WojoodOntology: Ontology-Driven LLM Prompting for Unified Information Extraction Tasks

Alaa Aljabari [λ*] **Nagham Hamad** [λ*] **Mohammed Khalilia**[λ] **Mustafa Jarrar**[σ,λ]
[λ] Birzeit University, Palestine
[σ] Hamad Bin Khalifa University, Qatar
{aaljabari, nhamad,mkhalilia,mjarrar}@birzeit.edu

Abstract

Information Extraction tasks such as Named Entity Recognition and Relation Extraction are often developed using diverse tagsets and annotation guidelines. This presents major challenges for model generalization, cross-dataset evaluation, tool interoperability, and broader industry adoption. To address these issues, we propose an information extraction ontology, *WojoodOntology*, which covers a wide range of named entity types and relations. *WojoodOntology* serves as a semantic mediation framework that facilitates alignment across heterogeneous tagsets and annotation guidelines. We propose two ontology-based mapping methods: (i) as a set of mapping rules for uni-directional tagset alignment; and (ii) as ontology-based prompting, which incorporates the ontology concepts directly into prompts, enabling large language models (LLMs) to perform more effective and bi-directional mappings. Our experiments show a 15% improvement in out-of-domain mapping accuracy when using ontology-based prompting compared to rule-based methods. Furthermore, *WojoodOntology* is aligned with Schema.org and Wikidata, enabling interoperability with knowledge graphs and facilitating broader industry adoption. The *WojoodOntology* is open source and available at https://sina.birzeit.edu/wojood.

1 Introduction

Information extraction tasks—such as Named Entity Recognition (NER) and Relation Extraction (RE)—are essential for extracting structured data from text. These tasks play a critical role in applications like information retrieval (Marinov et al., 2024), word sense disambiguation (Jarrar et al., 2023b; Al-Hajj and Jarrar, 2021), data extraction (Barbon Junior et al., 2024), language understanding (Khalilia et al., 2024), interoperability (Jarrar et al., 2011), among others.

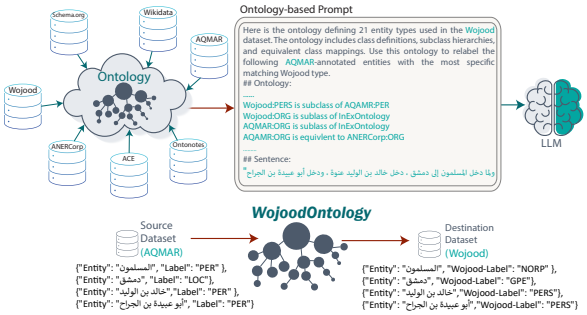

Figure 1: Ontology-guided prompting for mapping between datasets using LLMs. The model maps sentences and entity annotations from a source dataset to a destination dataset based on the defined in the ontology.

Although many NER and RE datasets have been developed, they cannot be combined due to differing annotation guidelines and schemas (Yang et al., 2025). This heterogeneity presents significant challenges. For instance, in Wojood NER dataset (Jarrar et al., 2022), الخليج العربي /Arabian Gulf is labeled as LOC and مدينة دمشق /Damascus City as GPE, whereas both are tagged as LOC in ANER dataset (Benajiba et al., 2007a). In addition, different boundary span definitions across datasets pose significant challenges. For instance, according to Wojood's guidelines, مدينة دمشق /Damascus City is annotated as a GPE, whereas in ANER, only دمشق /Damascus is tagged as GPE, and مدينة/City is labeled as O. Similarly, الملك عبدالله /King Abdallah is tagged as PERS in Wojood, but only عبدالله /Abdallah span is considered PERS in ANER and Ontonotes (Weischedel et al., 2017). In relation extraction, inconsistencies also emerge. For example, in Wikidata, the *hasConflictWith* relationship is defined between PERS and EVENT entities, whereas in *Wojood*^{Relations} often annotate it either between two PERS entities or between two ORG entities (Aljabari et al., 2025).

Furthermore, such inconsistencies prevent NLP tool interoperability. For instance, SinaTools and CaMLTools are incompatible, as each uses differ-

* Equal contribution.

Proceedings of The Third Arabic Natural Language Processing Conference, pages 179–193
November 8-9, 2025

ent tagsets and annotation guidelines. SinaTools supports 21 entity types and 40 relation types (Aljabari et al., 2024, 2025), while CaMLTools supports only 4 entity types (Obeid et al., 2020). Thus, the bidirectional mappings between these different tagsets are infeasible due to schema mismatches and annotation differences (See Section 3).

Schema.org provides shared data schemas widely used by industry and search engines for products, jobs, events, people, organizations, reviews, and more. Similarly, Wikidata covers most real-world entities and relationships in a multilingual knowledge graph. Yet, these standards are rarely considered in NER and RE tagset design, limiting real-world use. Aligning tagsets with standards like Wikidata and Schema.org would improve interoperability and ensure extracted data is immediately useful for industry applications.

Despite advances in Large Language Models (LLMs), they often misclassify entities due to ambiguity or unfamiliar schema labels (Potu et al., 2025). Studies have shown that LLMs may assign arbitrary labels, resulting in inconsistent outputs that are difficult to integrate (Feng et al., 2024).

To overcome these issues, we introduce *WojoodOntology*, a novel information extraction ontology that encompasses a wide range of named entity types and their relationships, including concepts and relations. The ontology defines 55 concepts (named entity types) and 40 relationships, including subclass and equivalent class relations. In addition, it is aligned with Schema.org and Wikidata, enabling interoperability with knowledge graphs and facilitating broader industry adoption. *WojoodOntology* serves multiple purposes. First, it provides a formal specification of concepts and relations (i.e., well-structured annotation guidelines). Second, it facilitates the alignment of heterogeneous tagsets and guidelines. We present two implementations of the ontology: (1) A Python library that provides uni-directional mapping rules for tagset alignment. (2) An ontology-based prompting method that integrates the ontology directly into LLM prompts, enabling effective bi-directional tagset mappings. As shown in Figure 1, these implementations allow users to re-annotate corpora labeled with one tagset (e.g., Wojood, OntoNotes, Wikidata) into another. We evaluated this prompting method by re-annotating the AQMAR corpus with Wojood guidelines. We achieved a 15% performance improvement compared with the rule-based mapping method.

The key contributions of this work are:

- *WojoodOntology*, a novel information extraction ontology.
- Python library for uni-directional mapping between IE tagsets.
- Novel ontology-based prompting method enabling LLMs to perform efficient bi-directional tagset mappings.

The paper is organized as follows: Section 2 reviews related work; Section 3 presents *WojoodOntology*; Section 5 presents the experiments; and we conclude in Section 7.

2 Related Work

2.1 NER and RE Datasets

Several Arabic NER corpora have been introduced with varying annotation schemes. *Wojood* (Jarrar et al., 2022) is a large-scale corpus of about 550k tokens annotated with 21 entity types, and its guidelines have become the basis for subsequent resources. *Wojood*$_{fine}$ expands *Wojood* with 30 fine-grained sub-entity types, yielding 51 categories in total (Liqreina et al., 2023; Jarrar et al., 2023a). *Wojood*^{Gaza}, a 60k-token corpus focusing on news about the Israeli War on Gaza and Nakba NLP, applies the same guidelines across 51 entity types and subtypes (Jarrar et al., 2024, 2025). *Konooz* is another large corpus encompassing $777K$ tokens across 10 domains and 16 dialects (Hamad et al., 2025). It is annotated with both flat and nested entities following the *Wojood* tagset. Other existing NER corpora focus on MSA, such as ANERCorp (Benajiba et al., 2007b), OntoNotes (Weischedel et al., 2017), and AQMAR (Mohit et al., 2012a).

