CoTOn: A Mid-Level Ontology for Theory-indexed Dataset Annotation and Retrieval in Cognitive Neuroscience

Anna Ravenschlag[1,*], Florian Hutzler[1] and Giancarlo Guizzardi[2]

[1]Salzburg University, Kapitelgasse 4/6, 5020 Salzburg, Austria
[2]University of Twente, Drienerlolaan 5, 7522 NB, Enschede, The Netherlands

Abstract

Cognitive neuroscience increasingly depends on rich, machine-actionable metadata to make datasets findable, interoperable, and reusable. However, metadata must describe not only files, formats, and experimental procedures, but also what a dataset is scientifically interpreted as being about. Neurocognitive datasets may contain neural measurements, behavioral responses, and task descriptions, but their scientific meaning depends on theoretical concepts such as memory, attention, or cognitive control. These concepts are latent constructs that are not directly observable in the data and may be conceptualized differently across theories. To address this problem, we showcase the Cognitive Theory Ontology (CoTOn) as a theory-aware mid-level ontology for dataset annotation and retrieval in cognitive neuroscience. Compared with an earlier alpha version, the current implementation extends CoTOn from a conceptual proposal into a modular, queryable ontology artifact grounded in Unified Foundational Ontology. The new version represents theory-indexed relations between cognitive terms, theory-specific concepts, cognitive capacity types, experimental settings, and dataset annotations across interconnected schema, theory, and annotation layers. We demonstrate its use with CoTOn-based annotations of two OpenNeuro datasets that share the term "cognitive control" but differ in the explicitness of their theory-specific interpretation. Five competency questions evaluate term-level retrieval, theory-indexed disambiguation, curated concept retrieval, and task-based candidate matching. SHACL validation identifies underspecified annotations, while SPARQL queries use theory-level operationalization knowledge to propose inspectable candidate annotations for curator review while keeping them distinct from curated assertions. The showcase illustrates how CoTOn supports semantic dataset discovery, validation of sparse metadata, and annotation enrichment beyond keyword-based retrieval while remaining compatible with heterogeneous and incomplete neurocognitive metadata.

Keywords

Cognitive Theory Ontology, research data management, metadata, dataset annotation, knowledge graph, neuroinformatics, dataset discovery

CEUR-WS.org/Vol-4342/paper3.pdf

1. Introduction

Cognitive neuroscience is a data-intensive and theory-driven discipline that seeks to explain human experience and behavior by relating cognitive capacities to physiological, behavioral, and neural measurements obtained in designed experimental settings. Over recent decades, the field has produced growing numbers of publicly available datasets, metadata schemas, vocabularies, and knowledge representation resources. These developments help make neurocognitive datasets findable, comparable, and reusable across studies. However, they so far do not solve a central semantic problem: neurocognitive data is not semantically self-contained, because what a dataset *is about* does not follow from measured values alone. It depends on the cognitive theory under which a target capacity is conceptualized, the experimental design through which this concept is assessed, and the indicators through which it is measured. Consequently, a dataset annotated with a common task name or cognitive term is often only partially interpretable. For example, metadata may describe a dataset as involving "cognitive control" or "working memory" without specifying the intended theory-dependent concept, its relations to other concepts, or the targeted capacity type.

Proceedings of FOIS 2026 Satellite events co-located with the 16th International Conference on Formal Ontology in Information Systems (FOIS 2026), September 21-25, 2026, Vitória, Brazil.
*Corresponding author.
annanatali.ravenschlag@plus.ac.at (A. Ravenschlag); florian.hutzler@plus.ac.at (F. Hutzler); g.guizzardi@utwente.nl (G. Guizzardi)
0000-0003-0272-1090 (A. Ravenschlag); 0000-0001-8195-4911 (F. Hutzler); 0000-0002-3452-553X (G. Guizzardi)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-27

This problem is not merely terminological. Cognitive concepts are latent constructs that are not directly observable and may be defined differently across theories [1]. Yet, neurocognitive communities routinely use the same linguistic terms to communicate about such (potentially diverging) concepts. This creates tension between the need for (1) shared terminology in the community, and (2) preserving theoretical disagreement as a driving force for scientific progress. If ignored, dataset annotations risk collapsing distinct conceptualizations into a single label. However, if represented explicitly, the same label can become an entry point into a structured semantic space. To address this problem, we present the *Cognitive Theory Ontology*[1] (CoTOn) as a mid-level ontology for theory-indexed dataset annotation retrieval. The implementation builds on an ontologically unpacked, multi-level reconstruction of the first CoTOn alpha version [2] – elevating it from a conceptual proposal into a modular, queryable ontology artifact. As such, the current version is ontologically grounded in a multi-level conceptual model and provides an RDF serialization, reusable theory modules, vocabulary integrations, dataset annotation files, SHACL validation shapes, and executable SPARQL evaluation queries.

The modular implementation distinguishes three interconnected layers (see also Figure 1). At the *schema layer*, CoTOn defines general classes and relations for representing relevant domain entities such as cognitive capacities, theories, concepts, experimental settings, measurements, and data artifacts. At the *theory layer*, separate modules represent theory-specific concepts and their relations to terms, capacity types, and experimental operationalizations. At the *annotation layer*, dataset-specific files connect repository datasets and local annotation entities to reusable schema and theory entities. These layers support—but should not be confused with—CoTOn's multilevel modeling distinctions, through which reusable types, such as task and capacity types, remain distinct from their instances. Metadata is thus integrated into an explicit ontological structure rather than represented as term-based labels.

