

# OAEI-LLM: A Benchmark Dataset for Understanding Large Language Model Hallucinations in Ontology Matching

Zhangcheng Qiang<sup>1,\*</sup>, Kerry Taylor<sup>1</sup>, Weiqing Wang<sup>2</sup> and Jing Jiang<sup>1</sup>

<sup>1</sup>Australian National University, School of Computing, 108 North Road, Acton, ACT 2601, Canberra, Australia

<sup>2</sup>Monash University, Faculty of Information Technology, 25 Exhibition Walk, Clayton, VIC 3800, Melbourne, Australia

## Abstract

Hallucinations of large language models (LLMs) commonly occur in domain-specific downstream tasks, with no exception in ontology matching (OM). The prevalence of using LLMs for OM raises the need for benchmarks to better understand LLM hallucinations. The OAEI-LLM dataset is an extended version of the Ontology Alignment Evaluation Initiative (OAEI) datasets that evaluate LLM-specific hallucinations in OM tasks. We outline the methodology used in dataset construction and schema extension, and provide examples of potential use cases.

## Keywords

ontology matching, large language models, LLM hallucinations

## 1. Motivation

Large language models (LLMs) have shown incredible capabilities in natural language generation (NLG) and question answering (QA). In the Semantic Web community, LLMs have recently been used in ontology matching (OM). Although LLMs offer a strong background knowledge base for OM, they may miss some true mappings and generate a number of false mappings. This phenomenon has been observed in several recent papers in [1, 2, 3, 4, 5, 6]. It is mainly caused by LLM hallucinations, where LLMs tend to generate synthesised answers when they do not have sufficient background knowledge or have biased domain knowledge [7].

The conference track is one of the Ontology Alignment Evaluation Initiative (OAEI) [8] campaigns that evaluates OM tasks in conference organisation. It includes several reference alignments between different ontologies related to research conference management [9, 10, 11, 12]. For example, “<http://cmt#Chairman>” and “<http://conference#Chair>” form a pair of valid match in the CMT-Conference alignment. LLMs can experience hallucinations when they mistakenly categorise a pair of matched entities to be unmatched entities. This often happens because it is very difficult for LLMs to understand polysemous words, such as where the word “chair” refers to the person who chairs the conference, and not to the furniture used at the conference. LLM hallucinations can also occur when one of the matched entities is mismatched with other entities. For example, LLMs may consider “chair” in the context of a research conference to have a similar meaning to “conference chair”. Therefore, “<http://conference#Chair>” can be mismatched to “<http://cmt#ConferenceChair>”. Both the former case of missing true mappings and the latter case of generating false mappings will have a negative impact on the matching performance. It seems feasible to mitigate LLM hallucinations using fine-tuning, but there are few datasets that can be used to provide baseline references.

Established in 2004, OAEI has become the largest community of specified ontology matching challenges. Over the years, OAEI has developed several multi-ontology tracks across a wide range of topic domains. We observe that no tracks have covered the topic of LLM hallucinations in OM tasks. For this reason, we propose to extend the current OAEI datasets to construct a new dataset called

---

ISWC 2024 Special Session on Harmonising Generative AI and Semantic Web Technologies, November 13, 2024, Baltimore, Maryland

\*Corresponding author.

✉ qzc438@gmail.com (Z. Qiang); kerry.taylor@anu.edu.au (K. Taylor); teresa.wang@monash.edu (W. Wang);

jing.jiang@anu.edu.au (J. Jiang)

🆔 0000-0001-5977-6506 (Z. Qiang); 0000-0003-2447-1088 (K. Taylor); 0000-0002-9578-819X (W. Wang); 0000-0002-3035-0074 (J. Jiang)



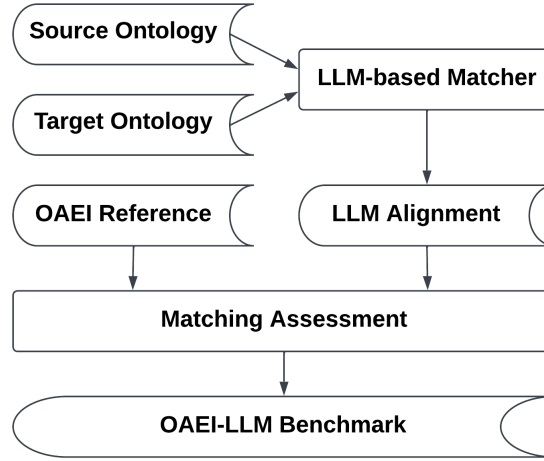
© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

OAEI-LLM, which serves to measure the degree of hallucination by LLMs in OM tasks. The new dataset compares the original human-labelled results with LLM-generated results, classifies different types of hallucinations made by different LLMs, and records the information with a new schema extension. We believe this dataset would benefit research in understanding LLM hallucinations in matching tasks and thereby research in improving LLM-driven OM.