Although several dialectal corpora with diverse types of linguistic annotations have been developed (Jarrar et al., 2023c; Nayouf et al., 2023), none include NER annotation, with the exception of the Palestinian and Lebanese *Curras+Baladi* corpora. Both corpora are part of the Wojood corpus (Haff et al., 2022; Jarrar et al., 2017). Beyond NER, they are also annotated with morphological tags and lemmatization, and further mapped to Qabas (Jarrar and Hammouda, 2024) and the Arabic Ontology (Jarrar, 2021).

For RE, existing Arabic relation extraction corpora include ACE05 (Doddington et al., 2004), a multilingual dataset covering English, Chinese, and

Arabic with 6 relations and 5 entity types. SMiLAR (Seganti et al., 2021), a multilingual joint entity and relation corpus with $9K$ Arabic sentences and 36 relation types. SRED^{FM} and RED^{FM} (Huguet Cabot et al., 2023), multilingual resources with automatic and human-verified annotations, including Arabic portions. Wojood^{Hadath} (Aljabari et al., 2024), an Arabic-specific event-argument extraction dataset with 3 relations and 21 entity types using a nested NER scheme. Last but not least, *Wojood*^{Relations} is the largest Arabic RE corpus, comprising $33K$ sentences annotated with 40 relation types and 21 entity types under a nested NER scheme (Aljabari et al., 2025).

2.2 Mapping

Recent studies show that fine-tuning LLMs on large-scale NER datasets improves their performance. However, direct training on existing datasets is hindered by the heterogeneity of entity and relation definitions, limiting the model's ability to generalize to unseen domains. To address the problem, ontology mapping has been explored using both manual and automatic approaches. Rizzo and Troncy (2012) proposed the NERD ontology as a common interface for entity annotation across different schemas. It consists of manually defined mappings between various named entity schemas, such as DBpedia Spotlight and OpenCalais. However, this manual approach lacks scalability when dealing with a wide range of entity types or adapting to new schemas. Nozza et al. (2021) introduced an automatic mapping approach by leveraging embedding representations of named entities to align taxonomies across domains, showing improvements over manual methods with an 86% F1 score. However, the method relies on BERT embeddings, which are less effective for entity representation.

The Open NER framework (Yang et al., 2025) has focused on improving entity recognition in English and Chinese by unifying entity definitions across datasets, demonstrating substantial improvements in NER performance. However, this approach lacks scalability for new entity types. It is mainly performed by holding out certain datasets from existing ones. Another approach proposes detailed annotation guidelines for entity and relation labeling (Sainz et al., 2024), but such guidelines are difficult to enforce consistently and challenging for models to interpret.

Fine-tuning NER models on multiple datasets, enabling LLMs to learn diverse entity definitions and enhance generalization (Gui et al., 2024; Sainz et al., 2024). However, this approach does not extend to RE, where inconsistent relation labels across datasets continue to hinder cross-domain performance. In addition, the absence of a unified taxonomy for both entities and relations remains a significant obstacle, preventing models from learning semantically consistent representations. Currently, no ontology is specifically designed for Arabic NER and RE datasets, nor one that effectively integrates external resources like Wikidata and Schema.org to support model generalization.

3 The *WojoodOntology*

WojoodOntology serves as a unified framework for mapping entity and relation types across diverse Named Entity Recognition (NER) and Relation Extraction (RE) datasets. It is constructed through a comprehensive review of existing Arabic information extraction datasets, spanning both named entity recognition and relation extraction. To ensure broad coverage, we include all entity and relation types identified in the literature. Furthermore, we integrate related concepts and hierarchical structures from external knowledge bases, such as Wikidata and Schema.org, to enhance semantic alignment and interoperability. The resulting ontology consists of 55 entity types (Figure 2) and 40 relation types (Figure **??**, Appendix §4), with sample relations shown in Figure 3.

To enable automated reasoning, consistency checking, and integration with external knowledge resources, we formalize the ontology using OWL, standard Web Ontology Language. The formalization captures both the structural and semantic properties of entity and relation types, as detailed in the following subsection.

3.1 Formalizing Ontology for NER and RE

WojoodOntology is a hierarchy of entity types and relationships between them. Entity types (e.g., `ORG`, `LOC`) are OWL classes , while relation types are defined as object properties connecting pairs of classes (e.g., `Located_In (ORG, LOC)`). The ontology is a formalization of these components using standard OWL axioms, including equivalentClass, subClassOf, and domain-range constraints.

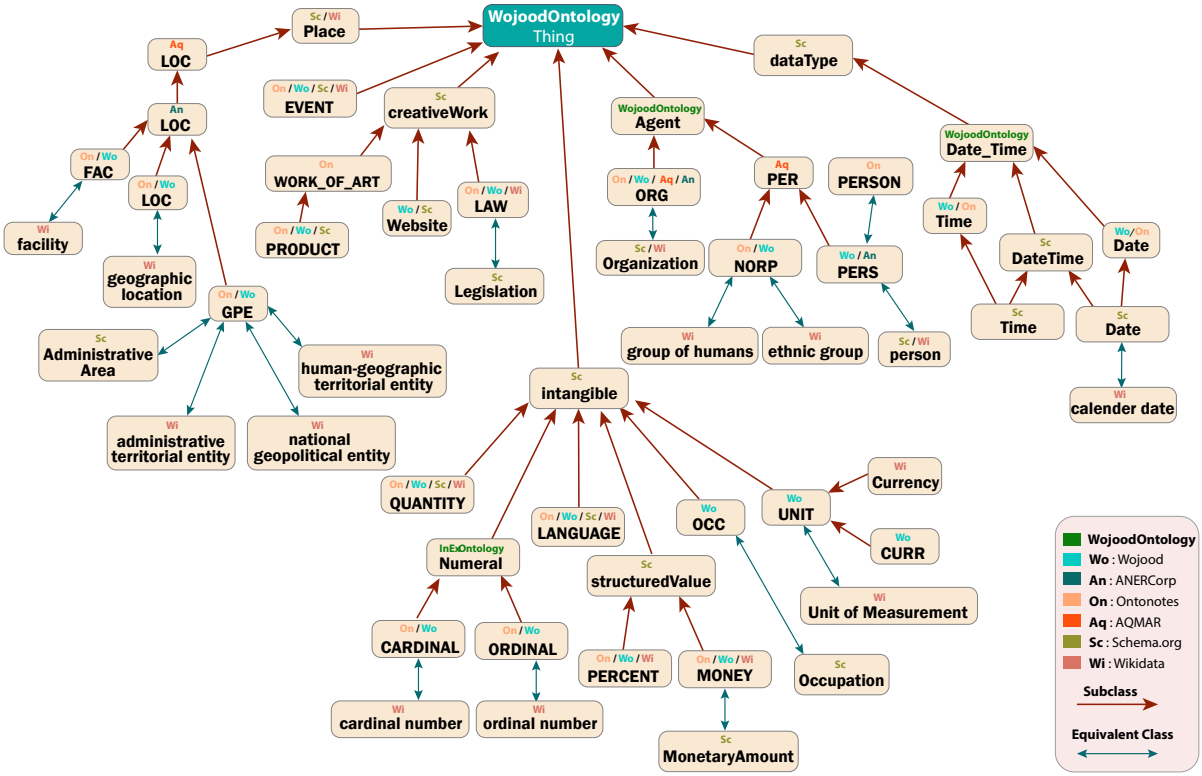

Figure 2: *WojoodOntology* (Class Hierarchy)

