom fields that carry IRIs to resolvable persistent identifiers, including *Header.AtomIRI* (conformance), *What.LabelIRI* (behaviour concept), and *How.ClassificationModelIRI* (classification model), using a minimal vocabulary layer based on four established vocabularies: XML Schema Definition Language¹⁵ [48], PROV-O [25], DCMI Metadata Terms [26], and SKOS [27]. The resulting JSON-LD projection preserved the original nested JSON structure while enabling Atoms to be expanded into RDF graphs and queried using the SPARQL query language¹⁶ within the R environment [49, 50]. The JSON-LD representation is treated as an optional implementation artefact for export and analysis and does not replace the normative COEL Behavioural Atom v2.0 JSON data object specified in Supplementary B.

3.5 System evaluation

Accelerometer data were obtained from an internal ActivInsights dataset comprising eight anonymised, two-week participant recordings collected between 6 June 2025 and 26 June 2025. Each binary file contained raw tri-axial acceleration sampled at 100 Hz together with concurrent light and ambient temperature measurements, recorded using the GENEActiv 1.2 device (ActivInsights Ltd., UK), a wrist-worn micro electro mechanical system accelerometer specified with a ± 8 g range, 12-bit resolution (3.9 milli-g), configurable sampling at 10 to 100 Hz, and a quartz real-time clock of ± 20 ppm. Data were collected by ActivInsights employees to support research for technology and service adoption. Participation was optional and employees were informed of the intended uses and consented to the data collection. Secondary analysis of the anonymised dataset for this study was approved by ActivInsights via written permission.

Each binary file was processed in RStudio (Version 2025.09.2+418) [51] using GENEAcCore (Version 1.1.3) [52] to derive event-level outputs, which were exported as COEL Atoms using an Atom JSON builder R function within the extended COEL

¹¹<https://w3id.org/coel/models/coel/2.0/>

¹²<https://www.w3.org/TR/rdf-schema/>

¹³<https://w3id.org/coel/models/activinsights/>

¹⁴<https://w3id.org/>

¹⁵<https://www.w3.org/TR/xmlschema11-2/>

¹⁶<https://www.w3.org/TR/sparql11-query/>

digital system. Each Atom encoded a time-bounded interval linked to a pseudonymous participant identifier and evidence metadata, with multiple classification labels included for cross-model analyses. The classification algorithms implemented were ActivInsights' methods intended for open-access release via GENEAcare and are supplied here as artefacts of the extended COEL digital system.

Evaluation comprised two layers aligned to the CQs. Validation CQs (CQ1–CQ5) were assessed using an error-seeding and detection pipeline. For each CQ-focused run, a reproducible random 10–30% subset of Atoms was perturbed using an injector that applied one enabled error type per selected Atom (field removals, data type changes, invalid labels, injected personal names, duplicates, or temporal gaps). All injected modifications were recorded in a changelog, which served as ground truth. The seeded payload was then passed to a validator configured to run only the test(s) associated with the target CQ, producing an Atom-level violations table. Appendix Table 1 presents the description of each error type seeded and the test performed for each validation CQ assessed. Performance was quantified at the Atom level using confusion-matrix metrics (True Positives, False Positives, False Negatives, True Negatives) and derived accuracy and F1-score by direct comparison with the changelog. For CQ3 (direct identifier screening), screening was implemented as a conservative name-token heuristic intended to minimise false positives and is therefore evaluated as a governance risk indicator rather than a comprehensive privacy safeguard. Baseline payloads were also validated without injection across CQ1–CQ5 to confirm zero violations prior to downstream scoping analyses.

Scoping CQs (CQ6–CQ8) were evaluated on baseline payloads using executable R functions. For CQ6, Atoms were retrieved for a specified participant and time window using an interval-overlap query (start < window end and end > window start), returning a tabular extract (*Header.AtomID*, start and end time, duration, label, overlap type) and an optional timeline visualisation. For CQ7, durations were aggregated by the *What.Label* with optional stratification by *How.EvidenceType* and *How.ClassificationModel*. For CQ8, the roll-up to COEL categories was performed exclusively via SKOS mapping triples. Baseline JSON Atom payload files were exported to JSON-LD, with export-time augmentation adding Atom URNs (@id), *EndTimeISO* derived from *TimeUTC* and *Duration*, and URNs for participant and evidence. JSON-LD payloads were then used to aggregate per-day durations across classification models into shared COEL categories by referencing the IRIs for *What.Label* with published mapping tables and resolving COEL codes.

4 Results

This section first reports the requirements and traceability outcomes (Objective 1). It then describes the development and implementation of the extended COEL digital system (Objective 2). Finally, it evaluates querying, validation, interoperability, and routine measurement and analytical use cases (Objective 3).

4.1 Requirements and traceability (Objective 1)

The RTM contains 32 requirements (RTM-001 to RTM-032) that link stakeholder-informed design needs and evaluation targets to concrete artefacts produced in this study and to explicit verification targets. All 32 requirements are implemented within this study and were clustered into three themes (Table 3). Conceptual requirements address stable representation of COEL concepts and identifiers to support adoption across systems. Technical requirements operationalise these needs through machine-actionable artefacts, including registries, ontology serialisations, JSON Schema constraints, mapping resources, and evaluation tooling. Governance requirements specify privacy-by-design constraints and associated checks. Relative to COEL v1.0, most requirements represent additions or tightened constraints needed for automated validation and interoperability: 20 requirements added new capability, 9 changed v1.0 behaviour or constraints, and 3 explicitly preserve v1.0 structures (Table 3).

Table 3. Requirements traceability matrix (RTM) theme summary (full RTM in Supplementary A).

Theme	Requirements (n)	Change type vs COEL v1.0	Primary artefact focus (count with theme)	Primary verification targets
Conceptual	9	Preserve: 3, Change: 6, Add: 0	Atom schema and payload (5); Model registry and ontology (1); JSON-LD projection (1); Registries (1); Identifiers and IRIs (1)	CQ1, CQ2, CQ5, CQ7, CQ8
Technical	20	Add: 19, Change: 1, Preserve: 0	Evaluation tooling (9); Model registry and ontology (4); Registries (2); Atom schema and payload (1); JSON-LD projection (1); Mappings (1); Publishing and hosting (1); Identifiers and IRIs (1)	CQ1, CQ2, CQ4, CQ5, CQ6, CQ7, CQ8
Governance	3	Add: 1, Change: 2, Preserve: 0	Atom schema and payload (1); Identifiers and IRIs (1); Evaluation tooling (1)	CQ1, CQ3, CQ4

CQ – competency question, RTM - requirements traceability matrix, JSON-LD – JavaScript Object Notation Linked Data, IRIs – Internationalized Resource Identifiers

Verification targets for the RTM entry and requirement theme were linked to CQs. CQ1 (structural conformance) and CQ2 (label integrity) account for most RTM-to-CQ links (Appendix Table 2), reflecting the emphasis on schema conformance and registry-grounded identifiers for classification labels. CQ5 to CQ7 extend verification to temporal integrity

and duration summaries, while CQ8 acts as the main interoperability demonstrator by testing cross-model aggregation via declared mappings. Several requirements map to more than one competency question; Table 3 therefore summarises unique requirements per theme with its primary artefact focus and verification targets, while Appendix Table 2 summarises RTM-to-CQ links rather than unique requirements.