CoTOn is positioned as a mid-level ontology between the *Unified Foundational Ontology* (UFO [3]) and domain-level resources such as the *Cognitive Atlas* [4], the *Cognitive Task Ontology* (CogiTO [5]), and *Hierarchical Event Descriptors* (HED [6]). Importantly, CoTOn does not aim to replace these resources. Rather, it provides a semantic integration layer that makes explicit how terms, theory-specific concepts, capacity types, experimental settings, measurements, and datasets are related. In particular, CoTOn distinguishes a theory-specific cognitive concept from (1) the capacity it conceptualizes and (2) the tasks through which it is operationalized. This ontological distinction is central to CoTOn: a cognitive *concept* is a theory-dependent construct, whereas a cognitive *capacity* is an individual mode inhering in an agent and classified by a cognitive *capacity type*. An *operationalization* connects a cognitive theory and theory-specific concept to an experimental setting type. Cognitive tasks may form parts of such a setting, while their executions, measurements, and resulting data remain ontologically distinct.

As a prototypical use case, we introduce the running example of *cognitive control*. Cognitive control is a widely used term in cognitive neuroscience, referring broadly to the capacity to guide thought and action in a goal-directed manner. However, it has no single uncontested conceptualization [7]. In the *Dual Mechanisms of Cognitive Control* (DMCC) theory, cognitive control is conceptualized as having component concepts of *proactive control* and *reactive control* [8]. In another framing, called here the *Cognitive Control Component Framing* (CCCF), it is decomposed into the concepts of *inhibition*, *shifting*, and *working memory* [9]. A flat metadata model may annotate datasets associated with these diverging conceptualizations using the same term "cognitive control", thereby collapsing distinct conceptualizations into a single label. CoTOn goes beyond common keyword indexing by representing that such annotations may refer to different theory-specific concepts while retaining term-level retrieval.

Critically, CoTOn does not favor one conceptualization over another. Instead, it preserves their different theoretical commitments and makes them available for comparison and retrieval. This helps prevent datasets with different theoretical interpretations from being treated as semantically equivalent merely because they share the same term. Conversely, CoTOn can identify datasets that use different terminology but connect to related cognitive concepts or capacity types. Represented theory and operationalization relations may also support dataset comparison, evidence synthesis, or possible reanalysis under an alternative theoretical framing.

[1] https://gitlab.com/anrave/coton

The remainder of the paper is structured as follows. Section 2 discusses related resources and the remaining gap. Section 3 presents CoTOn's ontological foundation, core modeling pattern, and implementation. Section 4 introduces the running example of cognitive control example its dataset annotation files. Section 5 presents the competency questions, SPARQL evaluation, results, and validation mechanisms. Finally, Section 6 discusses scientific relevance, limitations, and future work.

2. Related Resources and Remaining Gap

Several existing resources recognize the need for structured knowledge representation in cognitive neuroscience and provide building blocks for dataset annotation and semantic search. However, each resource focuses on a specific part of the problem. CoTOn reuses and integrates selected resources while adding an explicit theory-aware interpretation layer.

The Cognitive Atlas [4] systematically represents cognitive concepts and experimental tasks. A central motivation is that tasks and cognitive concepts are often conflated: tasks are treated as if they directly identified or measured concepts, although task-concept relations are not self-evident. Rather, a task is connected to a theory-specific cognitive concept through an operationalization involving an experimental setting type. Beyond distinguishing between tasks and concepts, CoTOn also enables terminological disambiguation between concepts. Communities may use the same term, such as "cognitive control" for different concepts, while the same or related concepts may be referred to by different terms. CoTOn therefore distinguishes cognitive terms, theory-specific cognitive concepts, and cognitive capacity types and relates them through theory-indexed semantic relations. While the Cognitive Atlas represents a cognitive concept through an entry with a single definition, CoTOn supports multiple theory-specific conceptualizations associated with the same term.

Hierarchical Event Descriptors (HED) [6] address a complementary problem. HED provides a taxonomic structure for annotating laboratory and real-world events. It supports fine-grained descriptions of stimuli, participant actions, presentation modalities, and environmental contexts, thus enabling event-level search, extraction, and comparison across studies. However, an event description does not itself specify which cognitive concept a task operationalizes or under which theory. HED can describe an event and its observable features, while CoTOn connects such experimental structures to theory-indexed cognitive concepts.

The Cognitive Task Ontology (CogiTO) [5] bridges task- and event-level descriptions by applying HED tags to Cognitive Atlas tasks. This operational connection between reusable task descriptions and event annotations is reused in CoTOn. The implementation therefore does not reproduce the task and event taxonomies supplied by these resources, but connects their entities to CoTOn's theory, concept, capacity, and operationalization structure.

Relevant resources also exist within the *Open Biological and Biomedical Ontologies Foundry*. The Ontology for Biomedical Investigations (OBI) [10] extensively represents biological and clinical investigations, including protocols, devices, specimens, generated data, and analyses. OBI thus overlaps partly with CoTOn's scope and provides entities that may be reused or aligned when more detailed representations of investigations and measurements are needed. However, CoTOn as a mid-level ontology has a different focus: rather than providing a general account of biomedical investigations, it represents how experimental structures and data annotations connect to theory-specific cognitive concepts and capacity types.

The Mental Functioning Ontology (MFO) [11] is an overarching ontology for mental functioning grounded in the Basic Formal Ontology [12]. It includes cognition and other aspects of mental functioning and is directly relevant to CoTOn's domain. Its realist orientation provides an account of mental processes and related entities. CoTOn addresses a complementary problem by distinguishing such domain entities from the theory-dependent concepts through which researchers conceptualize them. In CoTOn, a cognitive concept is not itself asserted to be a mental process or capacity, but is a theory-dependent conceptual entity related to a capacity type and the experimental structures through which that capacity is investigated in a particular theoretical context.

CoTon does not yet import or formally align OBI and MFO. This is not because they are considered unsuitable, but because the current implementation focuses on establishing and evaluating CoTOn's core theory-indexing and multi-level annotation patterns. Future alignment work will investigate which entities can be reused directly and where mappings are needed between CoTOn's UFO-grounded distinctions and the BFO-grounded distinctions used in OBI and MFO.