NER Formalization: The *equivalentClass* axiom is used to define semantic equivalence between entity types originating from different datasets or ontologies. Specifically, if an entity type C_i is declared equivalent to another type C_j, then any named entity assigned to C_i is also considered an instance of C_j, and vice versa. Formally, let:

$$\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$$

be the set of entity types in the ontology, where each C_i represents a class (e.g., `ORG`, `LOC`, `PERS`). Then the equivalence is defined as:

$$\text{equivalentClass}(C_i, C_j) \Leftrightarrow C_i = C_j$$

This axiom enables semantic interoperability by allowing entity types with consistent meaning and annotation boundaries to be treated interchangeably across datasets. In Wojood and OntoNotes, places are categorized into three types: `GPE`, `LOC`, and `FAC`, whereas ANERCorp and AQMAR use a single broad category, `LOC`. For example, the entity القدس/Jerusalem is labeled as `GPE` in Wojood and OntoNotes, whereas in AQMAR and ANERCorp it is labeled as `LOC`. Therefore, the `GPE` types in OntoNotes and Wojood can be treated as *equivalent classes*, whereas the `LOC` type in ANERCorp and AQMAR is not equivalent to `GPE` in Wojood.

The *subClassOf* axiom is used to define hierarchical relations between entity types. Specifically, if an entity type C_i is a subclass of another type C_j, then every named entity assigned to C_i is also implicitly assigned to C_j, but not vice versa. Formally, the subclass relation is defined as:

$$\text{subClassOf}(C_i, C_j) \Rightarrow C_i \subseteq C_j$$

This formalization enables mapping between entity types with different granularity or format constraints. For instance, `wojood:DATE` supports temporal instances expressed in natural language (e.g., ١٠ سبتمبر ٢٠١٨, عام ٢٠١٨) including standardized representations like the ISO 8601 formats. However, `schema:Date` is limited to ISO 8601. Therefore, we defined `schema:Date` as a subclass of `wojood:DATE`. This enables precise and consistent integration across datasets.

Figure 2 illustrates the class hierarchy, where arrows denote subclass relations (e.g., `ORG` → `Agent`), and bidirectional links indicate class equivalence (e.g., `NORP` ↔ `Ethnic Group`). This structure ensures coherent label integration across NER datasets, which are critical for supporting semantic interoperability and cross-dataset generalization.

Relation Formalization: In OWL, *object proper-*

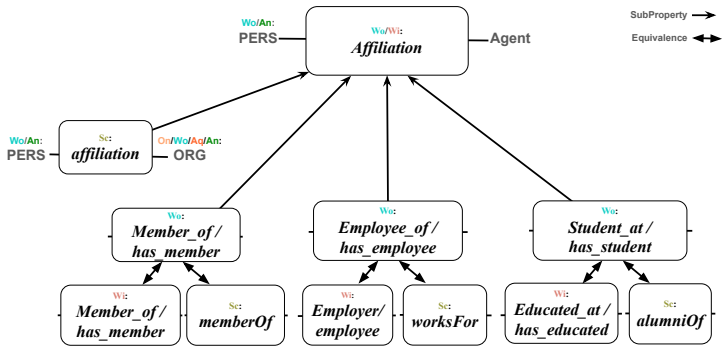

Figure 3: Example of *relationship hierarchy* - See the full hierarchy of relations in Appendix A2.

ties are relations between classes. Each relation type is an object property linking a subject class (domain) to an object class (range). Let the set of relation types be: $\mathcal{R} = \{R_1, R_2, \ldots, R_k\}$. Each relation $R_l \in \mathcal{R}$ is formally defined with domain and range constraints; $R_l : (C_a, C_b) \Rightarrow R_l \subseteq C_a \times C_b$, indicating that R_l holds between instances of class C_a (subject) and class C_b (object). For example, the relation `Located_In` is defined as $R_{\text{Located\_In}} : (\texttt{ORG}, \texttt{GPE})$, allowing assertions such as $(i_{\text{Google}}, i_{\text{USA}}) \in R_{\text{Located\_In}}$.

Relations in *WojoodOntology* are structured hierarchically using *subproperty* and *equivalence* axioms to enable consistent reasoning and cross-ontology mapping. A subproperty axiom defines a relation as a specialization of another, inheriting its semantics while providing more specificity:

$$\text{SubPropertyOf}(R_1, R_2) \Rightarrow \forall x, y \, (x \, R_1 \, y \Rightarrow x \, R_2 \, y)$$

In Figure 3, the (`Wo:employee_of` $\sqsubseteq$ `Wo:affiliation`) means that employment is a specific type of organizational affiliation. Equivalence axioms assert semantic identity between relations, potentially across ontologies:

$$\text{EquivalentObjectProperties}(R_1, R_2) \Rightarrow \forall x, y \, (x \, R_1 \, y \Leftrightarrow x \, R_2 \, y)$$

In Figure 3, (`Wo:employee_of` $\equiv$ `Sc:worksFor`) states that `employee_of` in $Wojood^{Relations}$ is equivalent to `worksFor` in Schema.org. These equivalences are essential for ensuring interoperability across heterogeneous datasets and external knowledge graphs.

Overall, these axioms (i) enforce inheritance of domain-range constraints and (ii) support unified reasoning over heterogeneous resources.

Trained Model	Inference Dataset	F1 Score	
		Macro	Micro
Wojood	ANERCorp	10%	44%
	OntoNotes	**33%**	**58%**
	AQMAR	8%	41%
ANERCorp	Wojood	8%	48%
	OntoNotes	9%	50%
	AQMAR	**25%**	**60%**
OntoNotes	Wojood	**22%**	**55%**
	ANERCorp	11%	52%
	AQMAR	9%	44%
AQMAR	Wojood	8%	48%
	ANERCorp	**29%**	**72%**
	OntoNotes	8%	48%

Table 1: Cross-dataset NER evaluations: each model is trained on one dataset and tested on others.

3.2 *WojoodOntology* Construction

WojoodOntology is constructed in multiple steps:
Step 1: Cross-dataset Validation of Entity Types. To examine the annotation differences across NER datasets, we conducted cross-dataset validation experiments using four datasets: Wojood (Jarrar et al., 2022), ANERCorp (Benajiba et al., 2007a), AQMAR (Mohit et al., 2012b), and OntoNotes (Weischedel et al., 2017). BERT-based models were trained on each dataset and evaluated on the others to examine the consistency of entity definitions and annotation guidelines. As shown in Table 1, all models experienced substantial performance degradation when tested on unseen datasets, highlighting the impact of annotation divergence. However, higher cross-dataset scores were observed between ANERCorp and AQMAR, as well as between OntoNotes and Wojood. This is attributed to the shared tagsets and similar annotation practices within each pair, suggesting that annotation alignment plays a key role in cross-domain generalization.