4.2 Digital system specification and implementation (Objective 2)

The extended COEL digital system was implemented as a set of versioned, machine-actionable artefacts supporting both non-semantic exchange workflows and semantic-web style interoperability. Core deliverables include (i) a versioned COEL Model registry and OWL serialisation preserving the four-level hierarchy, (ii) registries and OWL serialisations for the two ActivInsights classification taxonomies used in the evaluation, (iii) mapping resources linking the minimal classification labels to COEL Model codes, (iv) a COEL Behavioural Atom v2.0 JSON Schema and extension registry, and (v) example data and coding utilities including a JSON-LD context enabling projection of Atoms to RDF for analysis.

4.2.1 Model ontologies

The COEL Model was serialised as an OWL ontology and published under a versioned namespace. The ontology represents the COEL taxonomy with each COEL activity of daily living represented as owl:Class entities arranged in an asserted rdfs:subClassOf hierarchy and annotated with rdfs:label. The COEL Model ontology comprises 32 Clusters, 197 Classes, 1054 SubClasses, and 5,047 Elements, resulting in 6,330 unique behavioural event classification labels with unique IRIs for cross-referencing. Figure 1 illustrates the four-level hierarchy using the “Sleep” Cluster of the COEL Model. A top-level view of the 32 COEL Model Clusters and an expanded path within one Cluster is provided in Appendix Figure 1.




Figure 1. Protégé generated OntoGraf rendering of the COEL Model expanded for Cluster:Sleep, Class:Metrics, SubClass:Sleep stages, illustrating four-level branching via rdfs:subClassOf relations. + indicates nodes that can be expanded further.

The two ActivInsights behaviour classification models (Behavioural Bout Model and Rest Activity Model) were also serialised as lightweight OWL ontologies and published under versioned namespaces using the same annotation conventions as the COEL Model. Each model represents its behavioural state space as an asserted rdfs:subClassOf hierarchy of owl:Class entities with human-readable rdfs:label annotations, supporting consistent registry alignment and cross-model mapping. Protégé OntoGraf renderings (Figure 2 and Figure 3) illustrate all branches of each model, confirming that the intended label sets and hierarchical groupings are preserved in the OWL serialisation and are structurally compatible with the COEL Model.




Figure 2. Protégé generated OntoGraf rendering of the Behavioural Bout Model v1.0 expanded, illustrating three-level branching via `rdfs:subClassOf` relations.




Figure 3. Protégé generated OntoGraf rendering of the Rest Activity Model v1.0 expanded, illustrating single-level branching via `rdfs:subClassOf` relations.

The ontology summary metrics show that all three model ontologies were intentionally modelled as lightweight, taxonomy-oriented OWL class hierarchies (Table 4). The COEL Model ontology is dominated by annotations, declarations, and logical axioms comprising `rdfs:subClassOf` axioms, with no object properties, data properties, or individual axioms, and only a minimal annotation layer to support labels and registry alignment. The Behavioural Bout Model and Rest Activity Model follow the same pattern at smaller scale, comprising compact class sets with correspondingly low axiom counts and no asserted properties or individuals.

Table 4. Ontology summary metrics for the COEL Model, Behavioural Bout Model, and Rest Activity Model ontologies.

Ontology metrics	COEL Model Count	Behavioural Bout Model Count	Rest Activity Model Count
Total axioms	25,323	45	20
Annotation axioms	12,661	27	12
Declaration axioms	6,332	10	5
Logical axioms	6,330	8	3
Object property axioms	0	0	0
Data property axioms	0	0	0
Individual axioms	0	0	0
Class count	6,331	9	4
SubClass Of	6,330	8	3
Annotation Property count	2	3	3

Axiom refers to fundamental rules underlying the ontology, which can be annotation, declaration, logical, object property, data property, or individual rules. Class count refers to the number of `owl:Class` entities, not the number of COEL Classes in the four-level COEL Model. The SubClass Of count is the number of entities with `rdfs:subClassOf` relations.

4.2.2 COEL Behavioural Atom v2.0

The COEL Behavioural Atom v2.0 specification with the normative JSON Schema are provided in Supplementary B, and additionally in the extended COEL digital system. Relative to Atom v1.0, Atom v2.0 generalises the actor to a participant or environment and shifts event labelling from numeric COEL Model categories to multi-model labelling. The revision additionally updates the schema baseline, tightens temporal representation, and formalises evidence and model provenance via explicit evidence descriptors and classification model identifiers. Lastly, Atom v2.0 introduces optional Atom-level identifiers, updates consent metadata to the ISO-27560 standard [53] and replaces the tag-based extension mechanism with an open, registry-backed Extension design. Appendix Table 3 summarises the key changes and rationale for the changes between

COEL Behavioural Atom v1.0 and v2.0. Appendix Figure 2 illustrates the Atom v2.0 schema using an instance generated from the Behavioural Bout classification algorithm as an example. Appendix Figure 3 presents the JSON-LD context file used to project Atoms from the normative JSON representation to JSON-LD.

4.3 System evaluation (Objective 3)

Across eight participants' 14-day recordings (2,688 hours), the Atom payload generated using the Rest Activity classification model produced 230 Atoms (median Atom duration 9.3 hours) and the Behavioural Bout classification model produced 116,305 Atoms (median Atom duration 35 seconds) (Table 5). Mapping coverage for the Atom payloads was <100% as the "NonWear" classification label was intentionally left unmapped to the COEL Model. The Behavioural Bout Atom payload comprised 1.86 megabytes of data per day and the Rest Activity Atom payload comprised 2 kilobytes of data per day. Gzip compression reduced file size by 94-96% across payloads, equivalent to a 17- to 27-fold size reduction.

Table 5. Atom payload summary data with mapping coverage to COEL Model activities of daily living.