Taken together, these resources support the representation of cognitive constructs, experimental tasks, event structures, mental functioning, and biomedical investigations. What remains insufficiently represented for this use case is an explicit theory-aware layer relating linguistic terms, potentially diverging theory-specific concepts, cognitive capacity types, experimental operationalizations, and dataset annotations. CoTOn fills this gap by incrementally integrating existing neurocognitive resources and supporting future alignment with biomedical ontologies. This enables retrieval through shared terminology without treating that terminology as conceptually unambiguous and allows task- and event-level metadata to be interpreted within explicit theoretical contexts.

3. The Cognitive Theory Ontology

CoTOn is designed as a mid-level ontology for neurocognitive metadata semantics. Its aim is not to provide an exhaustive domain vocabulary, but a semantic structure connecting existing vocabularies, cognitive theories, experimental tasks, measurements, and dataset annotations. The implementation is based on the ontological unpacking [13] of the alpha version of CoTOn [2]. It is grounded in UFO [3], represented in OntoUML [14], and serialized in RDF/Turtle using gUFO [15] (*gentle UFO*, a lightweight UFO implementation for knowledge graphs semantic web applications).

3.1. Ontological Foundation and Core Modeling Pattern

Grounding CoTOn in UFO [3] distinguishes relevant domain entities of different ontological kinds. Without this grounding, flat metadata may compress these distinctions: for example, a cognitive concept may be treated as equivalent to a linguistic label, cognitive capacity, or experimental task. CoTOn instead represents them as distinct but related entities. A `Cognitive Agent` is a substantial individual capable of bearing cognitive capacities. A `Cognitive Capacity` is an intrinsic mode inhering in a cognitive agent, whereas a `Cognitive Capacity Type` classifies capacities of the same kind. Cognitive theories, concepts, and terms belong to a social-semiotic part of the model. A `Cognitive Theory` provides the context in which capacities are conceptualized. A `Cognitive Concept` is a theory-dependent hypothetical entity that abstracts from a capacity, and a `Cognitive Term` is a linguistic entity through which such a concept may be expressed. A cognitive concept is therefore not a capacity inhering in an agent. Rather, it is entertains relations to the capacity type it conceptualizes and the term representing it within a particular theoretical context.

Most importantly, relations between terms, concepts, and capacity types are not simple label relations but are indexed by a cognitive theory. This structure can be understood as a theory-indexed semiotic triangle [16], where (1) the term provides a linguistic access point, (2) the concept provides a theory-specific construct, and (3) the capacity type represents the kind of capacity addressed by the concept. CoTOn represents these connections through relator patterns involving theory, term, concept, and capacity type. Derived relations such as `represents`, `conceptualizes`, and `abstracts from` provide direct access to the resulting RDF connections while explicitly representing their underlying truthmaking structures. This design makes it possible to represent the same linguistic term as referring to different concepts in different theories. Simultaneously, it allows different concepts to be associated with the same capacity type without treating the concepts as identical. In the running example, *proactive control* and *reactive control* are parts of the DMCC cognitive control conceptualization, whereas *inhibition*, *shifting*, and *working memory* are parts of the CCCF conceptualization.

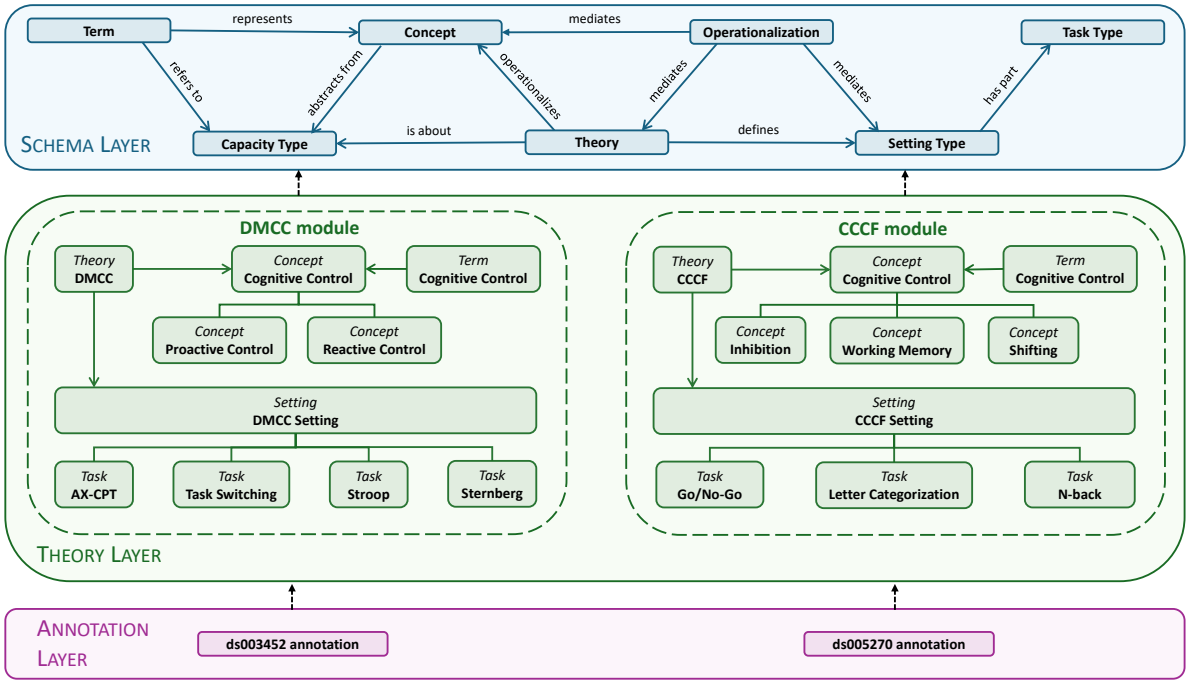

Figure 1: Core theory-indexed conceptualization and operationalization pattern in CoTOn.