## 2. Methodology

### 2.1. Dataset Construction

Figure 1 illustrates the procedure for constructing the dataset. The original OAEI datasets provide three files: the source ontology ( $O_s$ ), the target ontology ( $O_t$ ), and their OAEI Reference ( $R_{oaei}$ ). We implement an LLM-based OM system to generate the alignment file LLM Alignment ( $A_{llm}$ ). The system takes  $O_s$  and  $O_t$  as inputs and generates a set of predicted mappings. We constrain our mappings to the case of only one-to-one mappings, that is, every entity in the source ontology is mapped to at most one entity in the target, and each entity in the target ontology occurs in at most one mapping. The matching assessment procedure detailed below will compare  $R_{oaei}$  with  $A_{llm}$  and find their differences.



**Figure 1:** Procedure for constructing the dataset: an LLM Alignment of the source and target ontologies is compared with the OAEI Reference and recorded in the OAEI-LLM Benchmark.

For each mapping  $(e1, e2) \in R_{oaei}$ , we consider that it is not an LLM hallucination if there exists  $(e1', e2') \in A_{llm}$  such that  $e1 = e1'$  and  $e2 = e2'$ . We categorise LLM mapping errors in three ways, as shown in Table 1.

**Table 1**  
Different types of LLM mapping errors.

Type	Definition
<b>Missing from LLM</b>	For a specific pair $(e1, e2) \in R_{oaei}$ , there is no pair $(e1', e2') \in A_{llm}$ that satisfies $(e1 = e1')$ exclusive-or $(e2 = e2')$ .
<b>Missing from OAEI</b>	For a specific pair $(e1', e2') \in A_{llm}$ , there is no pair $(e1, e2) \in R_{oaei}$ that satisfies $(e1 = e1')$ exclusive-or $(e2 = e2')$ .
<b>Incorrect</b>	For a specific pair $(e1, e2) \in R_{oaei}$ , there exists a pair $(e1', e2') \in A_{llm}$ that either satisfies $e1 = e1'$ and $e2 \neq e2'$ or satisfies $e1 \neq e1'$ and $e2 = e2'$ .

The matching assessment process identifies the alignment from which a missing mapping is missing. It also determines a categorisation of incorrect mappings as follows:

- False-mapping: LLMs map the entity to an irrelevant entity.
- Disputed-mapping: LLMs map the entity to a relevant entity, but it does not align with the OAEI Reference. We call this mapping type “disputed” because we cannot guarantee the OAEI Reference is always correct. LLMs may discover a more precise matching entity.
- Align-up: LLMs map the entity to its intended superclass.
- Align-down: LLMs map the entity into its intended subclass.

In practice, the notions of *relevant* and *intended* are implemented by the chosen matching assessment process. We use an LLM-based evaluator for matching assessment, and directly query the LLM for relevance and superclass and subclass relationships. At this time we cannot recommend a particular LLM to be used.

## 2.2. Schema Extension

The OAEI tracks use a general alignment format derived from the Expressive and Declarative Ontology Alignment Language (EDOAL) [13]. An example to represent a valid match of “<http://cmt#Chairman>” and “<http://conference#Chair>” is illustrated as follows: Each mapping is recorded in a tag `<Cell>`. The schema uses the tags `<entity1>` and `<entity2>` to record the pair of entities, the tag `<measure>` to present their similarity, and the tag `<relation>` to indicate their relations.

We extend the current EDOAL mapping schema to record the new information related to LLM hallucinations. The Simple Standard for Sharing Ontological Mappings (SSSOM) [14] could be an alternative base to extend for our purpose. Code Snippet 1 illustrates an example of an OAEI-LLM reference alignment using our proposed extended EDOAL mapping schema. For each mapping, different types of hallucinations can occur in different LLMs. Each of them is recorded in the tag `<hallucination>`. In the extended version of the example (shown in blue, green, and purple), the first segment depicts the LLM hallucinations that occurred in claude-3-sonnet [15], which is a case of a mapping missing from the LLM. The second segment illustrates the LLM hallucinations that occurred in llama-3-8b [16], indicating an incorrect mapping to the entity “<http://cmt#ConferenceChair>”. The matching assessment classifies this mapping as a disputed mapping.