Atom payload	Atoms (n)	Partici-pants (n)	Event labels (n)	COEL codes (n)	COEL Mapping Coverage (%)	Total dur (h)	JSON (MB)	Gzip (MB)
Rest Activity	230	8	3	2	66.7	2688	0.27	0.01
Behavioural Bout	116,305	8	7	6	85.7	2688	208.33	12.53

Mapping coverage is calculated over unique event labels relative to the curated mapping to the COEL codes.

4.3.1 Validation

Validation CQs (CQ1 to CQ5) were executed on the baseline Atom payloads and on controlled error seeded variants to test detection of structural, semantic, governance, duplicate, and temporal integrity violations. For CQ1 (structural conformance), CQ2 (label integrity), CQ4 (duplicate detection), and CQ5 (temporal-gap detection), all seeded violations were detected in both payloads, yielding perfect classification performance at the Atom level (Table 6). In contrast, CQ3 (direct identifier screening) showed materially lower recall while preserving zero false positives (Appendix Table 4), consistent with a conservative heuristic design that prioritises avoiding incorrect flagging at the expense of missing some injected name tokens. Baseline payloads showed zero CQ1 to CQ5 violations, confirming conformance prior to scoping analyses.

Table 6. Summary of validation performance by CQ and payload (Atom-level; error-seeded runs).

CQ	Capability	Rest Activity (F1, %)	Behavioural Bout (F1, %)	Key observation
CQ1	Structural conformance	100	100	Perfect detection, no false positives
CQ2	Label integrity	100	100	Perfect detection, no false positives
CQ3	Direct identifier screening	42.4	30.9	Lower recall, conservative heuristics
CQ4	Duplicate detection	100	100	Perfect detection, no false positives
CQ5	Temporal-gap detection	100	100	Perfect detection, no false positives

CQ – competency question

4.3.2 Scoping

Scoping CQs (CQ6-CQ8) were executed on baseline (non-error seeded) JSON Atom payload files to demonstrate window-based retrieval, duration summarisation, and cross-model aggregation using declared mappings. Scoping outputs are illustrated for a single participant. For CQ6 (window-based retrieval), an executable interval-overlap query retrieved all Atoms directly from each baseline payload file whose event intervals intersected the query window. For visualisation, results are shown for a 48-hr window to make individual event durations interpretable within a dense event stream. Retrieved Atoms were returned as a tabular extract (AtomID, start time, end time, duration, label, evidence type, and classification model) and visualised as a timeline in which each horizontal segment represents one Atom spanning its start and end times, with dashed vertical lines denoting the query window boundaries and small vertical offsets added to separate overlapping, short-duration bouts at this plotting scale (Figure 4 and Figure 5). The same utility additionally supports export of the retrieved Atom subset and the associated table for reproducible downstream analysis. Across the Rest Activity and Behavioural Bout Atom payloads, CQ6 retrieval behaved as expected given the underlying classification models. The Rest Activity model produced long, contiguous intervals consistent with a low-granularity behavioural state representation (Figure 4), whereas the Behavioural Bout model produced many short intervals and higher label granularity within the same 48-hour window (Figure 5). The next section covers CQ7 (duration summarisation) and CQ8 (cross-model aggregation), demonstrated on the JSON-LD projection of the Atoms.




Figure 4. Rest Activity timeline retrieval of Atoms for a single participant over a 48-hour query window beginning at the start time of that participant’s first recorded Atom. Each horizontal segment represents one Atom, positioned by its label and spanning its start and end times; dashed vertical lines denote the retrieval window boundaries.





Figure 5. Behavioural Bout timeline retrieval of Atoms for a single participant over a 48-hour query window beginning at the start time of that participant’s first recorded Atom. Each horizontal segment represents one Atom, positioned by its label and spanning its start and end times; dashed vertical lines denote the retrieval window boundaries. At this plotting scale, dense short-duration bouts can appear as points; small vertical offsets separate overlapping segments.

4.3.3 Semantic interoperability with Linked Data

Atoms were projected to JSON-LD using the accompanying context file (Appendix Figure 3), and JSON-LD expansion produced an RDF graph in which Atom instances are identified by export-minted URNs and linked to model label IRIs via `dcterms:subject`. These model concepts were then connected to COEL v2.0 concepts via curated SKOS mappings (Table 7). Figure 6 illustrates an example of the semantic equivalence pathway used for cross-model roll-up for the COEL label “Active”.

Table 7. Simple Knowledge Organization System (SKOS) mappings for semantic equivalence of event labels derived from Rest Activity and Behavioural Bout classification models and COEL activities of daily living labels.

Model and label	COEL Code	COEL Label	Mapping to COEL
Rest Activity			
Rest	13.1	Asleep	skos:broadMatch
Wake	14.9.3.2	Active	skos:broadMatch
NonWear	-	-	Unmapped
Behavioural Bout			
Sedentary	14.9.3.1	Sedentary	skos:relatedMatch
Active	14.9.3.2	Active	skos:relatedMatch
SlowWalk	14.5	Walking	skos:broadMatch
FastWalk	14.5	Walking	skos:broadMatch
Run	14.4	Running	skos:relatedMatch
Sleep	13.1	Asleep	skos:broadMatch
NonWear	-	-	Unmapped

**Figure 6.** Example semantic equivalence pathway for cross-model roll-up to COEL model categories. Atom instances reference native model label concepts via `dcterms:subject`; native model concepts are linked to COEL v2.0 concepts via SKOS mapping predicates (for example, `skos:relatedMatch` and `skos:broadMatch`).

Although Atoms may embed multiple label IRIs (including COEL codes) to support immediate COEL-coded reporting, CQ7 and CQ8 are presented as a semantic mapping demonstration in which COEL categories are obtained solely through SPARQL graph traversal over RDF-based SKOS equivalence mappings and visualised in Figure 7. A native-label semantic summary was computed by SPARQL aggregation over the RDF projection (Appendix Figure 4), and the mapping-based roll-up was computed by SPARQL aggregation after SKOS traversal to COEL concepts (Appendix Figure 5). Atom instances reference native model label concepts via `dcterms:subject` and are linked to COEL v2.0 concepts using SKOS mapping predicates (for example, `skos:relatedMatch` and `skos:broadMatch`).

Behavioural Bout stream — participant]%@v55o:ae!l01jSGV.; first 48-hour window (observed span 48.1 h)




Figure 7. Semantic summaries computed from the JSON-LD projection for the Behavioural Bout stream (single participant; first 48-hour window). Panel A (left): CQ7 native-label duration summary (SPARQL aggregation by native label IRI). Panel B (right): CQ8 roll-up to COEL categories via SKOS traversal and SPARQL aggregation.