An `Operationalization` represents how a cognitive concept is connected to an `Experimental Setting Type` under a given cognitive theory. An experimental setting type describes a reusable configuration and may have task, event, condition, contrast, and indicator types as parts. CoTOn further distinguishes experimental design types from concrete executions. An `Experimental Setting Execution` is an event in which an `Experimental Setting` participates, which in turn instantiates an `Experimental Setting Type`. Likewise, reusable task, event, condition, contrast, indicator, and measurement classes remain distinct from the concrete executions in an investigation. These distinctions reflect CoTOn's multilevel character: reusable tasks and capacites occur at the type level, while dataset-local tasks and individual capacities occur at the instance level. In RDF, external task classes from the Cognitive Atlas and CogiTO remain reusable types, while annotation files introduce dataset-local individuals connected to them. Queries can thus traverse from a dataset-local annotation to a reusable task type and then to theory-level operationalization knowledge without collapsing types and instances.

Figure 1 presents the the three interconnected layers of CoTOn. At the schema layer, CoTOn defines general classes and relations for representing relevant domain entities. At the theory layer, these types are instantiated by theory-specific classes. At the annotation layer, dataset-specific annotation files connect repository datasets to reusable entities from the theory layer. Note that the figure shows the main distinctions used in the running example without visualizing every class in the CoTOn schema.

3.2. Modular Implementation

The implementation follows a modular structure (Table 1). The file `coton-schema.ttl` contains CoTOn's general ontological commitments and provides ontology classes and relationd based on CoTOn's conceptual model. The model was serialized in RDF through gUFO [15]. The main import module, `coton-core.ttl`, connects the schema to internal and external vocabularies. These include Cognitive Atlas [4], HED [6], and CogiTO [5] entities, mappings, and multilevel links connecting reusable task types to dataset-local individuals. Separate theory modules import the core ontology. For the running example, they represent the DMCC [8] and CCCF [9] cognitive control conceptualizations, including their conceptt components, terminological relations, and operationalization patterns.

Table 1
Selected modules and implementation artifacts of CoTOn.

Module	File or folder	Purpose
Schema	`ontology/ coton-schema.ttl`	Classes and properties derived from the conceptual model
Core	`ontology/ coton-core.ttl`	Main import module for the schema and vocabulary resources
Vocabularies	`vocabularies/`	External vocabularies, CoTOn vocabularies, mappings, and integrations
Theories	`theories/`	Theory-specific concepts, conceptualizations, and operationalization knowledge
Dataset annotations	`annotations/ openneuro/`	Annotation files of selected datasets from OpenNeuro
Knowledge graph profile	`kg-profile/ coton-kg-profile.ttl`	Query-oriented relations for sparse metadata and candidate annotations
Validation	`shapes/`	SHACL shapes for identifying invalid or underspecified annotations
Integrated build	`build/ coton-integrated.ttl`	Loadable graph for SPARQL querying and evaluation

Dataset-specific annotation files are stored separately from schema and theory modules. They identify repository datasets and connect them to cognitive terms, concepts, theories, tasks, and other annotation entities. This separation shows which assertions originate from CoTOn's schema, belong to a cognitive theory, or annotate a dataset. It also allows an annotation to remain incomplete without altering theory or schema modules. Because public neurocognitive metadata is often heterogeneous and incomplete, the implementation includes a knowledge graph (KG) profile in `coton-kg-profile.ttl`. The profile provides query-oriented relations for dataset-level retrieval over sparse metadata. It also distinguishes curated annotations from candidates generated by matching dataset-local task profiles against theory-level operationalization knowledge. Candidate relations are curation suggestions, not equivalent to assertions accepted in an annotation file.

The integrated build is stored in `coton-integrated.ttl`. It imports the schema, vocabularies, theory modules, dataset annotations, and KG profile into one graph loadable in GraphDB. Separate SHACL shapes identify structurally invalid or underspecified annotations.

Figure 2 illustrates the general workflow: external cognitive resources are connected through theoretical contexts, vocabulary mappings, and the KG profile. The CoTOn core, theory modules, and dataset annotations are loaded into a KG. The resulting graph is evaluated through SHACL validation and SPARQL querying to support theory-aware dataset discovery, terminology disambiguation, and traceable candidate annotation suggestions. Further, the modular structure supports a range of CoTOn's intended applications where users can inspect the conceptual schema, load selected theory modules, add and validate dataset annotations, or query the integrated build.

4. Cognitive Control as a Running Example

Cognitive control is a useful running example because it is frequently investigated in cognitive neuroscience but has no single uncontested conceptualization [7]. Depending on the theoretical context, it may be conceptualized through proactive and reactive control [8], or decomposed into inhibition, shifting, and working memory [9]. A dataset annotation containing only the term "cognitive control" therefore does not determine which theory-specific concept is intended.