For example, Figure 4a highlights major inconsistencies for the LOC category, with F1 scores dropping significantly across datasets. This stems

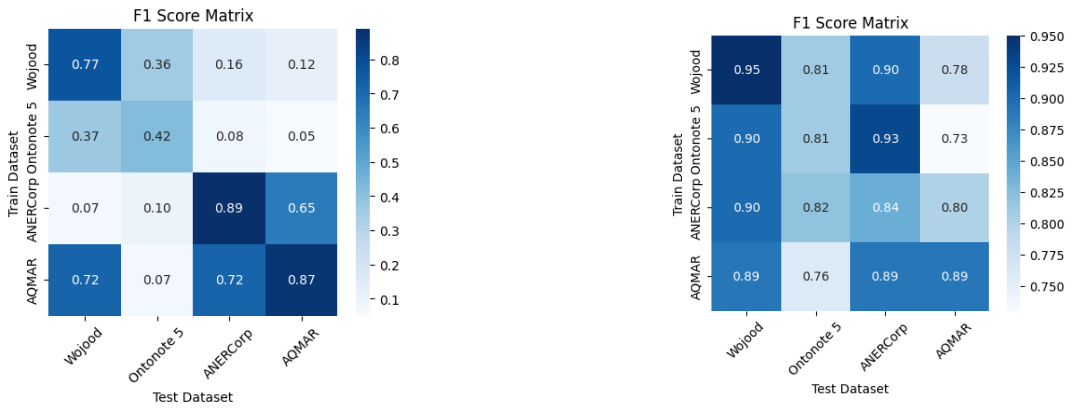

(a) LOC Heat Map

(b) PERS Heat Map

Figure 4: Heatmaps of Cross-dataset Predictions for LOC and PERS Entities

from schema mismatches, where some datasets distinguish between geopolitical entities and physical landmarks, while others merge them. In contrast, Figure 4b shows strong alignment for the PERS entity type between Wojood and other datasets, but weaker alignment between AQMAR and others. This discrepancy arises because AQMAR merges PERS and NORP into a single category, whereas other datasets maintain a finer-grained distinction, resulting in label mismatches. Consequently, Wojood's model underperforms on AQMAR (F1 = 0.78), while the AQMAR model performs better on Wojood (F1 = 0.89), reflecting Wojood's more detailed entity taxonomy. Furthermore, OntoNotes exhibits notable annotation inconsistencies, which further complicate cross-dataset generalization.

Step 2: Comparative Analysis of Entity Definitions and Annotations. To further investigate the causes of cross-dataset variability, we performed a comparative analysis of entity definitions and annotation schemes. In addition, we integrated external knowledge sources, including `Schema.org` and `Wikidata`, to provide broader semantic coverage. We systematically examined each entity type across all datasets and knowledge graphs to identify variations in annotation scope, label naming conventions, and granularity. The identified discrepancies in definitions and annotation guidelines between entity types across the NER datasets and knowledge graphs are summarized in Table 5.

Step 3: Ontology Construction and Schema Mapping Based on the comparative analysis, we identify equivalent and subclass relationships among entity types to construct a unified ontology. This step captures the hierarchical structure and semantic alignment between labels. For instance, PERS in Wojood and ANERCorp, and PERSON in OntoNotes, are identified as equivalent classes, all of which are modeled as subclasses of the broader PER class in AQMAR. The ontology supports reverse mapping by leveraging subclass relations to align each entity mention with its most specific fine-grained type. The class hierarchy of the ontology is presented in Figure 2.

Step 4: Relation Identification and Alignment. We identify relation types that connect the named entities defined in the constructed ontology and align them with external schemas such as `Schema.org` and `Wikidata`. This alignment follows the domain and range constraints formalized in Section 3.1, ensuring semantic consistency across sources.

To construct the relations ontology and establish the hierarchy among relations, we first compare the formal definitions of each relation across RE datasets and knowledge graphs. Two relations are considered equivalent when their definitions are semantically identical and their domain and range specifications are equivalent classes.

In contrast, a relation is defined as a sub-relation of another relation if two conditions are satisfied. First, semantic inclusion must hold, meaning that all instances of the first relation are also valid instances of the second relation, but not vice versa. Second, the domain and range of the first relation must be either equivalent to, or subclasses of, the domain and range of the second relation. When both conditions are met, a hierarchical dependency between the two relations is established, with the first relation formally designated as a sub-relation of the second. For example, *headquartered_in* is a sub-relation of *located_in*. The former specifies the

location of an organization's central office, while the latter denotes the place of any agent. Every instance of *headquartered_in* implies an instance of *located_in*, but not all instances of *located_in* (e.g., a branch or individual located in a place) satisfy the stricter definition of *headquartered_in*. Moreover, since the domain of *headquartered_in* (`Organization`) is a subclass of the domain of *located_in* (`Agent`) and both share the same range (`Place`), *headquartered_in* is formally identified as a sub-relation of *located_in*.

The resulting relation schema is presented in Appendix 4, and a representative snapshot is shown in Figure 3. Mapping details are summarized in Table 6 for Wikidata and Table 7 for Schema.org.

4 Mapping between Datasets

Mapping between datasets is challenging due to differences in annotation guidelines, as well as label granularity and definitions. Mapping can be categorized as unidirectional or bidirectional. Unidirectional mapping projects datasets with finer-grained entity types (e.g., Wojood) onto coarser-grained ones (e.g., ANERCorp). However, bidirectional mapping enables mutual alignment. Automatic bidirectional mapping is challenging and remains largely underexplored due to inconsistencies in annotation guidelines.

We introduce *WojoodOntology* as a novel solution for cross-dataset interoperability, supporting both unidirectional and bidirectional mapping.

4.1 Uni-directional Ontology-based Mapping

For *uni-directional* mapping, we use the *WojoodOntology* to derive mapping rules. These rules are derived from the *equivalentClass* and *subClassOf* semantic relationships defined in the ontology. When two entity types are linked via an equivalency relation, they are mapped directly to each other, such as the `ORG` entity type in Wojood and AQMAR.

When an entity type in one dataset is defined as a subclass of a broader type in another dataset (a subClassOf relation), the mapping rule assigns the more specific type to its parent type. For instance, as shown in Figure 2, the `FAC`, `LOC`, and `GPE` types in Wojood are all defined as subclasses of AQMAR's broader `LOC` type. Accordingly, all mentions tagged as `FAC`, `LOC`, or `GPE` in Wojood are re-labeled as `LOC` to align with AQMAR's annotation schema.

4.2 Ontology-Driven Prompting for Bi-directional Mapping

To enable *bi-directional mapping*, we propose an ontology-guided prompting approach using LLMs to translate between different datasets, leveraging the *WojoodOntology* as a semantic reference.

We propose using LLM prompting to re-annotate datasets originally labeled with one tagset into a target tagset. The ontology is embedded in the prompt to provide contextual guidance, ensuring consistent interpretation of tags and enabling accurate translation across annotation schemes. In this approach, the ontology serves as an external semantic reference, helping the LLM disambiguate and align tag definitions across datasets. For example, the *WojoodOntology* guides the LLM to re-label the broader `LOC` category in AQMAR into the more specific types `GPE`, `FAC`, or `LOC` in Wojood. As discussed in the next section, we experimented with four prompts (Figures 5 and 6) and their results are summarized in Table 4.

