

eCBRo - Easy Case-Based Reasoning Ontology

Andreas Korger¹, Joachim Baumeister¹

¹University of Würzburg, Am Hubland, D-97074 Würzburg

Abstract

This work presents the first release of eCBRo, an ontology to dynamically semantify Case-Based Reasoning (CBR) scenarios. Structuring knowledge into cases is valuable across a wide range of domains. Often, the objective is not to build a complete CBR application from the outset. Instead, the case base and its underlying structure typically evolve incrementally, often drawing input from heterogeneous and distributed sources before being integrated into fully developed CBR systems. To support this early-stage development, as well as other multi-modal, distributed, and dynamic CBR scenarios, a flexible semantic framework for representing stacked semantic case features and CBR-related knowledge is essential. Therefore, eCBRo takes a lightweight case-centric semantic perspective for dynamic attribute-value representation. The ontology is incrementally introduced, underpinned by examples, and then demonstrated through LLM-assisted case acquisition from a nuclear fire safety document.

Keywords

Case-Based Reasoning, Ontology, Knowledge Graph, Semantics, Case Acquisition, Large Language Models, Semantic WEB, Distributed CBR, Case Structure, Retrieval Augmented Generation

1. Introduction

This paper presents the first (beta) version of eCBRo¹ [1], an easy Case-Based Reasoning [2] (CBR) ontology with a focus on *dynamic attribute value representation*. It provides a *minimal entry barrier* for *dynamically enhanceable* case-based knowledge structuring and openness for the *integration of distributed systems*. CBR is a paradigm based on the principle that similar problems have similar solutions, reusing past experiences encoded in a case structure of problem-solution pairs stored in a case base. Figure 1 illustrates the reasoning process in a 4R-CBR-cycle [3], which includes the components of case retrieval, reuse, revision, and retention to solve a new problem and capture the found solution as a new case in the case base. Typically, creating a CBR application involves designing a fixed case structure fine-tuned to improve performance. Similarity measures and the capacity to adapt existing solutions to new problems are provided. This method serves significant purposes and has been successfully implemented in countless scientific and industrial applications.




Figure 1: 4R CBR cycle by Aamodt and Plaza [3] vs. case collection and Input/Output process.

The far-sighted vision of eCBRo is to provide a semantic foundation for encapsulating case representation, reasoning, and learning with the support of AI models. The presented initial evaluation supports the human user by automatically suggesting attributes and values, capturing them with eCBRo, aiming to reduce the need for both expertise in ontology and CBR modeling.

KI'25: Künstliche Intelligenz, 2025, Potsdam

✉ a.korger@informatik.uni-wuerzburg.de (A. Korger); joba@uni-wuerzburg.de (J. Baumeister)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹we call it "Easy Bro"

Supporting AI models involves integrating textual explanations and meta-knowledge into the ontology's class and property definitions, which enable models to perform CBR tasks more effectively. An additional idea to combine with agentic AI is to incorporate a layer of input and output. These are fundamental to technical data processing protocols, well semantified, and can be used to integrate, e.g., with concepts like the Model Context Protocol [4].

The semantification starts from the case, as shown on the right side of Figure 1. From a knowledge organization perspective, the cases are alignable, similar to a self-organizing data structure, and the 4R concept aligns with "input" and "output" in terms of a collection of cases. The case base is not a "fixed" construct but can be formed by clustering cases, depending on the application scenario. The eCBRO concept simplifies classical CBR in three key areas: the problem, its solution, and the case base. Still, it allows space for ongoing integration and customization according to the 4R principle and the specific needs of the application.

- In many scenarios, the problem and its solution are not clearly defined or are only tacitly understood, even though both may be hidden in the case description or contextual knowledge.
- Problem and solution might evolve over time, and their characteristics manifest after many cases have been created. The problem and solution are dynamically dependent on the context.
- Cases often cannot be meaningfully integrated into a "functional case base" from the outset of a CBR system development. While acquiring cases, it is more of a collection of cases rather than a case base. For instance, a performant indexing strategy is integrated ex-post.
- Subsequently, the sub-tasks of the 4R cycle can only be clearly defined if the previous characteristics are established and well understood.
- Furthermore, challenges of technical and societal development require integrating diverse and distributed systems. For instance, in electronic healthcare, crisis management, international trade, and agentic AI automation.
- In critical systems, functional elements like centralized reasoning may fail; redundancy and flexibility in the residual elements aid in maintaining functionality.

The presented ontology aims to be used with the intention of the following advantages:

- Minimum efforts to start modeling, focusing on semantic knowledge representation.
- Modular semantic case design adaptable to any arbitrary case structure.
- Integrates with any other CBR modeling ontology or framework.
- Attribute-Value pairs can be reused and shared.
- Cases can be incrementally extended with new attribute-value pairs.
- Merging of different CBR systems in distributed and disaggregated settings.
- Agent-oriented modeling support by integrating multi-modal and reasoning models.
- Free and flexible case order, such as sequential, hierarchical, or graph-based.
- Preprocess workflow for later CBR integration.