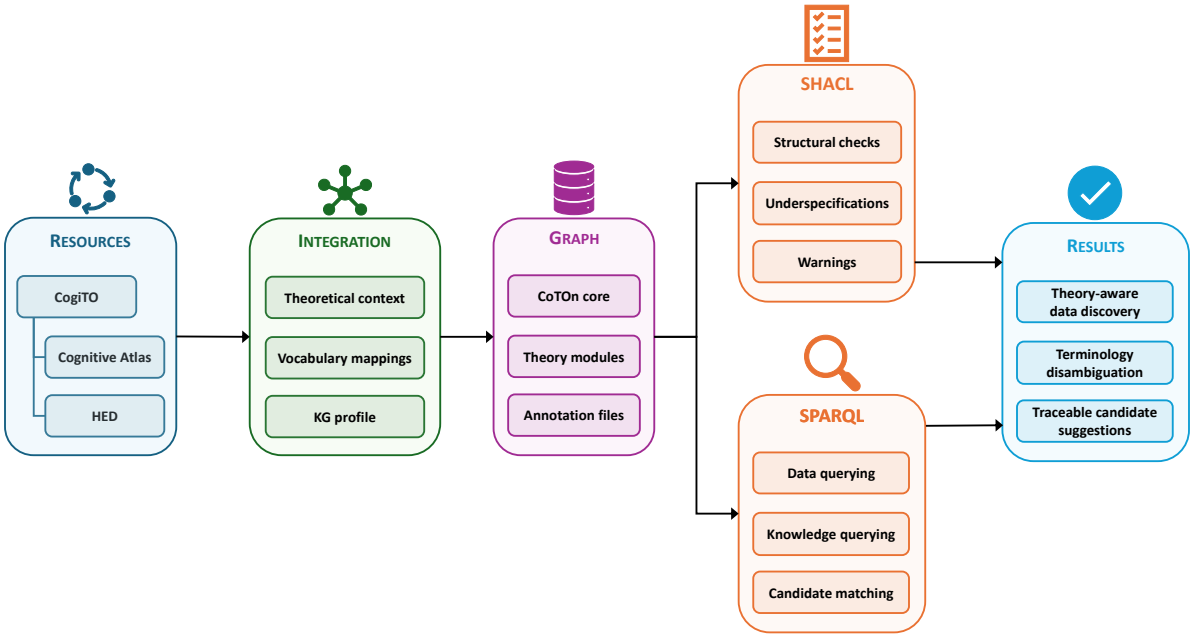

Figure 2:

Figure 3: CoTOn implementation and evaluation workflow.

For this showcase, we created CoTOn-based annotation files for two exisitng OpenNeuro[2] [17] datasets, `ds003452` [18] and `ds005270` [19]. OpenNeuro is an established repository for sharing human and animal brain-imaging data. Note that the original datasets were not modified for this work. The created annotations files represent selected conceptual and experimental information about these datasets using CoTOn. This distinction between a *source dataset* and its CoTOn-based *annotation file* is important: the source dataset comprises the original repository files and metadata, whereas the annotation file is a separate RDF artifact relating the dataset to terms, concepts, theories, capacity types, and local task annotations.

The annotation of `ds003452` represents a theory-explicit case. The dataset originates from the DMCC project [8] and employs a theoretically guided task battery. Its CoTOn-based annotation explicitly associates the dataset with the *term* "cognitive control", the DMCC *theory*, and the DMCC-specific *concepts* cognitive control, proactive control, and reactive control. It also contains local annotations for the *AX Continuous Performance Task*, *color-word Stroop task*, *cued task switching*, and *Sternberg item recognition*. These local task entities are individuals linked to reusable task types in the integrated ontology. The annotation of `ds005270` represents a deliberately sparse case. For the purpose of this showcase, its CoTOn annotation was limited to the *term* "cognitive control" and local annotations for the *Go/No-Go*, *letter-categorization*, and *N-back* tasks. No theory-specific concept, cognitive theory, or capacity type is asserted at dataset level. Table 2 summarizes the assertions in both annotation files.

The DMCC module represents the concept of cognitive control as having the concept parts proactive and reactive control, represented as theory-specific concepts. Their relations to empirical procedures are modeled separately through operationalizations and experimental settings. The CCCF module represents a different structure: cognitive control has the component concepts inhibition, shifting, and working memory. An operationalization connects this CCCF concept to an experimental setting type whose task type parts include Go/No-Go, letter categorization, and N-back tasks. Theory-specific component concepts are likewise connected to corresponding task types: *inhibition* to a *Go/No-Go task type*, *shifting* to a *letter-categorization task type*, and *working memory* to an *N-back task type*. Critically,

[2] https://openneuro.org/

Table 2
Selected assertions in the CoTOn annotation files of the two OpenNeuro datasets `ds003452` and `ds005270`.

Annotation element	`ds003452`	`ds005270`
Term	"cognitive control"	"cognitive control"
Theory	DMCC	Not asserted at dataset level
Concepts	Cognitive Control; Proactive Control; Reactive Control	Not asserted at dataset level
Tasks	AX-CPT; color-word Stroop; cued task switching; Sternberg item recognition	Go/No-Go; letter categorization; N-back

these theory-module assertions do not state that performing one of these tasks proves that a dataset instantiates a theory. A task may serve different purposes and be interpreted under different frameworks. Instead, the assertions encode the CCCF operationalization profile, which can be compared with local task annotations to identify a possible theory-aware interpretation.

The two annotation files demonstrate complementary uses of CoTOn. First, retrieval only through the term "cognitive control" would hide their theoretical differences. CoTOn instead represents a cognitive term as a linguistic entity associated with distinct theory-specific concepts, allowing `ds003452` to be retrieved with its DMCC interpretation. Second, when an annotation is sparse – which is a recurring issue with neurocognitive metadata – and lacks theory or concept assertions, CoTOn can compare its task profile with theory-level operationalization knowledge. For `ds005270`, the local Go/No-Go, letter-categorization, and N-back annotations correspond to the CCCF task-type requirements. CoTOn can therefore return CCCF and its cognitive-control concept as candidate annotations for curator review. Note that such a result is not treated as a deductive conclusion that the original researchers indeed intended a conceptualization of cognitive control according to CCCF. Theory and concept annotations remain absent from the curated annotation until reviewed and accepted. Candidate annotations use separate knowledge-graph-profile properties and remain distinguishable from assertions in the dataset annotation file. The example therefore evaluates whether CoTOn can make a possible interpretation transparent and inspectable, not whether it can infer a uniquely correct theory from a task list.

5. Dataset Annotation, Querying, and Validation

CoTon as a computational artifact is evaluated through competency questions (Table 3) over the integrated CoTOn knowledge graph. The graph is created by loading `build/coton-integrated.ttl` into GraphDB. This file imports the schema, vocabularies, theory modules, dataset annotations, and query-oriented KG profile. The authors formulated the competency questions based on expertise in cognitive neuroscience and ontology engineering. They derive from the use case in Section 4: retrieving datasets that share a term, distinguishing their theory-specific interpretations, and identifying possible interpretations of sparse annotations from task-level evidence. The questions therefore serve as implementation-oriented capability tests.