5 Experiments and Results

WojoodOntology provides a framework for mapping entities across heterogeneous NER and RE datasets. To evaluate its effectiveness, we use the mapping between Wojood and AQMAR datasets as a case study. Wojood supports 21 tags, while AQMAR is only 4, with differences in tag labels and annotation guidelines. We evaluate unidirectional and bidirectional mapping using the ontology.

In our experiments, we used the GPT-4o engine with carefully controlled hyperparameters. The temperature was set to 0.0 to ensure deterministic outputs, while the maximum token length was limited to $4,096$. We set Top_p to 1.

5.1 Uni-directional Ontology-based Mapping

To demonstrate the effectiveness of the mapping rules discussed in Section 4 (also summarized in Table 5), we apply these rules to map the entity types from Wojood to the corresponding AQMAR labels: the `PERS` and `NORP` labels in Wojood are considered `PER` in AQMAR; the `LOC` in Wojood is mapped to `LOC` in AQMAR; the `GPE` and `FAC` in Wojood are mapped to `LOC` in AQMAR; the `ORG` is considered `ORG` in AQMAR; and, all other labels in Wojood are considered to `O`.

In Table 2, we illustrate the impact of our mapping rules. First, we train a model on Wojood and evaluate it directly on AQMAR without applying

any mapping rule. This model achieves only an 8% F1 score. However, when the unidirectional mapping rules are used, performance increases to 40%. To verify that the low performance is due to domain shift rather than discrepancies in the mapping rules, we conducted an additional experiment. We trained a model on Wojood combined with 10% of AQMAR. This setup achieves a 52% F1 score on the remaining 90% of AQMAR, indicating that the performance degradation is due to domain shift rather than inconsistencies in the mapping rules.

Experimental Setting	F1	Improv.
Baseline (No Mapping)		
Wojood → AQMAR	8%	–
Ontology-Based Mapping		
Wojood (mapped to AQMAR)	40%	+32%
Wojood + 10% AQMAR (fine-tuned)	52%	+44%

Table 2: Experiments on ontology-based unidirectional mapping rules (Wojood → AQMAR).

5.2 Ontology-Driven Prompting for Bi-directional Mapping

To conduct bi-directional mapping experiments, we first re-annotated the AQMAR corpus manually following the Wojood guidelines. We call the new version of AQMAR as AQMAR^{W}. Table 3 presents the entity distribution of this version.

Second, we used AQMAR^{W} to evaluate LLMs' performance under two experimental setups: zero-shot and few-shot prompting, and with and without the *WojoodOntology*.

Tag	Count	Tag	Count
PERS	1,148	NORP	747
OCC	342	ORG	907
GPE	697	LOC	242
FAC	391	PRODUCT	317
EVENT	352	DATE	799
TIME	58	LANGUAGE	20
WEBSITE	7	LAW	4
CARDINAL	670	ORDINAL	440
PERCENT	29	QUANTITY	101
UNIT	20	MONEY	27
CURR	1	-	-
Total	7,319 **entity mentions**		

Table 3: AQMAR^{W} Dataset Statistics

Zero Shot Prompting: In the zero-shot setting, we conducted two experiments (Figure 5), both incorporating the *WojoodOntology* into the prompt to guide re-annotation of AQMAR entities. In the first experiment, the original AQMAR labels were provided, enabling the model to re-annotate them (`LOC`, `ORG`, `PER`, `MISC`) according to the Wojood tagset. However, it failed to capture entity types present in Wojood but absent in AQMAR (e.g., `GPE`, `PRODUCT`, `CURR`). In the second experiment, the ontology was used without AQMAR labels, yielding slightly better performance.

Overall, as shown in Table 4, both experiments demonstrate that incorporating the ontology substantially improves model performance compared to the baseline that did not use the ontology (29% vs. 8% F1-score). **Few-Shot Prompting:**

We further evaluated the effectiveness of *WojoodOntology* in a few-shot setting through two experiments (Figure 6). In the first experiment, we did not embed the ontology in the prompt, but we added seven demonstration examples. These examples were selected from AQMAR^{W} based on entity types that LLMs often misannotate (e.g., TIME, DATE, EVENT, CARDINAL, ORDINAL). This improved performance relative to the zero-shot setting, achieving 49% F1 compared to 29%. In the second experiment, we incorporated the ontology into the prompt alongside the same seven examples, which further improved performance to 55% F1 (Table 4).

Overall, the zero-shot and few-shot results—with and without the ontology—underscore that embedding the ontology as an external semantic reference substantially enhances model performance in AQMAR re-annotation.

Setting	Precision	Recall	F1-score
Zero-shot			
Ontology (w/ ent.)	0.3194	0.2388	0.2733
Ontology (w/o ent.)	0.3319	0.2595	0.2913
Few-shot			
Without Ontology	0.5109	0.4879	0.4991
With Ontology	**0.5730**	**0.5294**	**0.5504**

Table 4: Ontology-based prompting performance in zero-shot and few-shot bi-directional entity mapping.

6 Discussion

The result emphasizes the challenge posed by inconsistent annotation guidelines across NER datasets. LLMs struggle to infer fine-grained mappings between schemes when no ontology is given. In zero-shot settings, using the ontology improves performance slightly when entities are not explicitly provided, indicating that structural knowledge from the ontology offers better guidance than entity mention cues alone. However, the overall F1

(A) Ontology-based Prompt (With Provided AQMAR Dataset)

Map this sentence and its entities from AQAMR to Wojood using the given ontology. Infer from the OWL all possible entities in the sentence that are not annotated in AQMAR, but considered as entities in Wojood Only use entity type tags that exist in the Wojood dataset. Do not include any dataset prefix (e.g., return ORG instead of wojood#ORG). Your answer should be in JSON format as a list of dictionaries with this structure: [Entity Span: ENTITY_SPAN, Entity Type: ENTITY_TYPE]
Ontology: *[Ontology in OWL]* **Sentence:***[sentence]* **Entities in AQMAR:** *[AQMAR entities]*

(B) Ontology-based Prompt (without provided AQMAR dataset entities)

Map this sentence and its entities from AQAMR to Wojood using the given ontology. Infer from the OWL all possible entities in the sentence that are not annotated in AQMAR, but considered as entities in Wojood Only use entity type tags that exist in the Wojood dataset. Do not include any dataset prefix (e.g., return ORG instead of wojood#ORG). Your answer should be in JSON format as a list of dictionaries with this structure: [Entity Span: ENTITY_SPAN, Entity Type: ENTITY_TYPE]
Ontology: *[Ontology in OWL]* **Sentence:***[sentence]*

Figure 5: Zero-shot LLM prompts using ontology-guided named entity mapping

Few-shot without ontology-based prompting

Here is the 21 entity types used in the Wojood dataset. The tagsets are [PERS, ORG, NORP, LOC, OCC, DATE, TIME, EVENT, CARDINAL, ORDINAL, CURR, LAW, WEBSITE, GPE, FAC, PRODUCT, LANGUAGE, QUANTITY, PERCENT, UNIT]. Please use labels to relabel the following AQMAR-annotated entities with the most specific matching Wojood type. Ignore the AQMAR entity type — base your decision only on the span and sentence context. If you cannot confidently assign a type, return "None".
Sentence:*[sentence]* **Examples:***[7 Examples]*

Few-shot without ontology-based prompting

Here is the 21 entity types used in the Wojood dataset. The tagsets are [PERS, ORG, NORP, LOC, OCC, DATE, TIME, EVENT, CARDINAL, ORDINAL, CURR, LAW, WEBSITE, GPE, FAC, PRODUCT, LANGUAGE, QUANTITY, PERCENT, UNIT]. Please use labels to relabel the following AQMAR-annotated entities with the most specific matching Wojood type. Ignore the AQMAR entity type — base your decision only on the span and sentence context. If you cannot confidently assign a type, return "None".
Ontology: *[Ontology in OWL]* **Sentence:***[sentence]* **Examples:***[7 Examples]*

Figure 6: Few-shot LLM Prompt with (and without) ontology

score remains low in both zero-shot variants, reflecting the difficulty of schema mapping without demonstrations, with F1 below 0.30.