5 Discussion

This study specifies and delivers on requirements for representing behavioural health data in a FAIR format to support exchange across sectors and different behaviour classification methods. The primary contribution is an integrated set of models, data specifications, and implementation workflows demonstrating how behavioural events can be standardised from routine measurement methods, validated and inspected, and projected to Linked Data for downstream Semantic Web integration using knowledge graphs. Three primary outputs are provided: (i) ontology serialisations of two minimal behaviour classification taxonomies and the COEL Model, (ii) a COEL Behavioural Atom v2.0 JSON schema for representing behavioural events, with a JSON-LD context file for Linked Data projection, and (iii) a workflow implementation that operationalises validation, data integrity checks, scoping inspection, and Linked Data export. Together, these outputs present COEL v2.0 as a digital system comprising a versioned public repository and persistent identifiers for accessing and linking its machine-actionable artefacts. A requirements traceability matrix linking stakeholder-informed needs to implementation targets is provided as a reporting instrument and living traceability asset for future updates [42] (Supplementary A).

Prior Semantic Web and digital health efforts have typically progressed along two partially disconnected tracks: ontology-driven conceptual models of behaviour [29, 30], and exchange-oriented standards for healthcare information exchange and interoperability [54]. Applied integration studies for healthcare and wearable devices show that bridging these tracks often requires explicit alignment and mapping across different evidence systems to enable unified access and reuse [5]. The present study advances this complementary operational bridge using a dual artefact strategy: a behaviour-first event data object that remains compatible with routine measurement and analytics pipelines while enabling Semantic Web integration, when needed. This dual representation aligns with a layered systems view of semantic health monitoring, where semantics are often introduced as an integration and reasoning layer rather than replacing existing data-capture and data-exchange formats [8]. Evaluation showed that queries and validation could be executed directly over both evaluated JSON payloads (230 Rest Activity Atoms and 116,305 Behavioural Bout Atoms) while retaining provenance information needed to support downstream clinical and regulatory use [16, 18].

In this study, we specify and demonstrate the COEL Behavioural Atom v2.0 as a stable, provenance-aware data object that can carry behavioural events of any duration produced by any upstream evidence system. Because the Atom encodes time-bounded events, it can represent short and longer behavioural events that may be interpreted as behavioural bouts. This work intentionally avoids specifying a formal behavioural bout definition because bout separation will depend on the upstream evidence and analysis pipeline [55]. Instead, the contribution is a stable, provenance-aware data object that can carry behavioural events produced by different measurement and analytical methods, alongside validation and integrity checks for governance, and a Linked Data projection workflow for Semantic Web integration. Privacy-by-design is treated as an implementer responsibility rather than a guarantee of the data object, since governance depends on deployment context and the handling of sensitive metadata. In practice, governance must also consider that provenance and operational metadata can contain sensitive information and require protection strategies beyond simple string screening [18].

The contributions of this study address gaps repeatedly reported in device-based personal health monitoring systems, where heterogeneity of sensor data and limited Semantic Web integration hinder interoperability and downstream use across settings [8]. We demonstrated a standardised method to structure behavioural events that remains compatible with routine JSON-centric production workflows while enabling Semantic Web integration through persistent identifiers, registry-grounded labels, and a projection for Linked Data querying. This approach supports reuse of behavioural event records across systems

and classification models without requiring a single universal behaviour classification taxonomy, since different classification outputs can remain traceable in their native labels and mappings enable shared COEL Model-coded summaries for comparison and synthesis. There is also a clearer separation between the record of “what happened” (time-bounded events) and the contextual metadata required for audit and reuse, which directly targets the provenance shortcomings noted for behavioural monitoring systems [18].

Building on this separation of event representation, provenance metadata, and optional Semantic Web integration, the remaining adoption challenge is less about adding new core constructs and more about governing the resources across real deployments. Stakeholder adoption is therefore facilitated by treating COEL v2.0 as a set of versioned, machine-actionable resources with persistent identifiers, and by aligning next development steps to the responsibilities of stakeholder roles:

- Digital health technology providers: implement Atom export at the processing stage, run validation queries as a release gate, and publish registries alongside software and algorithm releases.
- Data platform providers and custodians: ingest Atoms as immutable data objects, enforce conformance, and support harmonising behaviour labels.
- Researchers and analysts: use Atoms as an implementation bridge between sensor-level data and movement behaviour endpoints, store provenance information, and publicly publish Atoms generated from a study to make data findable and accessible.
- Public health and surveillance users: use Atoms as a lower-burden transfer format than sensor-level data to reduce storage and network costs, while retaining provenance for auditable, reproducible summaries, and methods updates when algorithms change.
- Standards organisations and regulators: provide formal status and long-term governance for change control, stewardship, and conformance certification, so implementers can confidently use COEL v2.0.
- Commercial sponsors: require COEL-aligned, version-pinned artefacts and reproducible COEL-coded summary outputs in contracts and data-sharing plans.
- Insurers and payers: request COEL-coded indicators plus basic provenance fields for audit, avoiding raw data access unless essential.

For clinical patients and the public generally, the primary implication is indirect but important: standardised, micro-structured, behavioural event data, combined with deployment-appropriate privacy controls and explicit provenance, increase the feasibility of trustworthy secondary data use, supporting confident participation in research and scalable population monitoring of movement behaviour.

Several limitations constrain generalisability and should be made explicit. First, evaluation used a small internal dataset, limiting inference across broader populations and alternative devices. Second, the validation assessment relied on controlled error injection, which strengthens internal validity for the targeted defect classes but does not capture the full range of real-world data issues such as partial file corruption, systematic clock drift, or mixed time zones. Third, the privacy screening demonstration was intentionally restricted to personal-name detection and therefore cannot be interpreted as comprehensive privacy protection. In practical deployments, privacy-by-design remains an implementer responsibility, and residual false negatives or false positives are expected given the diversity of identifiers, languages, and contexts. Fourth, cross-model aggregation depends on curated mapping tables, so interoperability claims are conditional on mapping completeness, quality control, and version governance. Finally, the Linked Data projection was demonstrated as an optional analysis workflow rather than as a fully deployed cross-system infrastructure, so claims regarding end-to-end interoperability should be treated as enabled capability demonstrated within the study implementation. The evaluation assessed the implementation machinery, including schema conformance, validation, querying, and mappings, but did not assess whether the COEL Model taxonomy is complete or fit for all behavioural-health applications; taxonomy fitness-for-purpose therefore remains for external, domain-specific evaluation.