The rest of this paper is organized as follows. The next section covers related work. Then, the methods for defining the ontology model with classes and properties are described step by step. The model is demonstrated using examples from the fire safety domain. The Results section will present a simple application scenario based on a nuclear fire safety document. It will show how a case structure is created with the help of an LLM in both zero- and few-shot scenarios. The paper concludes with a discussion and potential directions for future research.

For complete semantic definitions, please see the ontology reference under www.ecbro.org [1].

2. Related Work

The growing capabilities of AI models influence ontology and CBR modeling. To gain an introductory overview of the open challenges and research directions in this regard, the CBR research community aggregated and structured the topic in the research manifesto by Bach et al. [5]. A noteworthy complementary contribution is the overview of state-of-the-art CBR systems that was presented by Schultheis et al. [6]. The integration of AI models and CBR is crucial for leveraging the eCBRo ontology, as it supports CBR subtasks such as case acquisition, attribute evaluation, similarity assessment, and case maintenance. In this context, several researchers have explored improving the flexibility and application of the CBR paradigm with technical support. The Python toolkit CBRkit for realizing CBR applications was presented by Lenz et al. [7]. Sourati et al. [8] integrated CBR and LLMs for logical fallacies. Gould et al. [9] elaborated on preference-based CBR for argumentation. Dannenhauer et al. [10] investigated the use of CBR for creating few-shot prompts. Hatalis et al. [11] discussed the influence of LLMs and CBR Agents. Florio et al. [12] presented their work on inferring legal CBR. All these approaches can be mirrored in eCBRo. The approach to defining cases based on attributes and their values has been enhanced by additional strategies of case design. The flexible definition of eCBRo allows adaptation to such methods. Related works in this context were presented by Malburg et al. [13] in CBR for Time Series beyond attribute-value representations.

The integration of CBR and knowledge graphs is a well-investigated field, dating back to the development of the Semantic Web [14]. There exist ontology models and knowledge graph integrations with CBR that aim in the same direction as this work. CBR_{Onto}² presented by Diaz [15] provides the most sophisticated of the related ontology models with a specific focus and precise semantification of CBR methods. This solution is recommended for particular fine-tuned CBR systems. CBR_{Onto} develops from the case base, whereas eCBRo aims to develop from the case and exploits relation definitions between cases, which can be used to model a case base by relating cases with relational classes. OntoCase, presented by Blomqvist [16], is an approach that employs CBR for ontology construction. The work does not address case semantification, yet it provides interesting discussions from another knowledge modeling perspective. S-CBR was presented by Thakker et al. [17]. They created a knowledge-based framework to make cases accessible as web services. They focus on the travel domain but leave space for a generic use of their knowledge graph design. Further work towards decentralized CBR using the semantic web was presented by Aquin et al. [18]. They present interesting ideas on how to achieve CBR with distributed description logics and decentralized CBR. Case Retrieval Nets as a Model for Building Flexible Information Systems is a noteworthy contribution by Lenz [19]. The work on CBR in knowledge-intensive domains by Aamodt [20] provides foundational ideas on how to integrate such domains into a knowledge graph.

The focus on CBR for diverse and dynamic case representation was given by the following authors in various aspects. These works inspired the semantification process by highlighting pitfalls in practical application and real-world scenarios of CBR modeling. A work towards dynamic case maintenance and evolving cases was presented by Smiti et al. [21]. Xi et al. [22] developed an approach for dynamic CBR with a process-oriented perspective. CBR for dynamic choices was investigated by Thomas et al. [23]. Dynamic case creation and expansion was the subject of an early work of Mostek et al. [24]. Using CBR for building retrofit is a noteworthy strategy presented by Li et al. [25]. Navigating CBR for case fidelity presented an interesting concept by Parsodkar et al. [26]. Dynamic CBR case base by [27]. The concept of a CBR cloud was discussed by Nisi et al. [28] and developed into a framework that employs a microsystem environment to implement different parts of the CBR process. Their work makes an interesting contribution to a dynamic and evolving CBR strategy.

²The namespace `cbr:` was claimed by CBR_{Onto} and is used solely in that context in the following examples.

3. Methods

The eCBRO model is formally defined as an OWL ontology and expressed as RDF triples. The complete eCBRO reference can be accessed under <http://www.ecbro.org>. A case is considered an arbitrary knowledge entity that is simply defined by its creation with a unique name. The case is then incrementally enriched with attributes and values. This enables the integration of arbitrary case characterizations, such as structured information, natural language, process data, multimodal data, and semantic embeddings. For such integration, mapping of attributes, integrity maintenance, and more, it relies on application-based alignment with third-party data models. Examples include the SKOS (Simple Knowledge Organization System) ontology [29], which can be used to organize cases in a hierarchical manner, and the BPMO (Business Process Modeling) ontology [30], which can integrate process-oriented case definition and organization.

3.1. The Case Class

Figure 2 illustrates the core class `ecbro:Case`. The aggregation of cases can be achieved by aligning cases directly to each other. A collection of cases is selected over the relating attributes, for instance, by selecting one specific case and all following cases linked to this case with the property `ecbro:next`. A versatile approach is to align weights or semantic concepts to the edges relating cases. Table 1 shows a selection of properties of common use, and Listing 1 provides an exemplification.




Figure 2: The case class and a relating property.

Table 1

Selected properties relating cases to other cases.