Code Snippet 1: An example of an OAEI-LLM reference alignment using our proposed extended EDOAL mapping schema. Text generated by AI tool Agent-OM [4] (an LLM-based ontology matching agent) is shown in green. Text generated by AI tool LLM-based evaluator is shown in purple.

---

```

<map>
  <Cell>
    <entity1 rdf:resource="http://cmt#Chairman"/>
    <entity2 rdf:resource="http://conference#Chair"/>
    <measure rdf:datatype="xsd:float">1.0</measure>
    <relation>=</relation>
    <hallucination>
      <llm>claude-3-sonnet</llm>
      <category>missing</category>
      <source>llm</source>
    </hallucination>
    <hallucination>
      <llm>llama-3-8b</llm>
      <category>incorrect</category>
      <entity1 rdf:resource="http://cmt#ConferenceChair"/>
      <type>disputed-mapping</type>
    </hallucination>
  </Cell>
</map>

```

---

## 3. Potential Use Cases

### 3.1. Benchmarking LLMs for OM Tasks

Different LLMs can perform differently on the same OM task. However, running an LLM-based matcher on a traditional OAEI dataset can only observe the differences in precision, recall, and F1 score. Using the extended version of OAEI-LLM, we can additionally quantify the LLM errors made by counting each of the different types of errors we have identified. Therefore, it becomes possible to understand the tendency of LLMs to generate incorrect answers or to provide relevant but not precise answers. This could also be a valuable supplement to enhance an LLM leaderboard for OM tasks. Indeed, this same additional information on matching failures could also be helpful if adopted more widely by non-LLM ontology matchers.

### 3.2. A Dataset for Fine-tuning LLMs Used in OM Tasks

Fine-tuning is a common approach used in LLMs to inject domain knowledge, but a significant amount of data typically needs to be provided. Unlike the traditional OAEI datasets, which only provide information on matched entities, OAEI-LLM provides extensive information identifying when and how LLM hallucinations occur, and even the mismatched entities caused by LLM hallucinations. This could be used to generate high-quality training data to fine-tune LLMs.

## 4. Limitations

In this paper, we only demonstrate the simplest one-to-one equivalent mappings. More generally, the reference alignment can be incomplete or contain one-to-many mappings, while the matching type can also go beyond equivalence to include subsumption matching and complex matching. Such sophisticated cases may cause additional types of LLM hallucinations. Some LLM hallucinations could be mitigated by providing more ontology structure as context to the LLM, for example, by employing the locality principle of traditional matchers [17]. Nevertheless, the proposed benchmark dataset will be useful until LLM-based matchers achieve perfection.

In our work to date we use a moderately simple LLM-based evaluator to perform the matching assessment. The assessment may vary depending on the LLM used. A more sophisticated matching assessment could use a voting process by multiple LLMs or human expert judgement to generate a more reliable OAEI-LLM benchmark dataset.

## 5. Further Work

The OAEI tracks can be divided into three categories: schema matching (i.e. TBox matching), instance matching (i.e. ABox matching), and interactive matching (i.e. matching that involves user interaction). Correspondingly, OAEI-LLM will also have three variants.

- OAEI-LLM-T will focus on comprehending LLM hallucinations in the TBox, with a specific interest in the formal use of terminologies in ontology design and implementation.
- OAEI-LLM-A will focus on addressing LLM hallucinations in the ABox, in which an ontology contains a schema together with many instance entities (similar to ontology-based knowledge graphs).
- OAEI-LLM-I will focus on understanding how user interaction could be used to mitigate LLM hallucinations and estimating the level of user interaction required for LLM-based OM tasks.

**Acknowledgements** The authors thank the reviewers for providing helpful comments. The authors thank the Commonwealth Scientific and Industrial Research Organisation (CSIRO) for supporting this project. According to the OAEI data policy (retrieved December 1, 2024), “OAEI results and datasets, are publicly available, but subject to a use policy similar to the one defined by NIST for TREC. These rules apply to anyone using these data.” Please find more details from the official website: <https://oaei.ontologymatching.org/doc/oaei-deontology.2.html>.