In contrast, few-shot prompting substantially improves performance, reaching an F1-score of 50%. Incorporating a small set of annotated demonstrations, particularly those containing challenging entities, allows the model to generalize more effectively. Importantly, the inclusion of ontology information alongside these demonstrations produces the highest performance, achieving an F1-score of 55%. This highlights the critical role of ontological knowledge in guiding the model. By providing structured semantic axioms, the ontology enhances few-shot learning and enables the LLM to perform more accurate cross-schema entity alignment.

7 Conclusion

The *WojoodOntology* provides a formal semantic framework that facilitates interoperability across heterogeneous datasets. Our results indicate that even straightforward, rule-based mappings, when guided by the ontology, improve model performance. Evaluation of zero-shot and few-shot experiments further demonstrates that ontology-guided prompting yields consistent improvements in model performance. These findings highlight the potential of ontology-driven methods for developing unified information extraction systems across diverse annotated resources.

8 Limitation

One limitation of this work is that the MISC tag in both ANERcorp and AQMAR datasets is not included in the ontology due to inconsistencies in its definition across the two resources. In ANERcorp, MISC includes entities that do not fall under standard types like `PER`, `LOC`, or `ORG`, while in AQMAR it often overlaps with other categories or lacks a clear scope. This discrepancy makes alignment challenging and may affect overall coverage. Additionally, all experiments were conducted using GPT-4o. While it shows strong performance, evaluating multiple LLMs would provide a more comprehensive understanding of model behavior and generalization across different architectures.

References

Moustafa Al-Hajj and Mustafa Jarrar. 2021. ArabGloss-BERT: Fine-Tuning BERT on Context-Gloss Pairs for WSD. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 40–48, Online. INCOMA Ltd.

Alaa Aljabari, Lina Duaibes, Mustafa Jarrar, and Mohammed Khalilia. 2024. Event-Arguments Extraction Corpus and Modeling using BERT for Arabic. In *Proceedings of the Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok, Thailand. Association for Computational Linguistics.

Alaa Aljabari, Mohammed Khalilia, and Mustafa Jarrar. 2025. Wojood^{Relations}: Arabic Relation Extraction Corpus and Modeling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Suzhou, China. Association for Computational Linguistics.

Sylvio Barbon Junior, Paolo Ceravolo, Sven Groppe, Mustafa Jarrar, Samira Maghool, Florence Sèdes, Soror Sahri, and Maurice Van Keulen. 2024. Are Large Language Models the New Interface for Data Pipelines? In *Proceedings of the International Workshop on Big Data in Emergent Distributed Environments*, BiDEDE '24, New York, NY, USA. Association for Computing Machinery.

Yassine Benajiba, Paolo Rosso, and José Miguel BenedíRuiz. 2007a. Anersys: An arabic named entity recognition system based on maximum entropy. In *Computational Linguistics and Intelligent Text Processing*, pages 143–153, Berlin, Heidelberg. Springer Berlin Heidelberg.

Yassine Benajiba, Paolo Rosso, and José Miguel Benedíruiz. 2007b. Anersys: An arabic named entity recognition system based on maximum entropy. In *Computational Linguistics and Intelligent Text Processing: 8th International Conference, CICLing 2007, Mexico City, Mexico, February 18-24, 2007. Proceedings 8*, pages 143–153. Springer.

George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. The automatic content extraction (ACE) program – tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).

Xiaohan Feng, Xixin Wu, and Helen Meng. 2024. Ontology-grounded automatic knowledge graph construction by LLM under wikidata schema. In *Proceedings of the KDD Workshop on Human-Interpretable AI*, volume 3841 of *CEUR Workshop Proceedings*, pages 117–135. CEUR-WS.org.

Honghao Gui, Lin Yuan, Hongbin Ye, Ningyu Zhang, Mengshu Sun, Lei Liang, and Huajun Chen. 2024. IEPile: Unearthing large scale schema-conditioned information extraction corpus. In *Proceedings of the 62nd Annual Meeting of ACL (Volume 2: Short Papers)*, pages 127–146, Bangkok, Thailand. Association for Computational Linguistics.

Karim El Haff, Mustafa Jarrar, Tymaa Hammouda, and Fadi Zaraket. 2022. Curras + Baladi: Towards a Levantine Corpus. In *Proceedings of the International Conference on Language Resources and Evaluation(LREC 2022)*, Marseille, France.

Nagham Hamad, Mohammed Khalilia, and Mustafa Jarrar. 2025. Konooz: Multi-domain Multi-dialect Corpus for Named Entity Recognition. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, pages 0–0, Vienna, Austria. Association for Computational Linguistics.

Pere-Lluís Huguet Cabot, Simone Tedeschi, Axel-Cyrille Ngonga Ngomo, and Roberto Navigli. 2023. RED^{fm}: a filtered and multilingual relation extraction dataset. In *Proceedings of the 61st Annual Meeting of ACL (Volume 1: Long Papers)*, pages 4326–4343, Toronto, Canada. ACL.

Mustafa Jarrar. 2021. The Arabic Ontology - An Arabic Wordnet with Ontologically Clean Content. *Applied Ontology Journal*, 16(1):1–26.

Mustafa Jarrar, Muhammad Abdul-Mageed, Mohammed Khalilia, Bashar Talafha, AbdelRahim Elmadany, Nagham Hamad, and Alaa' Omar. 2023a. WojoodNER 2023: The First Arabic Named Entity Recognition Shared Task. In *Proceedings of the 1st Arabic Natural Language Processing Conference (ArabicNLP), Part of the EMNLP 2023*, pages 748–758. ACL.

Mustafa Jarrar, Anton Deik, and Bilal Faraj. 2011. Ontology-based data and process governance framework -the case of e-government interoperability in palestine. In *Proceedings of the IFIP International Symposium on Data-Driven Process Discovery and Analysis (SIMPDA'11)*, pages 83–98.

Mustafa Jarrar, Nizar Habash, Faeq Alrimawi, Diyam Akra, and Nasser Zalmout. 2017. Curras: An annotated corpus for the palestinian arabic dialect. *Journal Language Resources and Evaluation*, 51(3):2-s2.0-85001544989.

Mustafa Jarrar, Nizar Habash, Mo El-Haj, Amal Haddad Haddad, Zeina Jallad, Camille Mansour, Diana Allan, Paul Rayson, Tymaa Hammouda, and Sanad Malaysha, editors. 2025. *Proceedings of the first International Workshop on Nakba Narratives as Language Resources*. Association for Computational Linguistics, Abu Dhabi, UAE.

