

Towards a Gold Dataset for Cross-Perspective Ontology Alignment in the Food Domain

Divyasha Sunil Naik¹, Birgitta König-Ries¹

¹Heinz Nixdorf Chair for Distributed Information Systems, Friedrich Schiller University Jena, Jena, Germany

Abstract

Ontology alignment serves as the backbone of semantic interoperability among heterogeneous knowledge resources. However, many real-world domains are modeled from multiple perspectives, depending on the purpose, community, and context of use. Existing benchmarks focus primarily on ontologies with similar modeling goals. This limits the evaluation of alignment methods across different perspectives. To address this gap, this paper presents an ongoing effort to develop a small, expert-oriented benchmark dataset for cross-perspective ontology alignment in the food domain. The dataset is created from selected subsets of three existing ontologies representing nutrition and diet, food quality and public health, and food-service occupations using the ROBOT tool. These heterogeneous ontological commitments often produce meaningful semantic relations that go beyond equivalence matching.

The workflow includes seed-term-based concept selection, extraction of labels, definitions, superclass context, properties, and restrictions, followed by expert annotation of semantic bridge relations across perspectives. The resulting dataset is intended to support the evaluation of ontology alignment method cross different standpoints.

Keywords

Food ontologies, Semantic heterogeneity, Cross-perspective alignment dataset, Benchmark creation

1. Introduction

Meaning is rarely fixed. The same term can have different meanings in different contexts, communities, and use cases. When such meanings are formalized in ontologies, these choices become explicit conceptual commitments. Ontologies therefore do not only describe a domain. They also encode particular perspectives, assumptions, and modeling decisions [1, 2].

Ontology alignment supports interoperability between heterogeneous knowledge systems. Existing approaches and benchmarks mainly focus on equivalence or hierarchical relations between ontologies with similar modeling goals [3]. This is limiting when ontologies describe the same phenomenon from different viewpoints. In these cases, concepts may be meaningfully related even when they are not synonyms or subclasses of one another. For example, the concept *Food Contamination* may be linked to *Food Product* through the relation *affects*. In this work, such a relation is referred to as a semantic bridge, which links concepts from different ontology perspectives on the same phenomenon while preserving their distinct meanings.

Current benchmarks provide limited support for evaluating such semantic bridges. As a result, it remains difficult to evaluate whether ontology alignment systems, including LLM-based methods, can identify meaningful bridges across heterogeneous perspectives.

This paper presents an ongoing effort to develop a gold-standard dataset for cross-perspective ontology alignment in the food domain. The primary contribution is an expert-curated reference dataset intended as a controlled testbed for evaluating whether LLM-based ontology alignment systems can identify cross-perspective semantic bridges. The dataset is designed to support the evaluation of semantic matching beyond equivalence and to study how different ontology perspectives can be related while preserving their conceptual distinctions. This makes the dataset useful for studying alignments where related concepts can be linked without being forced into a single meaning.

Posters, Demos, Blue Sky, and Tutorials at SEMANTiCS 2026, Sep 2026, Ghent, Belgium

✉ divyasha.sunil.naik@uni-jena.de (D. S. Naik); birgitta.koenig-ries@uni-jena.de (B. König-Ries)

id 0009-0003-9374-6951 (D. S. Naik); 0000-0002-2382-9722 (B. König-Ries)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: Workflow for Constructing a Cross-Perspective Food Ontology Gold-Standard Benchmark

The remainder of the paper reviews related ontology alignment benchmarks, presents the proposed benchmark construction workflow, and concludes with discussion and outlook.

2. Related Work

Since 2004, OAEI has provided datasets and reference alignments for evaluating ontology matching systems [3]. Existing benchmark tracks mainly focus on equivalence and subsumption relations between ontologies with overlapping domains and related modeling objectives.

The OAEI Food Nutritional Composition track [4] provides a food-domain benchmark whose reference alignments cover equivalence and subclass mappings between food concepts.

The OAEI Complex track goes beyond one-to-one mappings by evaluating correspondences that involve logical constructors or transformation functions. However, its main focus is on complex formal correspondences rather than on semantic bridges across different conceptual perspectives.

The OAEI BeyondEquivalence track evaluates ontology matching systems across five mutually disjoint relation types: equivalence, superclass, subclass, overlap, and disjointness. This track recognises that real-world ontologies may differ in structure, granularity, and classification perspective, and is designed to evaluate typed class correspondences.

These tracks are close to our work, but they do not provide a gold-standard dataset for ontologies that describe the same phenomenon from different modeling perspectives. Such a dataset is needed to evaluate whether alignment systems can correctly identify correspondences along with their semantic relations. In our case, the aim is not only to assign a relation type between two classes, but also to record a bridge property label that explains the semantic connection between the concepts. Our work addresses this gap by building an expert-guided food-domain dataset across nutritional, quality, and food-service perspectives.