## References

- [1] Y. He, J. Chen, H. Dong, I. Horrocks, Exploring Large Language Models for Ontology Alignment, in: Proceedings of the ISWC 2023 Posters, Demos and Industry Tracks, volume 3632, CEUR-WS.org, Athens, Greece, 2023.
- [2] S. S. Norouzi, M. S. Mahdavinejad, P. Hitzler, Conversational Ontology Alignment with ChatGPT, in: Proceedings of the 18th International Workshop on Ontology Matching – ISWC 2023, volume 3591, CEUR-WS.org, Athens, Greece, 2023, pp. 61–66.
- [3] S. Hertling, H. Paulheim, OLaLa: Ontology Matching with Large Language Models, in: Proceedings of the 12th Knowledge Capture Conference 2023, ACM, Pensacola, FL, USA, 2023, pp. 131–139. doi:10.1145/3587259.3627571.
- [4] Z. Qiang, W. Wang, K. Taylor, Agent-OM: Leveraging LLM Agents for Ontology Matching, 2023. URL: <https://arxiv.org/abs/2312.00326>. arXiv:2312.00326.
- [5] H. B. Giglou, J. D’Souza, F. Engel, S. Auer, LLMs4OM: Matching Ontologies with Large Language Models, 2024. URL: <https://arxiv.org/abs/2404.10317>. arXiv:2404.10317.
- [6] R. Amini, S. S. Norouzi, P. Hitzler, R. Amini, Towards Complex Ontology Alignment using Large Language Models, 2024. URL: <https://arxiv.org/abs/2404.10329>. arXiv:2404.10329.
- [7] Y. Zhang, Y. Li, L. Cui, D. Cai, L. Liu, T. Fu, X. Huang, E. Zhao, Y. Zhang, Y. Chen, L. Wang, A. T. Luu, W. Bi, F. Shi, S. Shi, Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models, 2023. URL: <https://arxiv.org/abs/2309.01219>. arXiv:2309.01219.
- [8] OAEI Community, Ontology Alignment Evaluation Initiative (OAEI), 2024. URL: <https://oaei.ontologymatching.org>.
- [9] M. Cheatham, P. Hitzler, Conference v2.0: An Uncertain Version of the OAEI Conference Benchmark, in: The Semantic Web – ISWC 2014, Springer, Riva del Garda, Italy, 2014, pp. 33–48. doi:10.1007/978-3-319-11915-1\_3.
- [10] A. Solimando, E. Jiménez-Ruiz, G. Guerrini, Detecting and Correcting Conservativity Principle Violations in Ontology-to-Ontology Mappings, in: The Semantic Web – ISWC 2014, Springer, Riva del Garda, Italy, 2014, pp. 1–16. doi:10.1007/978-3-319-11915-1\_1.
- [11] A. Solimando, E. Jimenez-Ruiz, G. Guerrini, Minimizing Conservativity Violations in Ontology Alignments: Algorithms and Evaluation, Knowledge and Information Systems 51 (2017) 775–819. doi:10.1007/s10115-016-0983-3.
- [12] O. Zamazal, V. Svátek, The Ten-Year OntoFarm and its Fertilization within the Onto-Sphere, Journal of Web Semantics 43 (2017) 46–53. doi:10.1016/j.websem.2017.01.001.
- [13] J. David, J. Euzenat, F. Scharffe, C. Trojahn dos Santos, The Alignment API 4.0, Semantic Web 2 (2011) 3–10. doi:10.3233/SW-2011-0028.
- [14] N. Matentzoglou, J. P. Balhoff, S. M. Bello, C. Bizon, M. Brush, T. J. Callahan, C. G. Chute, W. D. Duncan, C. T. Evelo, D. Gabriel, J. Graybeal, A. Gray, B. M. Gyori, M. Haendel, H. Harmse, N. L. Harris, I. Harrow, H. B. Hegde, A. L. Hoyt, C. T. Hoyt, D. Jiao, E. Jiménez-Ruiz, S. Jupp, H. Kim, S. Koehler, T. Liener, Q. Long, J. Malone, J. A. McLaughlin, J. A. McMurry, S. Moxon, M. C. Munoz-Torres, D. Osumi-Sutherland, J. A. Overton, B. Peters, T. Putman, N. Queralt-Rosinach, K. Shefchek, H. Solbrig, A. Thessen, T. Tudorache, N. Vasilevsky, A. H. Wagner, C. J. Mungall, A Simple Standard for Sharing Ontological Mappings (SSSOM), Database 2022 (2022) baac035. doi:10.1093/database/baac035.
- [15] Anthropic, claude-3-sonnet, 2024. URL: <https://www.anthropic.com/news/claude-3-family>.
- [16] Meta, llama-3-8b, 2024. URL: <https://ai.meta.com/blog/meta-llama-3/>.
- [17] E. Jiménez-Ruiz, B. C. Grau, I. Horrocks, R. Berlanga, Logic-based Assessment of the Compatibility of UMLS Ontology Sources, Journal of Biomedical Semantics 2 (2011). doi:10.1186/2041-1480-2-S1-S2.