Mustafa Jarrar, Nagham Hamad, Mohammed Khalilia, Bashar Talafha, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2024. WojoodNER 2024: The Second Arabic Named Entity Recognition Shared Task. In *Proceedings of the Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok, Thailand. Association for Computational Linguistics.

Mustafa Jarrar and Tymaa Hasanain Hammouda. 2024. Qabas: An Open-Source Arabic Lexicographic Database. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13363–13370, Torino, Italy. ELRA and ICCL.

Mustafa Jarrar, Mohammed Khalilia, and Sana Ghanem. 2022. Wojood: Nested Arabic Named Entity Corpus and Recognition using BERT. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2022)*, Marseille, France.

Mustafa Jarrar, Sanad Malaysha, Tymaa Hammouda, and Mohammed Khalilia. 2023b. SALMA: Arabic Sense-annotated Corpus and WSD Benchmarks. In *Proceedings of the 1st Arabic Natural Language Processing Conference (ArabicNLP), Part of the EMNLP 2023*, pages 359–369. ACL.

Mustafa Jarrar, Fadi Zaraket, Tymaa Hammouda, Daanish Masood Alavi, and Martin Waahlisch. 2023c. Lisan: Yemeni, Irqi, Libyan, and Sudanese Arabic Dialect Copora with Morphological Annotations. In *The 20th IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*, pages 1–7. IEEE.

Mohammed Khalilia, Sanad Malaysha, Reem Suwaileh, Mustafa Jarrar, Alaa Aljabari, Tamer Elsayed, and Imed Zitouni. 2024. ArabicNLU 2024: The First Arabic Natural Language Understanding Shared Task. In *Proceedings of the Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok, Thailand. Association for Computational Linguistics.

Haneen Liqreina, Mustafa Jarrar, Mohammed Khalilia, Ahmed Oumar El-Shangiti, and Muhammad Abdul-Mageed. 2023. Arabic Fine-Grained Entity Recognition. In *Proceedings of the 1st Arabic Natural Language Processing Conference (ArabicNLP), Part of the EMNLP 2023*, pages 310–323. ACL.

Martin Marinov, Youcef Benkhedda, Goran Nenadic, and Riza Batista-Navarro. 2024. Relation extraction for constructing knowledge graphs: Enhancing the searchability of community-generated digital content (CGDC) collections. In *Workshop on Deep Learning and Large Language Models for Knowledge Graphs*.

Behrang Mohit, Nathan Schneider, Rishav Bhowmick, Kemal Oflazer, and Noah A Smith. 2012a. Recall-oriented learning of named entities in arabic wikipedia. In *Proceedings of EACL 2012*, pages 162–173.

Behrang Mohit, Nathan Schneider, Rishav Bhowmick, Kemal Oflazer, and Noah A. Smith. 2012b. Recall-oriented learning of named entities in arabic wikipedia. In *Proceedings of ECAL 2012*, EACL '12, page 162–173, USA. ACL.

Amal Nayouf, Mustafa Jarrar, Fadi zaraket, Tymaa Hammouda, and Mohamad-Bassam Kurdy. 2023. Nâbra: Syrian Arabic Dialects with Morphological Annotations. In *Proceedings of the 1st Arabic Natural Language Processing Conference (ArabicNLP), Part of the EMNLP 2023*, pages 12–23. ACL.

Debora Nozza, Pikakshi Manchanda, Elisabetta Fersini, Matteo Palmonari, and Enza Messina. 2021. Learningtoadapt with word embeddings: Domain adaptation of named entity recognition systems. *Inf. Process. Manag.*, 58:102537.

Ossama Obeid, Nasser Zalmout, Salam Khalifa, Dima Taji, Mai Oudah, Bashar Alhafni, Go Inoue, Fadhl Eryani, Alexander Erdmann, and Nizar Habash. 2020. CAMeL tools: An open source python toolkit for Arabic natural language processing. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7022–7032, Marseille, France. European Language Resources Association.

Sai Teja Potu, Rachana Niranjan Murthy, Akhil Thomas, Lokesh Mishra, Natalie Prange, and Ali Riza Durmaz. 2025. Ontology-conformal recognition of materials entities using language models. *Scientific Reports*, 15(1):1–16.

Giuseppe Rizzo and Raphaël Troncy. 2012. NERD: A framework for unifying named entity recognition and disambiguation extraction tools. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of ACL*, pages 73–76, Avignon, France. ACL.

Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, and Eneko Agirre. 2024. GoLLIE: Annotation guidelines improve zero-shot information-extraction. In *The Twelfth International Conference on Learning Representations*.

Alessandro Seganti, Klaudia Firląg, Helena Skowronska, Michał Satława, and Piotr Andruszkiewicz. 2021. Multilingual entity and relation extraction dataset and model. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1946–1955, Online. Association for Computational Linguistics.

Ralph M. Weischedel, Eduard H. Hovy, Mitchell P. Marcus, and Martha Palmer. 2017. Ontonotes : A large training corpus for enhanced processing.

Yuming Yang, Wantong Zhao, Caishuang Huang, Junjie Ye, Xiao Wang, Huiyuan Zheng, Yang Nan, Yuran Wang, Xueying Xu, Kaixin Huang, Yunke Zhang, Tao Gui, Qi Zhang, and Xuanjing Huang. 2025. Beyond boundaries: Learning a universal entity taxonomy across datasets and languages for open named entity recognition. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10902–10923, Abu Dhabi, UAE. ACL.

A Comparative Analysis of Entity Definitions and Annotations.

To support the mapping process and analyze the source of cross-dataset inconsistencies, we conducted a comparative analysis of entity definitions and annotation schemes across Wojood, OntoNotes, ANERCorp, AQMAR, Schema and Wikidata. Table 5 summarizes the entity labels used in each dataset, their corresponding Wikidata classes, and notable annotation notes.

The analysis reveals significant differences in label granularity and category definitions. For instance, while Wojood distinguishes between FAC, LOC, and GPE, AQMAR merges these into a single LOC category. Such discrepancies are common across several entity types and directly affect interoperability between datasets.

B Constructing Relation Ontology

B.1 Aligning *Wojood*^{Relations} with Knowledge Graphs

To ensure interoperability between the *Wojood*^{Relations} schema and widely adopted knowledge bases such as Wikidata and Schema.org, we align relation types in *Wojood*^{Relations} with semantically equivalent or hierarchically related properties in these external ontologies. This alignment is based on formal relation definitions and constrained by domain and range specifications.

To capture the granularity and semantic compatibility of relation types across datasets and knowledge graphs, we conduct a comparative analysis of their definitions. Two relations are considered equivalent if they convey the same semantic meaning and their domain and range types are ontologically equivalent. A a relation is considered a subrelation if its semantics are subsumed by a broader relation and its domain and range are subclasses (or equivalents) of those of the broader relation. These equivalence and subsumption mappings are used to construct a hierarchical relation ontology.

For example, as shown in Table 6, the relation manager_of in *Wojood*^{Relations} is semantically aligned with the Wikidata property manager/director (P1037). In Wikidata, this property connects instances of Human (Q5) to Organization (Q43229), while in Wojood, manager_of links entities of type PERS to ORG. According to the entity ontology defined in *WojoodOntology*, PERS is equivalent to Human, and ORG is equivalent to Organization. Therefore, the two relations are considered semantically equivalent.