Property	Domain	Range
<code>ecbro:related</code>	<code>ecbro:Case</code>	<code>ecbro:Case</code>
<code>ecbro:next</code>	<code>ecbro:Case</code>	<code>ecbro:Case</code>
<code>ecbro:previous</code>	<code>ecbro:Case</code>	<code>ecbro:Case</code>
<code>ecbro:partOf</code>	<code>ecbro:Case</code>	<code>ecbro:CaseBase</code>

```
ecbro:case1 a ecbro:Case ;
  rdfs:label "Inspection Case 1"@en ;
  rdfs:comment "School inspection found plant in front of door."@en ;
  dct:description "The fire door should be easily accessible."@en .

ecbro:case2 a ecbro:Case ;
  rdfs:label "Inspection Case 2"@en ;
  rdfs:comment "School inspection found bicycles badly parked."@en ;
  dct:description "The fire ladder should be easily accessible."@en .
ecbro:previous ecbro:case1 ;
```

Listing 1: Example for the definition of two cases and a relation.

3.2. Defining Attributes And Values

The next step of the semantification defines arbitrary attribute-value pairs that can be customized to fit the specific case-based application CBR concepts as `ecbro:CBRConcepts`. The case can now be integrated into a `ecbro:caseBase` with `ecbro:partOf` as illustrated by Figure 3. In this manner, attributes and their values can be assigned to integrate with any combination of third-party semantification. Table 2 shows an ontology stack suggesting third-party models for selected tasks and domains.

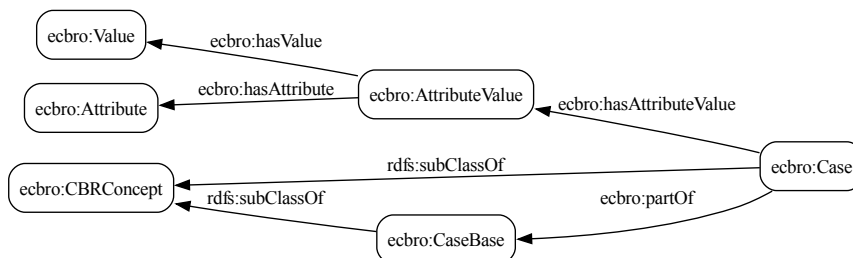


Figure 3: Defining attributes, values, and CBR concepts.

Table 2

Selected ontology stack for attribute-value semantification.

Domain	Ontology
Entities, concepts, knowledge	FOAF [31], PROV [32], SKOS [29]
Processes	OWL-S [33], PDL [34], BPMO [35], PROV [32]
Emdbdings, Machine Learning	ML-Schema [36], MLOnto [37]
Case-Based Reasoning	CBROnto [15], CaseML [38], OntoCase [16], S-CBR [17]
Time, Spatial, Location, Quantities	GEO [39], Geonames [40], TSO [41], QUDT [42]
Distributed systems	DUL [43], BLONDIE [44], TOSCA [45], ONS [46]

Listing 2 exemplifies this principle by extending the previously defined cases. The cases are now linked to a case base and a simplified selection of attributes with values.

```
ecbro:case2 a ecbro:Case ;
  rdfs:label "case 2"@en ;
  rdfs:comment "school inspection found bicycles parked badly."@en ;
  dct:description "The fire ladder should be easily accessible."@en .
ecbro:partOf ecbro:caseBase1 ;
ecbro:previous ecbro:case1 ;
ecbro:hasAttributeValue [foaf:Person, "Inspector Bob" ] ,
  [prov:Entity, "Fire safety inspection 1" ] ,
  [geo:Location, "eCBRO School" ] ,
  [skos:ConceptScheme, eCBRO:fireSafetyInspections] ;

ecbro:hasAttributeValue
[
  a ecbro:AttributeValue ; ecbro:hasAttribute foaf:Person ; ecbro:hasValue "inspector"
] .
```

Listing 2: Example for the definition of two multi-modal cases.

3.3. The CBR Vocabulary

Figure 4 illustrates the extendability of the core model to the CBR-typical terms and vocabulary. Cases include a description, a defined problem, and a solution. eCBRO pre-defines the corresponding subclasses of the class `ecbro:AttributeValue`, such as `ecbro:Solution`, along with the associated sub-properties, such as the property `ecbro:hasSolution`. This basic nomenclature can be adjusted to specific needs in a similar manner. These "official" CBR terms are also defined as subclasses of the CBR concept (only illustrated for `ecbro:Problem`); alternatively, they can be derived from other CBR ontologies, such as CBRonto [15]. Listing 3 extends the previously defined cases as described;