The first three competency questions address theory-indexed dataset discovery: CQ1 establishes baseline term-level retrieval. CQ2 retrieves the theory-specific concepts and concept parts represented by the term "cognitive control". CQ3 retrieves datasets with curated theory-specific concept annotations for that term. CQ4 and CQ5 address sparse annotations, which are common in neurocognitive datastes. CQ4 identifies datasets whose local task annotations cover the task-type requirements of a specific theory-level operationalization profile. CQ5 explains the candidate match by returning the local task annotations that provide the evidence.

Figure 4 shows the principal SPARQL patterns used in these queries. Panel (a) retrieves the theory-specific concept and experimental setting associated with the term "cognitive control." Panel (b) tests whether the local task annotations of the dataset cover all task types required by that setting. Panel (c) retrieves the task-level evidence connecting the dataset-local tasks to the operationalizations of

Table 3
Competency questions used to evaluate the current CoTOn implementation.

CQ	Natural-language query	Capability evaluated
CQ1	Which datasets are annotated with the term "cognitive control"?	Baseline term-level retrieval
CQ2	Which theory-specific concepts and concept parts are represented by the term "cognitive control"?	Theory-indexed terminological disambiguation
CQ3	Which datasets already have curated theory-specific concept annotations for the term "cognitive control"?	Curated theory-aware dataset interpretation
CQ4	Which sparse datasets with the term "cognitive control" match a theory-level operationalization profile through their task annotations?	Candidate theory-aware interpretation from sparse metadata
CQ5	Which local task annotations provide the evidence for the candidate match?	Interpretability of candidate annotations and multilevel task matching

the component concepts. Note that prefix declarations, label retrieval for output variables, and other presentationrelated parts are omitted for readability. The full example queries can be found in the CoTOn repository[3].

The nested `FILTER NOT EXISTS` pattern in Figure 4(b) expresses a coverage condition: every task type required by the theory-level setting must be matched by at least one local task annotation. The property path `cotonkg:datasetHasAnnotationTask/rdf:type/rdfs:subClassOf*` first follows the dataset annotation to a local task individual and uses `rdf:type` to retrieve its reusable task type. The final element, `rdfs:subClassOf*`, matches either that type or any superclass (the asterisk denotes zero or more applications of `rdfs:subClassOf`). Thus, a local task annotated with a more specific type can satisfy a requirement expressed through a more general type.

For each competency question, the expected result was determined from the assertions represented in the integrated graph and compared with the GraphDB result. Table 4 reports the comparison. Observed and expected results matched for all five questions. CQ1 most closely resembles keyword-style discovery. It retrieves both `ds003452` and `ds005270` because both annotations contain "cognitive control". Unlike in common keyword search, the term is an entity embedded in a semantic structure and can connect to different theory-specific concepts. CQ2 queries knowledge stored in the theory modules independently of a dataset. It retrieves both the DMCC conceptualization on cognitive control (with its two component concepts proactive and reactive control) and the CCCF conceptualization (with the three component concepts inhibition, shifting, and working memory). This shows that CoTOn treats the shared term as a linguistic access point rather than one unambiguous concept. CQ3 retrieves only `ds003452`, whose annotation explicitly asserts the DMCC-specific cognitive control concept; `ds005270` is excluded because no theory-specific concept is asserted. CQ4 identifies `ds005270` as matching the CCCF operationalization profile because the task annotations cover the task-type requirements in the context of this theory. Accordingly, CCCF is returned as a candidate theory together with cognitive control and its three components as candidate concepts. CQ5 makes this result inspectable by returning that in the context of CCCF, the Go/No-Go task is associated with the concept of inhibition, the letter categorization task is associated with the concept of shifitng, and the N-back task is associated with the concept of working memory.

A candidate suggestion is thus grounded in a match between the dataset-local annotation and a theory-specific operationalization profile. However, candidate annotation properties in the KG profile are intentionally separated from curated annotation properties. A candidate result is a curation suggestion, not a deductively entailed assertion. This distinction is necessary because a dataset's task profile may be compatible with an operationalization without uniquely determining the theory intended by the original researchers.

[3] https://gitlab.com/anrave/coton

(a) Theory traversal

```
?d cotonkg:
  datasetHasAnnotationTerm
  ?t .

?c coton:isRepresentedBy
  ?t .

?th coton:conceptualizes
  ?c .

?op a
  coton:Operationalization ;
  gufo:mediates
    ?th,
    ?c,
    ?s .
```

(b) Profile coverage

```
FILTER NOT EXISTS {
  ?s coton:
    hasTaskTypePart ?rt .

  FILTER NOT EXISTS {
    ?d cotonkg:
      datasetHasAnnotationTask/
      rdf:type/
      rdfs:subClassOf*
      ?rt .
  }
}

FILTER NOT EXISTS {
  ?d cotonkg:
    datasetHasAnnotationConcept
    [] .
}
```

(c) Evidence extraction

```
?c coton:hasConceptPart
  ?part .

?op2 a
  coton:Operationalization ;
  gufo:mediates
    ?th, ?part, ?ps .

?ps coton:hasTaskTypePart
  ?rt .

?d cotonkg:
  datasetHasAnnotationTask
  ?lt .

?lt rdf:type/
  rdfs:subClassOf*
  ?rt .
```

Figure 4: Principal SPARQL patterns for CQ4 and CQ5. Panel (a) retrieves the theory-specific concept and experimental setting associated with the term "cognitive control". Panel (b) tests whether the local task annotations of a sparse dataset cover all task types required by that setting. Panel (c) explains the candidate match by connecting each local task annotation to the component concept whose operationalization requires the corresponding task type.

Table 4
Expected and observed results of the competency-question evaluation.