6 Conclusion

Behavioural health data remain difficult to combine, check, and reuse because measurement methods are heterogeneous, formats are fragmented, and semantics are inconsistent for data exchange. This study developed and evaluated a standardised approach for representing behavioural events in a structure and format that remains usable in routine data workflows while also supporting Semantic Web integration. The evaluation showed that these data objects can be validated, summarised, and aggregated across methods while retaining the necessary information for transparent and open records. These findings address the heterogeneity and interoperability gaps by demonstrating that behavioural health data can be standardised and machine-audited. This work therefore provides a practical proof-of-concept step toward semantic interoperability in behavioural health data and exchange across research, clinical, public health, commercial, and data-management settings.

7 Declarations

7.1 Ethics approval and consent to participate

Data were collected by ActivInsights employees to support research for technology and service adoption. Participation was optional and employees were informed of the intended uses and consented to the data collection. Secondary analysis of the anonymised dataset for this study was approved by ActivInsights via written permission.

7.2 Availability of data and materials

The dataset and code supporting the findings of this study are made openly available in the COEL v2.0 GitHub repository versioned via a Zenodo record (<https://doi.org/10.5281/zenodo.22028457>).

7.3 Competing interests

JYC and AW are employees of ActivInsights Ltd. JL is Chief Technology Officer and co-founder of ActivInsights Ltd. The devices used in the study and the anonymised employee dataset used for system testing were provided through ActivInsights Ltd. JL was also an original contributor to the OASIS COEL v1.0 standard. MJT, APK, LBS, and DWE declare no competing interests.

7.4 Funding

The Learning Network for Advanced Behavioural Data Analysis (LABDA) project received funding from the European Union's Horizon Europe Research and Innovation Programme under grant agreement No 101072993. The funder had no role in the design and conduct of the study.

7.5 Author contributions

MJT – Conceptualization; Methodology; Software; Validation; Formal analysis; Investigation; Data curation; Visualization; Writing – original draft; Writing – review and editing; Project administration.

JYC – Software; Writing – review and editing; Resources.

APK – Methodology; Writing – review and editing; Resources; Supervision.

AW – Methodology; Resources; Supervision.

LBS – Methodology; Writing – review and editing; Resources; Supervision.

DWE – Methodology; Writing – review and editing; Resources; Supervision.

JL – Conceptualization; Methodology; Software; Resources; Writing – review and editing; Supervision; Project administration.

7.6 Acknowledgements

The authors thank David Alexander from Mydex CIC for contributing stakeholder perspectives during the discussions that informed this study.

7.7 Additional files

Supplementary A: Requirements traceability matrix

Supplementary B: COEL Behavioural Atom v2.0 specification

References

- [1] Davis R, Campbell R, Hildon Z, et al. Theories of behaviour and behaviour change across the social and behavioural sciences: a scoping review. *Health Psychol Rev* 2015; 9(3): 323–344. DOI: 10.1080/17437199.2014.941722.
- [2] Ahmed FU. Defining public health. *Indian J Public Health* 2011; 55(4): 241–245. DOI: 10.4103/0019-557X.92397.
- [3] World Health Organization. Physical activity, <https://www.who.int/news-room/fact-sheets/detail/physical-activity> (2024, accessed 20 August 2026).
- [4] Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 2016; 3: 160018. DOI: 10.1038/sdata.2016.18.
- [5] Hodges J, Kritzler M, Michahelles F, et al. Ontology alignment for wearable devices and bioinformatics in professional health care. Manuscript submitted for publication, *Semantic Web Journal*, 2014. <https://www.semantic-web-journal.net/system/files/swj826.pdf> (accessed 20 August 2026).
- [6] Ruta M, Scioscia F, Loseto G, et al. Machine learning in the Internet of Things: a semantic-enhanced approach. *Semantic Web* 2018; 10(1): 183–204. DOI: 10.3233/SW-180314.
- [7] Palmerini L, Reggi L, Bonci T, et al. Mobility recorded by wearable devices and gold standards: the Mobilise-D procedure for data standardization. *Sci Data* 2023; 10(1): 38. DOI: 10.1038/s41597-023-01930-9.
- [8] Nzomo M and Moodley D. Semantic Web technologies in sensor-based personal health monitoring systems: a systematic mapping study. *Semantic Web* 2026; 17(3). DOI: 10.1177/22104968261431426.
- [9] Chukwu E, Mechael P, Edelman JK, et al. *The state of digital health 2023*. Global Digital Health Monitor, 2023. https://digitalhealthmonitor.org/s/State-of-Digital-Health_2023.pdf (accessed 20 August 2026).
- [10] Cohen AB and Mathews SC. The digital outcome measure. *Digit Biomark* 2018; 2(3): 94–105. DOI: 10.1159/000492396.
- [11] Taylor KI, Staunton H, Lipsmeier F, et al. Outcome measures based on digital health technology sensor data: data- and patient-centric approaches. *NPJ Digit Med* 2020; 3: 97. DOI: 10.1038/s41746-020-0305-8.
- [12] Bestek M. *Democratizing health data semantics: a commons-based technology-enhanced activity space to support productive work with health data semantics*. Doctoral dissertation, IT University of Copenhagen, 2022. <https://pure.itu.dk/en/publications/democratizing-health-data-semantics-a-commons-based-technology-en/> (accessed 20 August 2026).
- [13] Torab-Miandoab A, Samad-Soltani T, Jodati A, et al. Interoperability of heterogeneous health information systems: a systematic literature review. *BMC Med Inform Decis Mak* 2023; 23: 18. DOI: 10.1186/s12911-023-02115-5.
- [14] Institute of Electrical and Electronics Engineers. *IEEE standard computer dictionary: a compilation of IEEE standard computer glossaries*. IEEE, 1990. DOI: 10.1109/IEEESTD.1991.106963.
- [15] Mollerus F, Lynch C and Bruining H. Data interoperability for a systems approach to developmental conditions. *Neurosci Biobehav Rev* 2025; 176: 106245. DOI: 10.1016/j.neubiorev.2025.106245.
- [16] Murray ML, Love SB, Carpenter JR, et al. Data provenance and integrity of health-care systems data for clinical trials. *Lancet Digit Health* 2022; 4(8): e567–e568. DOI: 10.1016/S2589-7500(22)00122-4.
- [17] Shull JG. Digital health and the state of interoperable electronic health records. *JMIR Med Inform* 2019; 7(4): e12712. DOI: 10.2196/12712.
- [18] Ahmed M, Dar AR, Helfert M, et al. Data provenance in healthcare: approaches, challenges, and future directions. *Sensors (Basel)* 2023; 23(14): 6495. DOI: 10.3390/s23146495.
- [19] Bruton P, Langford J, Reed M, et al. (eds). *Classification of Everyday Living Version 1.0*. OASIS Committee Specification 01. OASIS Open, 2018. <http://docs.oasis-open.org/coel/COEL/v1.0/cs01/COEL-v1.0-cs01.html> (accessed 20 August 2026).
- [20] Hitzler P. A review of the Semantic Web field. *Commun ACM* 2021; 64(2): 76–83. DOI: 10.1145/3397512.
- [21] Gruber TR. A translation approach to portable ontology specifications. *Knowl Acquis* 1993; 5(2): 199–220. DOI: 10.1006/knac.1993.1008.
- [22] Hitzler P and Van Harmelen F. A reasonable Semantic Web. *Semantic Web* 2010; 1(1–2): 39–44. DOI: 10.3233/SW-2010-0010.
- [23] Hogan A, Blomqvist E, Cochez M, et al. Knowledge graphs. *ACM Comput Surv* 2021; 54(4): 1–37. DOI: 10.1145/3447772.
- [24] Compton M, Barnaghi P, Bermudez L, et al. The SSN ontology of the W3C Semantic Sensor Network Incubator Group. *J Web Semant* 2012; 17: 25–32. DOI: 10.1016/j.websem.2012.05.003.