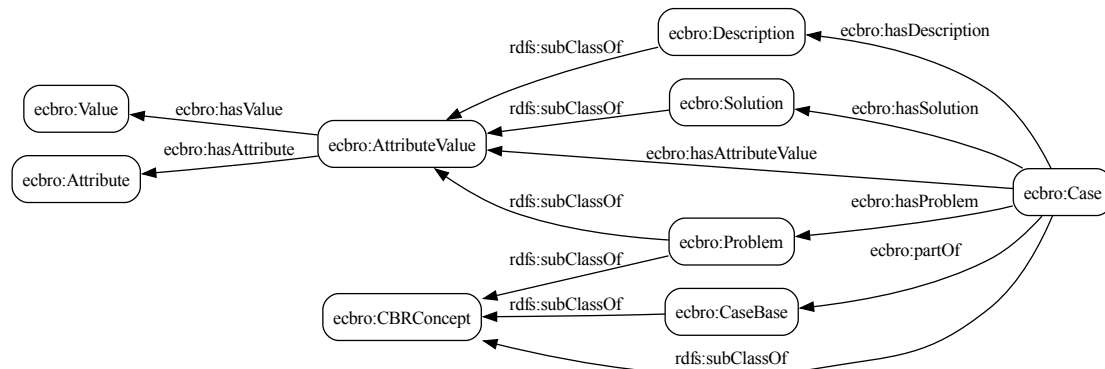


Figure 4: Defining the basic CBR concepts.

```

ecbro:case2 a ecbro:Case ;
  ecbro:partOf ecbro:CaseBase ;
  ecbro:hasDescription "If a fire happens, put it out with water." ;
  ecbro:hasProblem "fire" ;
  ecbro:hasSolution "water" ;
  ecbro:previous ecbro:case1 ;
  ecbro:hasAttributeValue [ecbro:SemanticEmbedding, ecbro:case2-embedding] ;
  ecbro:hasAttributeValue [cbr:SimilarityAssessment, ecbro:similarityAss1] .

ecbro:case2-embedding a ecbro:SemanticEmbedding ;
  ecbro:hasFormat "float32" ;
  ecbro:hasDimension "384"^^xsd:integer ;
  ecbro:hasRawRepresentation "[0.125, -0.023, 0.991, ...]" .

ecbro:similarityAss1 a cbr:SimilarityAssessment ;
  cbr:usesSimilarityMeasure
  [ a cbr:SimilarityMeasure ;
    cbr:appliesTo ecbro:case2-embedding ;
    cbr:hasFunction "cosine" ;
    cbr:hasWeight "0.7"^^xsd:float ;
    cbr:usedFor "semantic comparison" ] ,
  [ a cbr:SimilarityMeasure ;
    cbr:appliesTo foaf:person, prov:entity, geo:location ;
    cbr:hasFunction "string-equality" ;
    cbr:hasWeight "0.3"^^xsd:float ;
    cbr:usedFor "attribute comparison" ] ;
  cbr:usedInTask "hybrid retrieval" .
  
```

Listing 3: Example for a hybrid case definition with "classical" CBR semantification.

3.4. Client-Case-Communication

Until now, the modeling approach has been static. It is now extended with explicit definitions of activities and an emphasis on ensuring that each case remains self-contained while supporting interactions with the case base and external clients. Many established standards in computer science, particularly protocol and API specifications, are grounded in the concepts of input and output (I/O). To align with these standards, the 4R is abstracted into simpler I/O operations. This abstraction facilitates the modeling of scenarios involving both synchronous and asynchronous access to cases, while preserving semantic and functional correctness, even in situations where the case base may be temporarily unavailable. Figure 5 illustrates a very easy semantification, where a case is `ecbro:output` to a client and then `ecbro:input` back in. This actually refers to information contained in the case being output and input. Initially, a case can be created by a client.




Figure 5: Communicating directly with the case.

Figure 6 illustrates how the activities are modeled that can hold particular information. Then, Table 3 outlines which semantic components are involved in the 4R CBR components, followed by an exemplification in Listing 4.




Figure 6: Linking input and output activities to the case.

Table 3

Semantification of 4R concepts using I/O concepts.

Property	Involves
<code>ecbro:retrieve</code>	<code>ecbro:input</code> , <code>ecbro:output</code> , <code>ecbro:Client</code>
<code>ecbro:reuse</code>	<code>ecbro:Client</code> , <code>ecbro:Output</code>
<code>ecbro:revise</code>	<code>ecbro:Client</code> , <code>ecbro:Output</code>
<code>ecbro:retain</code>	<code>ecbro:input</code> , <code>ecbro:Input</code>

```
ecbro:clientA ecbro:input ecbro:case1 .
```

```
ecbro:solution1 a ecbro:Output .
```

```
ecbro:problem1 a ecbro:Input .
```

```
ecbro:case1 ecbro:hasInput ecbro:problem1 .
```

```
ecbro:case1 ecbro:hasOutput ecbro:solution1 .
```

Listing 4: Example for I/O activities on cases.

4. Results

In this section, an application scenario for instantiating the ontology model is demonstrated. Textual case acquisition from safety texts can be semantified with the support of an LLM (Mistral Large 2.1 [47]). Case-based representation exists in manifold domains, and thereby cases are often tacitly hidden in texts and other modalities of knowledge. This virtually unites into a cloud of cases. Depending on the application, cases are clustered and defined by attribute-value pairs, text, process data, and other relevant factors. Nuclear safety is of great importance and can be broken down into several subdomains, such as nuclear fire safety and the safe transportation of nuclear materials. This virtual cloud of cases is illustrated in Figure 7 for the domains of fire, transportation, and nuclear safety.




Figure 7: Cloud of virtual cases from the domains of fire, transportation, and nuclear safety.