CQ	Expected result	Observed result	Outcome
CQ1	Both datasets retrieved through the term annotation	ds003452 and ds005270, both associated with "cognitive control"	Pass
CQ2	DMCC and CCCF cognitive-control concepts and their respective concept parts	DMCC: Proactive Control and Reactive Control; CCCF: Inhibition, Task Switching, and Working Memory	Pass
CQ3	Only the theory-explicit dataset annotation	ds003452, DMCC, Cognitive Control	Pass
CQ4	The sparse annotation matched to the CCCF operationalization profile	ds005270, CCCF, and Cognitive Control returned as a candidate interpretation	Pass
CQ5	Three task-level evidence relations for the CCCF candidate	Go/No-Go – Inhibition; letter categorization – Task Switching; N-back – Working Memory	Pass

5.1. SHACL Validation

To account for commonly sparse and heterogeneous neurocognitive metadata, CoTOn includes a lightweight SHACL validation layer in shapes/dataset-annotation-shapes.ttl. The shapes do not require a complete theory-aware description before an annotation can be loaded or queried. Instead, they distinguish structural errors from usable but underspecified annotations. Missing recommended annotations are reported as warnings rather than violations, allowing incomplete metadata to remain operational while making its limitations explicit. This separation supports incremental curation, because curators can identify which theory-, concept-, capacity-, or task-level assertions are still absent without rejecting the annotation as a whole.

For example, the annotation file for ds005270 asserts the term "cognitive control" and three local task annotations but no cognitive theory, theory-specific concept, or capacity type. SHACL therefore reports missing-annotation warnings, while the annotation remains available for querying and candidate enrichment. The validation report can consequently serve as a structured checklist for subsequent review and extension of the annotation.

Validation and querying support complementary workflow stages. SHACL identifies structurally invalid or semantically underspecified annotations, while SPARQL uses theory and operationalization knowledge to generate inspectable candidates. Sparse annotations can thus be loaded, evaluated, and progressively refined without conflating generated suggestions with curated assertions.

6. Discussion and Outlook

CoTOn as an implemented computational artifact demonstrates how theory-aware semantic structures can support neurocognitive dataset annotation and discovery. Its central contribution is not to create yet another cognitive vocabulary, but to explicitly represent relations among terms, theory-specific concepts, capacity types, experimental settings, measurements, and datasets. This matters because (1) a shared term does not necessarily denote the same concept across theories, and (2) a dataset's task annotations do not uniquely determine the authors' intended theoretical interpretation.

Compared to CoTOn's alpha version [2], the implementation presented here introduces a more explicit ontological structure. The earlier model established the need to represent theory-dependent cognitive concepts; the current version unpacks this idea using UFO [3] and OntoUML [14]. It distinguishes cognitive concepts from capacity types and their individual manifestations, represents theory-term-concept and theory-concept-capacity type connections through dedicated relator patterns. It also introduces multi-level patterns for reusable types and their local or execution-level instances. The implementation further adds a modular RDF architecture, theory modules, dataset annotations, a query-oriented KG profile, and SHACL validation.

6.1. Scientific Relevance

The running example illustrates the scientific relevance of preserving theory-specific interpretations. Both datasets in this example can be discovered through the term "cognitive control", but the term alone does not determine whether cognitive control is conceptualized through proactive and reactive control, as in DMCC [8], or through inhibition, shifting, and working memory, as in CCCF [9]. CoTOn preserves these differences while allowing the shared term to remain an entry point for discovery.

This structure supports more transparent dataset comparison and secondary analysis. A researcher may begin with a broad term search and inspect the theories, concepts, component structures, and operationalizations associated with retrieved datasets. Conversely, a researcher interested in a theory-specific concept may search for datasets whose explicit annotations or task profiles are compatible with it. The semantic path remains inspectable: users can see whether retrieval follows an asserted concept, shared term, or task-derived candidate.

Importantly, CoTOn does not establish one theory as preferred or correct. Different theory modules can coexist in one knowledge graph, and the ontology does not collapse their concepts merely because they share a term or broad capacity type. This neutrality is essential for theory-aware infrastructure: the aim is to make theoretical commitments explicit and comparable, not to resolve scientific disagreement ontologically. The same principle applies to candidate annotations. An N-back task may be compatible with a certain operationalization of the concept of cognitive control, but this does not prove that the dataset authors intended a particular theory or conceptualization. Candidate matching is therefore semantic decision support: it indicates plausible interpretations and their supporting assertions, while final decisions remain with researchers or curators.

6.2. Limitations

The evaluation presented here is limited in scope as it only uses two datasets, two illustrative cognitive-control theory modules, and a small set of competency questions. The sparse `ds005270` annotation was intentionally constructed without an asserted theory or concept to test candidate matching. The CQ4 and CQ5 results show that the graph patterns behave as intended for this example – queries retrieve the expected semantic paths and SHACL identifies the intended missing annotations.

The candidate matching query also uses a strict coverage condition: every task type required by the setting profile must match a local task annotation. This gives an interpretable demonstration but may be too restrictive for heterogeneous metadata. Future implementations may distinguish required, optional, and alternative task patterns and incorporate provenance, confidence, or evidence scores without converting uncertain matches into asserted annotations.

6.3. Outlook

The next development stage of CoTOn will extend the number and diversity of theory modules, task mappings, and annotated datasets. Domains beyond the prototypical example of cognitive control are needed to test whether the patterns generalize to constructs with different theoretical structures and experimental traditions. Additional competency questions should address theory comparison, partial operationalization matches, indicator and measurement relations, and cross-infrastructure retrieval.

A second priority is external-user evaluation. Cognitive neuroscience researchers, data curators, ontology engineers, and infrastructure developers should assess whether CoTOn's distinctions are understandable, its annotation workflow practical, and candidate explanations transparent enough for curation. Such studies could compare term-only with theory-aware retrieval and test whether CoTOn identifies conceptually relevant datasets without misleading equivalences between shared labels.