3. Benchmark Construction Workflow

This section presents the workflow for constructing a cross-perspective ontology alignment benchmark. The workflow focuses on creating an expert-guided gold-standard dataset in the food domain. The proposed methodology is organized into six key stages. Figure 1 illustrates the overall workflow used to build the benchmark data set. The workflow starts with ontology selection and subset extraction, followed by metadata profiling, meta-tag assignment, relation assignment, and finalization of the gold-standard dataset.

Table 1

Selected ontology subsets and their primary perspectives.

Ontology subset	Perspective	Example concepts
ONS	Food material and diet	vegan diet, plant food product
OccO	Food-service occupations	chef or head cook, waiter, bartender
MeSH	Food quality and public health	food contamination, food labeling

Table 2

Statistics of the food-related ontology subsets.

Ontology subset	Classes	Obj. prop.	Data prop.	Definitions	Restrictions
ONS	26	4	0	26	12
OccO	52	12	1	45	71
MeSH	13	0	0	13	0

3.1. Searching and Selecting Ontologies

Food was selected as the initial benchmark domain because it represents a common real-world phenomenon that can be modelled from multiple perspectives without requiring highly specialized domain expertise. Food can be conceptualized as a dietary practice, a public-health concern, a quality-controlled product, or a food-service activity. This makes the domain suitable for studying cross-perspective ontology alignment.

To construct the benchmark, food-related ontologies were searched on BioPortal¹. Three ontologies representing complementary perspectives of the food domain were selected: **ONS** (Ontology for Nutritional Studies) [5], **MeSH** (Medical Subject Headings) [6], and **OCCO** (Ontology of Culinary Occupations) [7]. Table 1 summarizes the perspectives represented by the selected ontologies.

3.2. Ontology Subset Extraction

Small ontology subsets were extracted from the selected existing ontologies using ROBOT². Seed terms were manually selected to obtain meaningful subsets related to food perspectives. The BOT extraction method was applied with intermediates and individuals to preserve local semantic structure while reducing unrelated ontology content. Table 2 summarizes the metadata of the three ontology subsets. The subsets differ in size, annotation coverage, and logical expressivity.

3.3. Ontology Profiling

Each ontology subset was profiled by extracting class labels, definitions, hierarchy, properties, and restrictions. This information was used to characterize the food-domain perspective represented by each ontology.

3.4. Meta-Tag Assignment

To support systematic comparison across perspectives, concepts may optionally be grouped using neutral meta-tags that act as an intermediate conceptual layer independent of the original ontology structures. These tags support candidate generation between semantically compatible concept groups across perspectives. This step is not necessary for the MeSH subset because it is very small, but it is useful for the ONS and OccO subsets. Examples of meta-tags include *Dietary Pattern*, *Food preparation or serving role*.

¹<https://bioportal.bioontology.org/>

²<http://robot.obolibrary.org/>

Table 3

Summary of semantic bridge groups across ontology pairs.

Ontology pair	Strong	Weak	Few Strong semantic bridges
MeSH ↔ OccO	6	3	Food Industry ↔ Food preparation or serving role; Food Quality ↔ Food preparation or serving role
MeSH ↔ ONS	7	3	Nutritive Value ↔ Dietary Patterns; Food Quality ↔ Food material; Food Labeling ↔ food material
ONS ↔ OccO	4	1	Edible and Processed Food ↔ Food preparation role; Management and food preparation role; food material ↔ Food preparation role; Management and food preparation role

Table 4

Examples of concept-level semantic bridges across ontology pairs. Bridge labels are directed from Concept A to Concept B.

Ontology Pair	Concept A	Concept B	Bridge Label	Reason/Evidence
OccO–MeSH	Food preparation or serving related occupation holder	Food Industry	works_in	Food preparation or serving-related occupation holders work in the food industry.
MeSH–ONS	Food Contamination	Food product	affects	Food contamination affects food products.
ONS–OccO	Food product	Chef or head cook	prepared_by	Food products are prepared by chefs or head cooks.

3.5. Candidate Pair Selection and Relation Assignment

Candidate concept pairs were generated using compatible meta-tag groups when available, thereby reducing semantically unrelated comparisons. When meta-tags were not assigned, as in the MeSH subset, concepts were considered directly against relevant concepts from the other ontology subsets. The final reference alignment therefore contains only directed concept-to-concept correspondences. The selection process focused on concepts describing related phenomena in the food domain from different perspectives.

1. **Relation assignment:** Annotators assigned one of the following relation types: *equivalent*, *broadier*, *narrower*, *related*, or *no_relation*. The relation types were defined as follows:
 - *Equivalent*: two concepts refer to the same or nearly the same concept.
 - *Broadier*: the source concept is more general than the target concept.
 - *Narrower*: the source concept is more specific than the target concept.
 - *Related*: a semantic bridge is possible between two concepts.
 - *No relation*: no meaningful semantic relation was identified between the concepts.
2. **Assign Property label with short reason:** Annotators assigned a property label describing how the concepts are related across perspectives.