Authorities oversee this sector, and one outstanding publisher is the International Atomic Energy Agency [48], a subsidiary of the United Nations. The IAEA maintains a corpus through the efforts of human domain experts. The texts code the essence of experiences made in the past with the intention of regulation, reuse, and guidance. The speech is widely prototypical [49] as a blueprint for regulatory corpora in other affected domains. A related task is to recontextualize subcomponents for their specific application, for instance, from generic fire safety advice to specific advice for a particular type of nuclear reactor. Example 1 shows an excerpt of such a document for fire safety in nuclear power plants.

Example 1. *17 9.2. Specialized fire safety training should be established for designated plant staff involved in operations, maintenance, and firefighting at the plant, including contractors' personnel temporarily assigned to the plant, where applicable. The training programme should provide training to ensure that the staff have adequate technical skills and familiarity with the detailed procedures to be followed. Training should be sufficient to ensure that individuals understand the significance of their duties and the consequences of errors arising from misconceptions or lack of diligence. The specialized training programme should cover the following:*

- (a) The importance of maintaining the integrity and operability of plant fire protection features (both passive and active) by performing regularly scheduled inspections, routine and unplanned maintenance of equipment, and periodic functional tests of equipment and systems;*
- (b) The design and operating details of the specific fire protection features installed in the plant to permit effective maintenance of equipment for operability;*
- (c) The significance of planned design changes and plant modifications with respect to fire safety, including both direct and indirect impacts on fire safety and any effects on the integrity or operability of the fire protection features(both passive and active) as a result of the planned modifications;*

From this document, a few-shot scenario was created by providing an LLM with some cases defined manually according to the given text. Listing 5 shows the definition of two cases from paragraph 17.9.2, and its subparts (a) and (b). The cases were enriched with contextual information.

```
ecbro:case-iaea-fs-17-9-2-a a ecbro:Case ;
  rdfs:label "Case IAEA Fire Safety 17 9.2. (a)"@en ;
  rdfs:comment "Emphasizes integrity and operatibility."@en ;
  ecbro:partOf ecbro:IAEACaseBase, ecbro:IAEAFireSafetyCaseBase ;
  ecbro:next ecbro:case-iaea-fs-17-9-2-c ;
  ecbro:upper ecbro:case-iaea-fs-17-9-2 ;
  ecbro:hasDescription "The importance of maintaining the integrity and operability of plant fire
    protection features (both passive and active) by performing regularly scheduled inspections,
    routine and unplanned maintenance of equipment, and periodic functional tests of equipment and
    systems;"@en ;
  ecbro:hasProblem ecbro:UnknownProblemDescription ;
  ecbro:hasSolution "maintenance", "inspections", "functional tests" ;
  ecbro:hasAttributeValue [foaf:Person, "inspector"] ,
    [prov:Entity, "fire inspection document 1"] ,
    [geo:Location, "Nuclear Power Plant"] .
```

Listing 5: Example for the definition of two modular cases with eCBRO.

This handcrafted case is uniquely identified within its scope. A textual label and comment are provided for human understanding. The case's belonging to the potential case bases is indicated. As described, the original text is defined. The problem is identified as unknown; it cannot be directly derived from the text, possibly from the larger context. Textual solution components are extracted from the descriptive text, along with additional semantically annotated attributes with textual values.

Then, an LLM was prompted to structure the text into cases analogously to the manually provided ones. To define the attributes in a manner that fits the current textual characteristic, provide an in-context description of the case to make it reusable by machine processing, and align the cases hierarchically for the whole document. The stacked prompt is organized as follows:

1. **Agent role:** You assume the role of an intelligent agent responsible for extracting structured cases from input texts. These cases are intended for use in ontology-enhanced case-based reasoning systems, such as myCBR [50].
2. **Few-shot examples:** Several representative cases are provided as input to guide your task.
3. **Case structure and annotations:** The example cases use case-based reasoning terminology, semantic attribute-value pairs, and annotations from external ontologies (e.g., FOAF [31], PROV [32], GEO [32], PIRI [51]).
4. **Attribute selection:** When generating new cases, identify and extract attributes and values that best fit the case and context, leveraging both in-context cues and background knowledge.
5. **Indexing:** Assign each case an indexing structure conducive to similarity-based retrieval or ontology-based navigation.
6. **Data types:** Where possible, use heterogeneous attributes, including not only textual descriptors but also temporal, geospatial, numeric, chemical, or biomedical features.
7. **Inter-case relations:** Specify semantic links (e.g., causality, temporal sequence, similarity) between cases to support advanced retrieval strategies and case base navigation.
8. **Output:** Present a summary of the generated case(s) that supports human evaluation, for example, by listing the extracted attributes along with their semantic annotations. Additionally, provide the complete structured output using the Turtle [52] serialization of RDF, ensuring proper namespace declarations and ontology-aligned predicate usage.

The goal of this experiment is to demonstrate the usage of eCBRo for quickly providing a case structuring knowledge graph. The results of the LLM support were promising but still require more fine-grained prompt construction and revision cycles. An excerpt of the results is given by Listing 6.

Case: IAEA Fire Safety 17 9.2

Label: "Case IAEA Fire Safety 17 9.2"

Comment: "Establishes specialized fire safety training for plant staff."