CoTOn may also serve as a theory-aware interpretation layer in federated metadata environments. Integration with initiatives such as *Neurobagel*[4] [20] and alignment with international projects such the *EBrains Knowledge Graph*[5] could connect dataset resource and provenance metadata to theory-indexed cognitive concepts, operationalizations, tasks, and indicators. As a result, OpenNeuro [17] and other Neurobagel nodes could retain their metadata structures while exposing a shared layer for cognitive interpretation and search.

In conclusion, the current implementation of CoTOn shows that theory-dependent cognitive meaning can be represented without sacrificing shared terminology or interoperability. By separating terms from concepts, concepts from capacities, and curated annotations from candidates, CoTOn provides a basis for more transparent annotation, comparison, and discovery of cognitive neuroscience datasets.

Acknowledgments

We thank Pedro Paulo F. Barcelos for sharing his expertise in conceptual modeling and ontology engineering. This work was supported by the Federal State of Salzburg under the grant number F2500775-FPR (GO-CoTOn).

Declaration on Generative AI

During preparation of this work, the authors used DeepL and ChatGPT for: Grammar and spelling check, paraphrasing and rewording, language editing, and code editing. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] A. Ravenschlag, M. Denissen, L. Kronbichler, M. Kronbichler, Orthographic processing, 2 ed., Elsevier, 2025, pp. V2–358–V2–371. doi:10.1016/B978-0-12-820480-1.00139-X.

[2] A. Ravenschlag, B. Löhnert, G. Guizzardi, M. d. G. da Silva Teixeira, M. Denissen, F. Hutzler, Coton: A cognitive theory ontology for representing diverging conceptualizations of cognitive concepts, in: CEUR workshop proceedings, volume 3637, CEUR, 2023.

[3] G. Guizzardi, A. Botti Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, T. Prince Sales, Ufo: Unified foundational ontology, Applied ontology 17 (2022) 167–210.

[4] R. A. Poldrack, A. Kittur, D. Kalar, E. Miller, C. Seppa, Y. Gil, D. S. Parker, F. W. Sabb, R. M. Bilder, The cognitive atlas: toward a knowledge foundation for cognitive neuroscience, Frontiers in neuroinformatics 5 (2011) 17.

[4] https://neurobagel.org/

[5] https://kg.ebrains.eu

[5] B. Löhnert, B. H. Engler, N. Augsten, F. Hutzler, The cognitive task ontology (cogito), 2025. URL: https://doi.org/10.5281/zenodo.17828454. doi:10.5281/zenodo.17828454.
[6] K. Robbins, D. Truong, S. Appelhoff, A. Delorme, S. Makeig, Capturing the nature of events and event context using hierarchical event descriptors (hed), NeuroImage 245 (2021) 118766.
[7] D. Badre, Cognitive control, Annual Review of Psychology 76 (2025) 167–195.
[8] J. A. Etzel, R. E. Brough, M. C. Freund, A. Kizhner, Y. Lin, M. F. Singh, R. Tang, A. Tay, A. Wang, T. S. Braver, The dual mechanisms of cognitive control dataset, a theoretically-guided within-subject task fmri battery, Scientific Data 9 (2022) 114. URL: https://doi.org/10.1038/s41597-022-01226-4. doi:10.1038/s41597-022-01226-4.
[9] J. R. Rieck, G. Baracchini, C. L. Grady, Contributions of brain function and structure to three different domains of cognitive control in normal aging, Journal of Cognitive Neuroscience 33 (2021) 1811–1832. URL: https://doi.org/10.1162/jocn_a_01685. doi:10.1162/jocn_a_01685.
[10] A. Bandrowski, R. Brinkman, M. Brochhausen, M. H. Brush, B. Bug, M. C. Chibucos, K. Clancy, M. Courtot, D. Derom, M. Dumontier, et al., The ontology for biomedical investigations, PloS one 11 (2016) e0154556.
[11] J. Hastings, W. Ceusters, M. Jensen, K. Mulligan, B. Smith, Representing mental functioning: Ontologies for mental health and disease, in: J. Hastings, W. Ceusters, M. Jensen, K. Mulligan, B. Smith (Eds.), Towards an Ontology of Mental Functioning (ICBO Workshop), CEUR, 2012.
[12] J. N. Otte, J. Beverley, A. Ruttenberg, Bfo: Basic formal ontology, Applied ontology 17 (2022) 17–43.
[13] G. Guizzardi, N. Guarino, Explanation, semantics, and ontology, Data & Knowledge Engineering 153 (2024) 102325.
[14] G. Guizzardi, C. M. Fonseca, A. B. Benevides, J. P. A. Almeida, D. Porello, T. P. Sales, Endurant types in ontology-driven conceptual modeling: Towards ontouml 2.0, in: International conference on conceptual modeling, Springer, 2018, p. 136–150.
[15] J. P. A. Almeida, G. Guizzardi, T. P. Sales, C. M. Fonseca, gufo: A gentle foundational ontology for semantic web knowledge graphs, 2026. URL: https://arxiv.org/abs/2603.20948.
[16] C. K. Ogden, I. A. Richards, The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism, Harcourt, Brace, 1930.
[17] C. J. Markiewicz, K. J. Gorgolewski, F. Feingold, R. Blair, Y. O. Halchenko, E. Miller, N. Hardcastle, J. Wexler, O. Esteban, M. Goncavles, et al., The openneuro resource for sharing of neuroscience data, Elife 10 (2021) e71774.
[18] OpenNeuro, Dataset ds003452, https://openneuro.org/datasets/ds003452, 2020. OpenNeuro dataset, accession ds003452.
[19] OpenNeuro, Dataset ds005270, https://openneuro.org/datasets/ds005270, 2024. OpenNeuro dataset, accession ds005270.
[20] S. G. W. Urchs, A. Dai, A. Jahanpour, M. Wang, M. Dugré, N. Bhagwat, B. McPherson, J.-K. Pfarr, E. van Heese, M. A. Laansma, et al., Neurobagel: building an international network for distributed data discovery, 2026. URL: osf.io/preprints/metaarxiv/x64tn_v1.