- [25] Lebo T, Sahoo S and McGuinness D (eds). *PROV-O: The PROV Ontology*. World Wide Web Consortium, 2013. <https://www.w3.org/TR/prov-o/> (accessed 20 August 2026).
- [26] DCMi Usage Board. *DCMI Metadata Terms*. Dublin Core Metadata Initiative, 2020. <https://www.dublincore.org/specifications/dublin-core/dcmi-terms/> (accessed 20 August 2026).
- [27] Miles A and Bechhofer S (eds). *SKOS Simple Knowledge Organization System reference*. World Wide Web Consortium, 2009. <https://www.w3.org/TR/skos-reference/> (accessed 20 August 2026).
- [28] Esnaola-Gonzalez I, Bermúdez J, Fernandez I, et al. Ontologies for observations and actuations in buildings: a survey. *Semantic Web* 2020; 11(4): 593–621. DOI: 10.3233/SW-200378.
- [29] Schenk PM, West R, Castro O, et al. An ontological framework for organising and describing behaviours: the Human Behaviour Ontology. *Wellcome Open Res* 2025; 9: 237. DOI: 10.12688/wellcomeopenres.21252.2.
- [30] Kim H, Mentzer J and Taira R. Developing a Physical Activity Ontology to support the interoperability of physical activity data. *J Med Internet Res* 2019; 21(4): e12776. DOI: 10.2196/12776.
- [31] Silva P, Andrade MT, Carvalho P, et al. A structured and flexible language for physical activity assessment and characterization. *J Sports Med (Hindawi Publ Corp)* 2013; 2013: 420916. DOI: 10.1155/2013/420916.
- [32] National Center for Biomedical Ontology. SMASH ontology, Social Activity Ontology. NCBO BioPortal, <https://bioportal.bioontology.org/ontologies/SMASHSOCIAL> (2015, accessed 20 August 2026).
- [33] Herrmann SD, Willis EA, Ainsworth BE, et al. 2024 Adult Compendium of Physical Activities: a third update of the energy costs of human activities. *J Sport Health Sci* 2024; 13(1): 6–12. DOI: 10.1016/j.jshs.2023.10.010.
- [34] World Wide Web Consortium. Activity Streams 2.0, <https://www.w3.org/TR/activitystreams-core/> (2017, accessed 20 August 2026).
- [35] Open mHealth. Open mHealth schemas, <https://www.openmhealth.org/mhealth-schema/> (accessed 20 August 2026).
- [36] National Library of Medicine. NIH Common Data Elements Repository. National Institutes of Health, <https://cde.nlm.nih.gov/> (accessed 20 August 2026).
- [37] European Commission Joint Research Centre. EU RD Platform: set of common data elements for rare disease registration, https://eu-rd-platform.jrc.ec.europa.eu/system/files/public/CDS/EU_RD_Platform_CDS_Final.pdf (accessed 20 August 2026).
- [38] Observational Health Data Sciences and Informatics. Data standardization: the OMOP Common Data Model, <https://www.ohdsi.org/data-standardization/> (accessed 20 August 2026).
- [39] HL7 International. FHIR Release 5: Fast Healthcare Interoperability Resources, <https://www.hl7.org/fhir/> (2023, accessed 20 August 2026).
- [40] Uschold M and Gruninger M. Ontologies: principles, methods and applications. *Knowl Eng Rev* 1996; 11(2): 93–136. DOI: 10.1017/S0269888900007797.
- [41] Keet CM and Khan ZC. On the roles of competency questions in ontology engineering. In: Alam M, Rospocher M, Van Erp M, et al. (eds) *Knowledge engineering and knowledge management*. Springer, 2024, pp. 123–132. DOI: 10.1007/978-3-031-77792-9_8.
- [42] Institute of Electrical and Electronics Engineers. *IEEE recommended practice for software requirements specifications*. IEEE Std 830-1998. IEEE, 1998. DOI: 10.1109/IEEESTD.1998.88286.
- [43] Ramesh B, Powers T, Stubbs C, et al. Implementing requirements traceability: a case study. In: *Proceedings of the 1995 IEEE International Symposium on Requirements Engineering*. IEEE, 1995, pp. 89–95. DOI: 10.1109/ISRE.1995.512549.
- [44] Brickley D and Guha RV (eds). *RDF Schema 1.1*. World Wide Web Consortium, 2014. <https://www.w3.org/TR/rdf-schema/> (accessed 20 August 2026).
- [45] World Wide Web Consortium Permanent Identifier Community Group. w3id.org: permanent identifiers for the Web, <https://w3id.org/> (accessed 20 August 2026).
- [46] Musen MA. The Protégé project: a look back and a look forward. *AI Matters* 2015; 1(4): 4–12. DOI: 10.1145/2757001.2757003.
- [47] Kellogg G, Champin PA and Longley D (eds). *JSON-LD 1.1: a JSON-based serialization for Linked Data*. World Wide Web Consortium, 2020. <https://www.w3.org/TR/json-ld11/> (accessed 20 August 2026).
- [48] Peterson D, Gao S, Malhotra A, et al. (eds). *W3C XML Schema Definition Language (XSD) 1.1 Part 2: Datatypes*. World Wide Web Consortium, 2012. <https://www.w3.org/TR/xmlschema11-2/> (accessed 20 August 2026).