Part Of: IAEA Case Base, IAEA Fire Safety Case Base

Description: "Specialized fire safety training should be established for designated plant staff involved in operations, maintenance, and firefighting at the plant, including contractors' personnel temporarily assigned to the plant, where applicable. The training programme should provide training to ensure that the staff have adequate technical skills and familiarity with the detailed procedures to be followed. Training should be sufficient to ensure that individuals understand the significance of their duties and the consequences of errors arising from misconceptions or lack of diligence."

Problem: Unknown Problem

Solution: "training programme", "technical skills", "familiarity with procedures"

Attribute-Value Pairs:

```
[foaf:Person, "plant staff"]  
[foaf:Person, "contractors' personnel"]  
[prov:Entity, "training programme"]  
[prov:Activity, "fire safety training"]  
[geo:Location, "Nuclear Power Plant"]
```

Case: IAEA Fire Safety 17 9.2 (b)

Label: "Case IAEA Fire Safety 17 9.2 (b)"

Comment: "Details design and operating specifics of fire protection features."

Part Of: IAEA Case Base, IAEA Fire Safety Case Base

Next: case-iaea-fs-17-9-2-c

Upper: case-iaea-fs-17-9-2

Description: "The design and operating details of the specific fire protection features installed in the plant to permit effective maintenance of equipment for operability."

Problem: Unknown Problem

Solution: "design details", "operating details", "effective maintenance"

Attribute-Value Pairs:

```
[foaf:Person, "maintenance personnel"]  
[prov:Entity, "fire protection features"]  
[geo:Location, "Nuclear Power Plant"]
```

Case: IAEA Fire Safety 17 9.2 (c)

Label: "Case IAEA Fire Safety 17 9.2 (c)"

Comment: "Significance of planned design changes and plant modifications."

Part Of: IAEA Case Base, IAEA Fire Safety Case Base

Next: case-iaea-fs-17-9-2-d

Upper: case-iaea-fs-17-9-2

Description: "The significance of planned design changes and plant modifications with respect to fire safety, including both direct and indirect impacts on fire safety and any effects on the integrity or operability of the fire protection features (both passive and active) as a result of the planned modifications."

Problem: Unknown Problem

Solution: "planned design changes", "plant modifications", "fire safety impacts"

Attribute-Value Pairs:

```
[foaf:Person, "design reviewer"]  
[prov:Entity, "design changes"]  
[prov:Entity, "plant modifications"]  
[geo:Location, "Nuclear Power Plant"]
```

Listing 6: Example for the definition of two modular cases with eCBRo.

5. Conclusion and Future Work

This work presented an ontology to represent cases and CBR components semantically. The semantification approach had two aims. First, to allow modeling with a minimum entry barrier and maximum flexibility, and second, to provide a simplified semantification alignment with the classical 4R CBR cycle. After discussing related work, the core principles and methods were explained.

The layered structure of the ontology was demonstrated with examples of increasing complexity. This demonstrated an evolution from a simple "beginner" to a CBR expert. It could also be demonstrated that facilitating the CBR process can maintain core functionality while still allowing for the option to incorporate the full CBR methodology. It should be emphasized that this work presents the initial version of the ontology, leaving room for improvements and validation, particularly in the context of distributed applications. If the model is capable of mapping a complex CBR application installed in a CBR system to a distributed setting, further investigation is needed. Especially in terms of performance and maintenance, more practical experiments in different application domains should be conducted with the provision of convenient performance measures.

For future work, three research directions seem promising from the current perspective. Validating the current version of the model for robustness and integrity, and identifying potential improvements for the next released version. Exploring additional application scenarios to demonstrate the approach's flexibility is also beneficial. Enhancing the performance of automated and dynamic case acquisition, including agentic AI support and reasoning about the collected structure to predict shortcomings in later CBR use, should offer advantages.