Short textual justifications were collected to support the expert decisions and improve the interpretability of the resulting benchmark. Existing property labels were reused whenever possible. The preferred sources were Relation Ontology, upper ontologies such as BFO and DOLCE, or three selected ontologies. When no suitable property label was available, a custom property label was assigned.

3.6. Gold Dataset Finalization

The expert annotations were consolidated in three steps. First, candidate pairs proposed by the five experts were grouped and compared across all annotation records. The minimum 2-expert threshold is set to select candidate pairs across ontologies. For each repeated pair, the selected relation labels were analyzed to identify agreement and disagreement along with their short reason.

Second, the bridge property labels and short justifications were reviewed. Cases with the same candidate pair, relation label, and bridge property label were treated as strong agreement. Cases with different bridge property labels or justifications were manually reviewed to decide whether they reflected disagreement, different granularity, or complementary interpretations.

Third, the direction of each bridge property relation was checked. If experts annotated a bridge property relation in opposite directions, both directions were reviewed. When both directions were semantically valid, both property labels were retained. When one direction implied an inverse relation, the inverse property was recorded.

The final bridge property labels were selected based on expert agreement in a second review iteration. Table 3 summarizes the resulting strong and weak semantic bridge groups across the three ontology pairs. Strong bridges correspond to annotations supported by three to five experts, while weak bridges correspond to annotations supported by two experts. Table 4 provides examples of finalized bridge relations, including the selected bridge property labels and expert reason/evidence. These examples illustrate how the reference dataset captures not only whether two concepts are related, but also how the semantic connection is justified across ontology perspectives. The current version of the dataset, *CrossAligner-Food*³, is publicly available.

4. Discussion and Outlook

The proposed benchmark captures semantic bridges across nutritional, quality, and food-service perspectives. It preserves the conceptual distinctions between these perspectives. Bridge property labels and short expert justifications provide additional context for each correspondence. All finalized correspondences are classified as *related* and further described using directed bridge property labels. Labels such as *prepared_by* and *processed_by* were treated as custom properties unless a verified ontology IRI was available. Semantic bridges were identified across all three ontology pairs.

The benchmark is intentionally designed as a compact, expert-curated evaluation resource rather than as a method for manually aligning complete ontologies. Its limited scope supports detailed expert review and transparent annotation analysis. However, it also limits the scalability of the current workflow. Expert validation remains necessary because cross-perspective relations cannot always be inferred from equivalence or class hierarchies.

For larger extensions, candidate generation could be separated from gold-standard validation. One possible approach is to represent the selected ontology subsets as a combined graph and treat potential cross-ontology relations as missing links. Link-prediction methods could then rank potentially relevant concept pairs. This provides a more scalable way to identify promising candidate pairs before expert validation. This could reduce the annotation effort without replacing expert judgment. Future work will explore this approach, include more ontology perspectives, and evaluate ontology alignment systems using the benchmark.

5. Acknowledgments

This work is supported by the German Federal Ministry of Education and Research through ThWIC⁴. We thank Balaramakrishna Paritala, Suresh Kumar Bonagiri, Vishvapalsinhji Ramsinh Parmar, Sven Thiel, Pascal Scherreiks, Cornelia Fürstenau, and Tarek Al Mustafa for their contributions to the dataset creation.

6. Declaration on Generative AI

The author(s) have used Generative AI tools for sentence enhancement.

³<https://github.com/fusion-jena/CrossAligner-Food>

⁴<https://www.thwic.uni-jena.de>

References

- [1] T. R. Gruber, A translation approach to portable ontology specifications, *Knowledge acquisition* 5 (1993) 199–220.
- [2] N. Guarino, D. Oberle, S. Staab, What is an ontology?, in: *Handbook on ontologies*, Springer, 2009, pp. 1–17.
- [3] M. A. N. Pour, et al., Results of the ontology alignment evaluation initiative 2025, in: *OM 2025*, 2026.
- [4] P. Buche, et al., How to manage incompleteness of nutritional food sources?: A solution using FoodOn as pivot ontology, *International Journal of Agricultural and Environmental Information Systems* 12 (2021) 1–26. doi:10.4018/IJAEIS.20211001.0a4.
- [5] F. Vitali, et al., Ons: an ontology for a standardized description of interventions and observational studies in nutrition, *Genes & Nutrition* 13 (2018) 12.
- [6] C. E. Lipscomb, Medical subject headings (mesh), *Bulletin of the Medical Library Association* 88 (2000) 265.
- [7] S. Smith, D. M. Dooley, E. Merrell, J. Beverley, Y. He, Toward an occupation ontology, *occo.*, in: *ICBO*, 2022, p. 1.