-
- [49] Harris S and Seaborne A (eds). *SPARQL 1.1 query language*. World Wide Web Consortium, 2013. <https://www.w3.org/TR/sparql11-query/> (accessed 20 August 2026).
- [50] R Core Team. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, 2024. <https://www.R-project.org/> (accessed 20 August 2026).
- [51] Posit Team. *RStudio: integrated development environment for R*. Posit Software, PBC, 2025. <https://posit.co/> (accessed 20 August 2026).
- [52] Langford J, Chua JY, Long I, et al. Validation and optimisation of wearable accelerometer data pre-processing for digital measure implementation and development. *bioRxiv*. Preprint, 2026. DOI: 10.64898/2026.03.21.713324.
- [53] International Organization for Standardization and International Electrotechnical Commission. *Privacy technologies—Consent record information structure*. ISO/IEC TS 27560:2023. 2023. <https://www.iso.org/standard/80392.html> (accessed 20 August 2026).
- [54] Bender D and Sartipi K. HL7 FHIR: an agile and RESTful approach to healthcare information exchange. In: *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems*. IEEE, 2013, pp. 326–331. DOI: 10.1109/CBMS.2013.6627810.
- [55] Falligant JM, Hagopian LP and Newland MC. Bouts, pauses, and units of operant performance: a primer. *Perspect Behav Sci* 2024; 47(3): 643–674. DOI: 10.1007/s40614-024-00419-z.

Appendix

Table 1. Description of errors seeded and validation tests performed for specific competency questions under assessment.

CQ focus	Error number and type (on subset of randomly selected Atoms)	Validation Test performed
CQ1: Basic interoperability, validation	Error 1: Field removed – Deleted required elements or sub-elements.	Test 1: Each Atom MUST include Header, When, What, and Who. Test 2: Each Atom MUST include Header.AtomVersion; When.TimeUTC, When.Duration, When.UTCOffset; What.Label; and Who.ParticipantID or Who.EnvironmentID.
CQ1: Basic interoperability, validation	Error 2: Field type changed – Flipped data types for required elements or sub-elements.	Test 3: Required sub-elements and elements MUST satisfy the schema constraints for type (e.g. AtomVersion string, Duration integer, Label array of strings, ParticipantID string)
CQ2: Basic interoperability, validation	Error 3: Label string corrupted – selected string fields corrupted.	Test 4: Atoms MUST use Label names that exist in the corresponding classification model registry.
CQ3: Governance, validation	Error 4: Direct PII injected – replaces a selected string field with a realistic randomly selected name token generated using “charlatan” package in R.	Test 5: An Atom MUST NOT contain personal names.
CQ4: Governance, validation	Error 5: Atoms duplicated – appends exact copies of selected Atoms	Test 6: Each Atom MUST be unique in both structure and content within an Atom payload.
CQ5: Temporal, validation	Error 6: Temporal gaps inserted – removes blocks of consecutive Atoms within a participant stream to create end-to-start gaps above a threshold (set at 30 min).	Test 7: An Atom payload MUST not have temporal gaps between subsequent Atoms that exceed a specified threshold (30 min).

CQ – competency question, PII – personal identifiable information

Table 2. RTM-to-CQ coverage by requirement theme.

Theme	CQ1	CQ2	CQ3	CQ4	CQ5	CQ6	CQ7	CQ8	Row total
Conceptual	8	4	0	0	1	0	2	1	16
Technical	5	5	0	1	2	1	1	5	20
Governance	1	0	2	1	0	0	0	0	4
Column total	14	9	2	2	3	1	3	6	40

Counts represent requirement-to-CQ links rather than unique requirements. A requirement linked to multiple CQs is counted once in each applicable CQ column. RTM – requirements traceability matrix; CQ – competency question.




Figure 1. Protégé asserted class hierarchy view showing the COEL Model top-level Clusters and an expanded path within the Accident Cluster.

Table 3. Key changes and their rationale between COEL Behavioural Atom v1.0 and v2.0.

Area	Atom v1.0	Atom v2.0	Rationale
Schema ID	Draft-04; no stable schema IRI	Draft-07 with \$id (w3id)	Enables resolvable, versioned schema identification and modern validation tooling.
Header versioning	Version = [spec_major,spec_minor,model_major,model_minor]	AtomVersion = "major.minor"; optional AtomIRI, AtomID; optional ExtensionRegistry[]	Separates Atom versioning from model versioning; supports tooling and extension governance.
Who	ConsumerID either/or DeviceID	ParticipantID and/or EnvironmentID; pseudonymous only	Generalises beyond consumer-device framing and makes the privacy constraint explicit and testable.
What	Cluster/Class/SubClass/Element integers	Label[] strings; optional LabelIRI[]	Supports different taxonomies and multi-label events and direct linking to published concept identifiers.
When	Time required; Duration optional	TimeISO or TimeUTC + UTCOffset required; Duration required	Reduces ambiguity, supports ISO 8601 workflows, and enforces time-bounded events.
How	How, Certainty, Reliability	EvidenceType, EvidenceSystem, ClassificationModel[] (+ optional ...IRI[]), EvidenceID, AssessmentTimeISO	Shifts from confidence scoring to explicit provenance and model descriptors needed for reproducibility and cross-model aggregation.
Consent	Date (Unix seconds), RetentionPeriod (integer), Purpose (enum); PolicyURL; RecordService	EventTime (ISO 8601), RetentionPeriod (string), Purpose (identifier recommended), PrivacyNotice; RecordID + PartyID	Aligns consent metadata with consent record structures and identifier-based governance (ISO27560).
Extension	Fixed tag/value pairs (Ext*Tag/Value); closed object	Open Extension object; registry-declared keys via ExtensionRegistry; custom keys prefixed with _	Reduces tag proliferation and supports typed, publishable extensions while retaining implementer flexibility.

Table 4. Evaluation metrics of the validation for CQs1-5, with baseline and error-seeded Atom payload results, reported for correctly identified Atom Errors vs the changelog.