6. Declaration on Generative AI

During the preparation of this work, the authors used Grammarly and ChatGPT in order to check Grammar and spelling. The authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] A. Korgar, eCBRO Ontology Reference, 2025. <https://www.ecbro.org>, accessed 2025-06-24.
- [2] J. Kolodner, Case-based reasoning, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.
- [3] A. Aamodt, E. Plaza, Case-based reasoning: Foundational issues, methodological variations, and system approaches, *AI Communications* 7 (2001) 39–59.
- [4] Anthropic, Model context protocol (mcp), 2024. URL: <https://docs.anthropic.com/en/docs/agents-and-tools/mcp>, accessed 2025-06-24.
- [5] K. Bach, R. Bergmann, F. Brand, M. Caro-Martínez, V. Eisenstadt, M. W. Floyd, L. Jayawardena, D. Leake, M. Lenz, L. Malburg, D. H. Ménager, M. Minor, B. Schack, I. Watson, K. Wilkerson, N. Wiratunga, Case-Based Reasoning Meets Large Language Models: A Research Manifesto For Open Challenges and Research Directions, 2025. URL: <https://hal.science/hal-05006761>.
- [6] A. Schultheis, C. Zeyen, R. Bergmann, An overview and comparison of case-based reasoning frameworks, in: *Case-Based Reasoning Research and Development: 31st International Conference, ICCBR 2023, Aberdeen, UK, July 17–20, 2023, Proceedings*, Springer-Verlag, Berlin, Heidelberg, 2023, pp. 327–343.
- [7] M. Lenz, L. Malburg, R. Bergmann, CBRkit: An intuitive case-based reasoning toolkit for python, in: J. A. Recio-Garcia, M. G. Orozco-del Castillo, D. Bridge (Eds.), *Case-Based Reasoning Research and Development*, Springer Nature Switzerland, Cham, 2024, pp. 289–304.
- [8] Z. Sourati, F. Ilievski, H.-Å. Sandlin, A. Mermoud, Case-based reasoning with language models for classification of logical fallacies, in: E. Elkind (Ed.), *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23, International Joint Conferences on Artificial Intelligence Organization, 2023*, pp. 5188–5196.
- [9] A. Gould, G. Paulino-Passos, S. Dadhania, M. Williams, F. Toni, Preference-Based Abstract Argumentation for Case-Based Reasoning, in: *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning, 2024*, pp. 394–404.
- [10] D. C. K. H. Dustin Dannenhauer, Zohreh Dannenhauer, A case-based reasoning approach to dynamic few-shot prompting for code generation, in: *ICML 2024 Workshop on LLMs and Cognition, 2024*.
- [11] K. Hatalis, D. Christou, V. Kondapalli, Review of case-based reasoning for llm agents: Theoretical foundations, architectural components, and cognitive integration, 2025. URL: <https://arxiv.org/abs/2504.06943>. arXiv:2504.06943.
- [12] C. Florio, X. Liu, E. Lorini, A. Rotolo, G. Sartor, Inferring New Classifications in Legal Case-Based Reasoning, volume 379 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2023.
- [13] L. Malburg, A. Schultheis, Modeling and using complex iot time series data in case-based reasoning: From application scenarios to implementations, in: *Workshops Proceedings for the 31st International Conference on Case-Based Reasoning co-located with the 31st International Conference on Case-Based Reasoning (ICCBR 2023), Aberdeen, Scotland, UK, July 17, 2023At: Aberdeen, Scotland, 2023*.
- [14] T. Berners-Lee, J. Hendler, O. Lassila, The semantic web: A new form of web content that is meaningful to computers will unleash a revolution of new possibilities, *ScientificAmerican.com* 284 (2001) 34–43.
- [15] B. Díaz Agudo, CBRonto: A task/method ontology for CBR., in: *Proceedings of the Fifteenth International Florida Artificial Intelligence Research Society Conference, 2002*, pp. 101–105.
- [16] E. Blomqvist, Ontocase-automatic ontology enrichment based on ontology design patterns, in: A. Bernstein, D. R. Karger, T. Heath, L. Feigenbaum, D. Maynard, E. Motta, K. Thirunarayan (Eds.), *The Semantic Web - ISWC 2009, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009*, pp. 65–80.
- [17] D. Thakker, T. Osman, D. Al-Dabass, S-CBR: Semantic case based reasoner for web services discovery and matchmaking, in: *Proceedings 20th European Conference on Modelling and Simulation, 2006*.
- [18] M. d’Aquin, J. Lieber, A. Napoli, Case-based reasoning within semantic web technologies, in:

- J. Euzenat, J. Domingue (Eds.), *Artificial Intelligence: Methodology, Systems, and Applications*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 190–200.
- [19] M. Lenz, *Case Retrieval Nets as a Model for Building Flexible Information Systems*, Ph.D. thesis, Humboldt Universität zu Berlin, 1999.
- [20] A. Aamodt, Knowledge-intensive case-based reasoning in CREEK, in: P. Funk, P. A. González Calero (Eds.), *Advances in Case-Based Reasoning*, Springer, Berlin Heidelberg, 2004, pp. 1–15.
- [21] A. Smiti, Z. Elouedi, Dynamic maintenance case base using knowledge discovery techniques for case based reasoning systems, *Theoretical Computer Science* 817 (2020) 24–32.
- [22] Q. Xia, M. Rao, Dynamic case-based reasoning for process operation support systems, *Engineering Applications of Artificial Intelligence* 12 (1999) 343–361.
- [23] P. Thomas, T. Guilfoos, Case-based reasoning and dynamic choice modeling, *Land Economics* 99 (2022) 103–121.
- [24] T. Mostek, K. D. Forbus, C. Meverden, Dynamic case creation and expansion for analogical reasoning., in: *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, July 30 - August 3, 2000, Austin, Texas, USA., 2000, pp. 323–329.
- [25] Y. Li, H. Du, S. Kumaraswamy, Case-based reasoning approach for decision-making in building retrofit: A review, *Building and Environment* 248 (2023) 111030.
- [26] A. P. Parsodkar, D. P., S. Chakraborti, Navigating the landscape of case fidelity and competence in case-based reasoning, in: *Artificial Intelligence XLI: 44th SGAI International Conference on Artificial Intelligence*, AI 2024, Cambridge, UK, December 17–19, 2024, *Proceedings, Part I*, Springer-Verlag, Berlin, Heidelberg, 2024, pp. 235–249.
- [27] I. Ettahiri, K. Doumi, A novel dynamic enterprise architecture model: leveraging MAPE-K loop and case-based reasoning for context awareness, *International Journal of Electrical and Computer Engineering (IJECE)* 14 (2024) 1875–1889.
- [28] I. Watson, R. Weber (Eds.), *Cloud CBR: Towards Microservices Oriented Case-Based Reasoning*, Springer International Publishing, Cham, 2020.
- [29] W3C, SKOS Simple Knowledge Organization System Reference, 2009. <http://www.w3.org/TR/skos-reference> Accessed 2025-06-06.
- [30] X. Zhao, Y. Wu, J. Yu, Y. Zhang, BPMO: Semantic business process modeling and annotation, in: *Proceedings of the IEEE International Conference on Services Computing (SCC 2007)*, IEEE, 2007, pp. 442–449.
- [31] D. Brickley, L. Miller, Foaf vocabulary specification 0.99, <http://xmlns.com/foaf/0.1/>, 2014. Namespace Document describing the FOAF RDF vocabulary.
- [32] W3C, PROV-O: The PROV Ontology, 2013. <http://www.w3.org/TR/prov-o> Accessed 2025-06-06.
- [33] D. Martin, M. Paolucci, S. McIlraith, M. Burstein, D. McDermott, D. McGuinness, B. Parsia, T. Payne, et al., OWL-S: Semantic markup for web services, <https://www.w3.org/Submission/OWL-S/>, 2004. W3C Member Submission.
- [34] D. K. S. Initiative, *Process Description Language (PDL) Specification*, Technical Report, DARPA, 1997. URL: <https://www.daml.org/2001/03/pddl/>, accessed: 2025-06-27.
- [35] W. M. van der Aalst, A. H. ter Hofstede, M. Weske, Business process modeling ontology, in: *International Conference on Business Process Management*, Springer, 2005, pp. 205–219.
- [36] K. Janowicz, P. Hitzler, J. Wilkes, K. Taylor, D. Garijo, I. Toma, A.-C. N. Ngomo, S. Etzold, S. Auer, M.-E. Vidal, ML-Schema: An extensible schema for sharing machine learning metadata, in: *Proceedings of the 2019 International Semantic Web Conference (ISWC)*, 2019. URL: <https://w3id.org/mls/>, accessed: 2025-06-24.
- [37] J. Braga, J. Dias, F. Regateiro, A machine learning ontology, <https://osf.io/chu5q/>, 2020. Accessed 2025-06-24.
- [38] H. Chen, Z. Wu, On case-based knowledge sharing in semantic web, *Proceedings. 15th IEEE International Conference on Tools with Artificial Intelligence (2003)* 200–207.
- [39] W. W. W. Consortium, Wgs84 geo positioning: an rdf vocabulary, <https://www.w3.org/2003/01/>

- geo/, 2003. Accessed 2025-06-24.
- [40] G. Team, Geonames ontology, <http://www.geonames.org/ontology/documentation.html>, 2006. Accessed 2025-06-24.
- [41] D. Pufal, M. Hepp, A transport ontology for the semantic web, in: Proceedings of the 10th International Conference on Knowledge Engineering and Knowledge Management (EKAW), Springer, 2006, pp. 189–202.
- [42] M. Fox, R. Hodgson, P. Yim, QUDT: Quantities, units, dimensions and data types ontologies, <http://www.qudt.org/>, 2010. Accessed: 2025-06-24.
- [43] A. Gangemi, V. Presutti, Ontology design patterns for semantic web content, in: A. P. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. Finin, K. Thirunarayan (Eds.), The Semantic Web - ISWC 2008, volume 5318 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 262–276.
- [44] F. Johansson, O. Hammar, J. Stirna, BLONDIE: An ontology for describing blockchain data, in: Proceedings of the 16th International Conference on Semantic Systems (SEMANTiCS 2020), volume 12262 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 231–247.
- [45] OASIS, Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0, Technical Report, OASIS Standard, 2013. URL: <https://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf>, oASIS Standard.
- [46] P. Jain, M. P. Singh, Ontology-based information sharing in cyber-defense, in: Proceedings of the 2006 Semantic Web for Collaborative Defense Workshop (SWCD), Arlington, VA, USA, 2006. Describes the ONS (Ontology for Network Security).
- [47] MISTRAL, 2025. <https://mistral.ai/> Accessed 2025-06-06.
- [48] International Atomic Energy Association, 2025. <https://www.iaea.org> Accessed 2025-06-06.
- [49] A. Korger, J. Baumeister, Case-based sample generation using multi-armed bandits, in: S. Massie, S. Chakraborti (Eds.), Case-Based Reasoning Research and Development: 31st International Conference, ICCBR 2023, Aberdeen, UK, Springer-Verlag, Berlin, Heidelberg, 2023, pp. 118–133.
- [50] K. Bach, K.-D. Althoff, Developing case-based reasoning applications using mycbr 3, in: B. D. Agudo, I. Watson (Eds.), Case-Based Reasoning Research and Development, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 17–31.
- [51] PIRI Ontology Reference, 2025. <https://www.piri-safety.com> Accessed 2025-06-06.
- [52] W3C, Turtle Reference, 2014. <https://www.w3.org/TR/turtle/> Accessed 2025-06-06.