Payload	CQ focus	True Errors	True Positives	False Positives	False Negatives	True Negatives	Accuracy (%)	F1-score (%)
Rest Activity Atom payload								
Error-seeded 1-2	CQ1	26	26	0	0	204	100	100
Error-seeded 3	CQ2	26	26	0	0	204	100	100
Error-seeded 4	CQ3	26	7	0	19	204	91.7	42.4
Error-seeded 5	CQ4	26	26	0	0	230	100	100
Error-seeded 6	CQ5	48	48	0	0	156	100	100
Baseline	CQ1-5	0	0	0	0	230	100	—
Behavioural Bout Atom payload								
Error-seeded 1-2	CQ1	29,030	29,030	0	0	87,275	100	100
Error-seeded 3	CQ2	29,030	29,030	0	0	87,275	100	100
Error-seeded 4	CQ3	26,866	4,900	0	21,966	89,439	81.1	30.9
Error-seeded 5	CQ4	29,030	29,030	0	0	116,305	100	100
Error-seeded 6	CQ5	48	48	0	0	92,995	100	100
Baseline	CQ1-5	0	0	0	0	116,305	100	—

CQ – competency question; Error 1 – field removed; Error 2 – field type changed; Error 3 – Label string corrupted; Error 4 – personal names injected; Error 5 – Atoms duplicated; Error 6 – temporal gaps inserted

```

{
  "Header": {
    "AtomVersion": "2.0",
    "AtomIRI": "https://w3id.org/coel/atom/2.0",
    "AtomID": "111710_1749218067_1_4032059#1749214484#27#slow_walk",
    "ExtensionRegistry": [
      "https://w3id.org/coel/atom/2.0/extension-registry.csv"
    ]
  },
  "When": {
    "TimeUTC": 1749214484,
    "UTCOffset": 3600,
    "TimeISO": "2025-06-06T13:54:44+01:00",
    "Duration": 27
  },
  "What": {
    "Label": [
      "slow_walk",
      "14.5"
    ],
    "LabelIRI": [
      "https://w3id.org/coel/models/activinsights/behavioural_bout/1.0#slow_walk",
      "https://w3id.org/coel/models/coel/2.0#14.5"
    ]
  },
  "Who": {
    "ParticipantID": "{-t$EFb&)nvn-VQ#)P6"
  },
  "How": {
    "EvidenceType": 7,
    "EvidenceSystem": "GENEAcore Version: 1.1.3; Device: GENEActiv 1.2; DeviceID: 111710; FirmwareVersion: Ver06.17 15June23; Protocol: right wrist, 100 Hz",
    "ClassificationModel": [
      "Behavioural Bout Model v1.0",
      "COEL Model v2.0"
    ],
    "ClassificationModelIRI": [
      "https://w3id.org/coel/models/activinsights/behavioural_bout/1.0/",
      "https://w3id.org/coel/models/coel/2.0/"
    ],
    "AssessmentTimeISO": "2025-12-19T01:00:00+01:00",
    "EvidenceID": "111710_1749218067_1_4032059"
  },
  "Extension": {
    "xMean": -0.2053,
    "xSD": 0.8224,
    "yMean": -0.1288,
    "ySD": 0.4844,
    "zMean": -0.7133,
    "zSD": 0.4445,
    "LightMean": 122.2533,
    "LightMax": 700,
    "TempMean": 36.2,
    "TempSD": 0.4137,
    "AGSAMEan": 0.2681,
    "ENMOMean": 0.2017,
    "UpDownMean": -6.9948,
    "UpDownSD": 26.9858,
    "DegreesMean": 148.39,
    "DegreesSD": 42.5655,
    "StepCount": 27,
    "StepMean": 61.5502,
    "StepSD": 0.7841
  }
}

```

Figure 2. Example data record of a behavioural event generated by the Behavioural Bout classification algorithm in the COEL Behavioural Atom v2.0 normative format.

```

{
  "@context": {
    "prov": "http://www.w3.org/ns/prov#",
    "dcterms": "http://purl.org/dc/terms/",
    "skos": "http://www.w3.org/2004/02/skos/core#",
    "xsd": "http://www.w3.org/2001/XMLSchema#",

    "Header": {
      "@context": {
        "AtomID": "dcterms:identifier",
        "AtomIRI": { "@id": "dcterms:conformsTo", "@type": "@id" }
      }
    },
    "When": {
      "@context": {
        "TimeUTC": { "@id": "dcterms:temporal", "@type": "xsd:integer" },
        "Duration": { "@id": "dcterms:extent", "@type": "xsd:integer" },
        "UTCOffset": { "@id": "dcterms:spatial", "@type": "xsd:integer" }
      }
    },
    "What": {
      "@context": {
        "LabelIRI": { "@id": "dcterms:subject", "@type": "@id", "@container": "@set" }
      }
    },
    "Who": {
      "@context": {
        "ParticipantID": "dcterms:references",
        "EnvironmentID": "dcterms:references"
      }
    },
    "How": {
      "@id": "prov:wasGeneratedBy",
      "@type": "prov:Activity",
      "@context": {
        "EvidenceType": { "@id": "dcterms:type", "@type": "xsd:integer" },
        "EvidenceSystem": "dcterms:description",
        "ClassificationModelIRI": { "@id": "prov:used", "@type": "@id", "@container": "@set" },
        "AssessmentTimeISO": { "@id": "prov:endedAtTime", "@type": "xsd:dateTime" },
        "EvidenceID": "dcterms:identifier"
      }
    }
  }
}

```

Figure 3. JSON-LD context of COEL Behavioural Atom v2.0 for projection of JSON to JSON-LD format.

```
PREFIX dterms: <http://purl.org/dc/terms/>
PREFIX xsd:    <http://www.w3.org/2001/XMLSchema#>

SELECT ?src (SUM(xsd:integer(?dur)) AS ?duration_s)
WHERE {
  ?atom dterms:extent ?dur .
  ?atom dterms:subject ?src .
}
GROUP BY ?src
```

Figure 4. SPARQL query used to summarise time by native model labels from the RDF projection of an Atom payload. The query aggregates Atom duration values (seconds) across label concepts referenced via dterms:subject, returning a total duration per native label IRI for the selected participant window.

```
PREFIX dcterms: <http://purl.org/dc/terms/>
PREFIX skos:    <http://www.w3.org/2004/02/skos/core#>
PREFIX xsd:     <http://www.w3.org/2001/XMLSchema#>

SELECT ?coel (SUM(?d) AS ?duration_s)
WHERE {
  {
    SELECT ?coel (xsd:integer(?dur) AS ?d)
    WHERE {
      ?atom dcterms:extent ?dur .
      ?atom dcterms:subject ?src .

      FILTER(!STRSTARTS(STR(?src), "https://w3id.org/coel/models/coel/2.0")) .

      { ?src skos:broadMatch ?coel . }
      UNION
      { ?src skos:relatedMatch ?coel . }
    }
  }
}
GROUP BY ?coel
```

Figure 5. SPARQL query used to roll up native model labels to COEL v2.0 concepts via SKOS mapping traversal. The query aggregates Atom duration values (seconds) after traversing curated SKOS mapping predicates (for example, `skos:broadMatch` and `skos:relatedMatch`) from native label IRIs to COEL concept IRIs, returning a total duration per COEL IRI for the selected participant window.