Similarly, Table 7 extends this alignment to Schema.org, listing for each *Wojood*^{Relations} property its corresponding Schema.org property and the associated URI. This facilitates interoperability with applications and tools that adopt Schema.org as their semantic backbone, ensuring that the relational semantics of *Wojood*^{Relations} are preserved when integrated into web-scale knowledge graphs.

B.2 Relations Ontology

Based on the hierarchical mappings between *Wojood*^{Relations}, Wikidata, and Schema.org, we construct a unified relation ontology that integrates equivalence and subsumption relations across the three schemas. Each *Wojood*^{Relations} property is positioned within this hierarchy according to its semantic correspondence, ensuring that narrower relations are subsumed under broader ones while maintaining consistent domain and range constraints. The resulting ontology captures the alignment at multiple levels of abstraction, which serves as a bridge for interoperability across RE datasets and knowledge graphs. The complete relation ontology is shown in Figure 7.

Description	Wojood	OntoNote	ANERCorp	AQMAR	Schema.org	Wikidata	Notes
Person	PERS	PERSON	PERS	PER	Person	Person (Q215627)	AQMAR: PERS category also includes NORP (Nationalities and Religious/Political Groups).
Group of people	NORP	NORP	O	PER	-	Ethnic group (Q41710)	OntoNote: Includes nationalities (e.g., أمريكي/American).
Occupation	OCC	O	O	O	Occupation	Occupation (Q12737077)	
Organization	ORG	ORG	ORG	ORG	Organization	Organization (Q43229)	Wojood: ORG spans may include GPE or LOC of an organization, whereas other datasets do not, i.e. in Wojood جمعية سيدات الأعمال في مصر, while in others جمعية سيدات الأعمال.
Geopolitical Entities	GPE	GPE	LOC	LOC	Administrative Area	Geopolitical entity (Q15642541), National geopolitical entity (Q116052725), administrative territorial entity (Q56061), administrative territorial entity (Q56061)	ANERCorp and AQMAR: GPE is considered part of LOC category.
Location	LOC	LOC	LOC	LOC	-	Geographic Location (Q2221906)	ANERCorp: GPE and LOC are treated as the same category. AQMAR: GPE, LOC, and FAC all fall under LOC.
Facility	FAC	FAC	LOC	LOC	-	Architectural structure (Q811979), Facility (Q13226383)	AQMAR: Facilities (FAC) are classified under LOC.
Product	PRODUCT	PRODUCT	O	O	Product	Product (Q2424752)	ANERCorp and AQMAR: PRODUCT is classified under MISC.
Event	EVENT	EVENT	O	O	Event	Event (Q1656682)	ANERCorp and AQMAR: EVENT is classified under MISC.
Date	DATE	DATE	O	O	DATE	Point in time (Q186081)	AQMAR: Reference dates (e.g., العصر العباسي) are categorized as MISC, whereas actual dates are annotated as DATE.
Time	TIME	TIME	O	O	Time	Time (Q11471)	
Language	LANGUAGE	LANGUAGE	O	O	Language	Language (Q34770)	
Law	LAW	LAW	O	O	Legislation	Law (Q7748)	
Cardinal	CARDINAL	CARDINAL	O	O	-	Cardinal number (Q163875)	
Ordinal	ORDINAL	ORDINAL	O	O	-	Ordinal number (Q191780)	
Percent	PERCENT	O	O	O	Structured Value	Percentage (Q11229)	
Quantity	QUANTITY	QUANTITY	O	O	Quantity	Quantity (Q309314)	
Unit	UNIT	O	O	O	-	Unit of measurement (Q47574)	OntoNote: Currency (CURR) is part of QUANTITY (e.g., ٥٠ م), and no standalone units occur without a value (e.g., م alone).
Money	MONEY	MONEY	O	O	Monetary Amount	Money(Q1368)	
Currency	CURR	O	O	O	-	Currency (Q8142)	OntoNote: Currency (CURR) is considered part of MONEY (e.g., ١٥٠ دولار), and no standalone currencies occur without a value (e.g., دولار alone).

Table 5: Entity Granularity Across Different NER Datasets and Knowledge Graphs

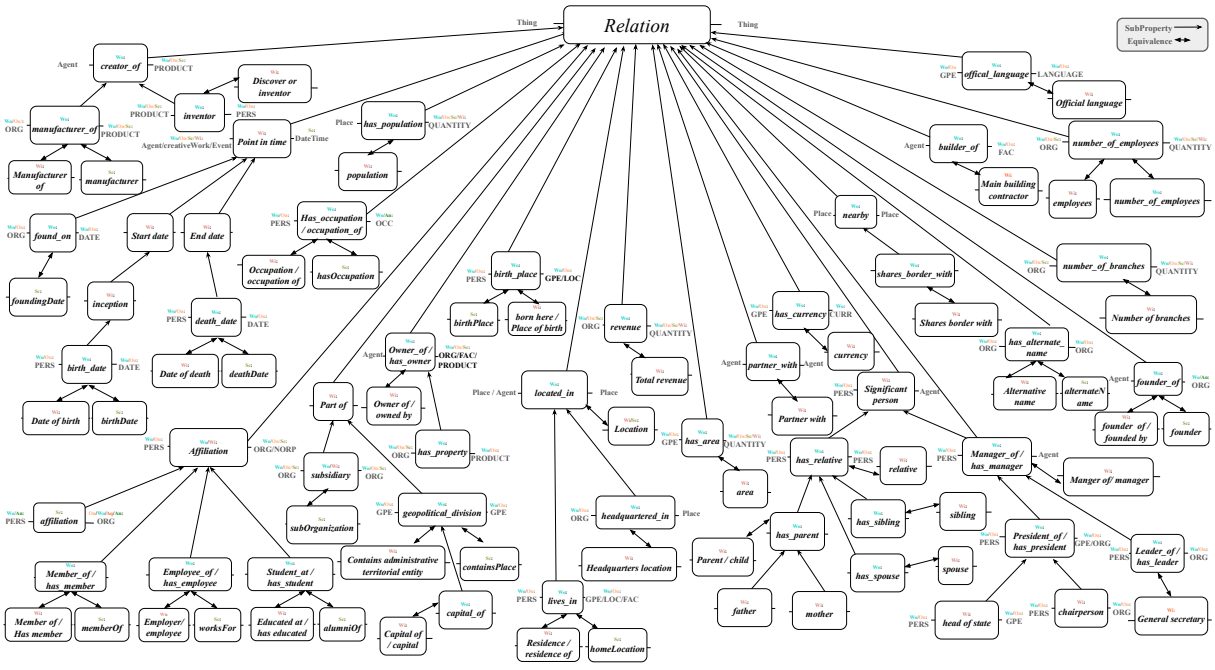

Figure 7: Relation Extraction Ontology

	WikiData			
Wojood Relations	**Property Name**	**Domain**	**Range**	**Subclass of**
has_parent	parent (P8810) / union of: father (P22), mother (P25)	Human (Q5)	Human (Q5)	relative (P1038)
has_spouse	P26: spouse	Human (Q5)	Human (Q5)	relative (P1038) / significant person (P3342)
has_sibling	P3373: sibling	person (Q215627)	person (Q215627)	relative (P1038)
has_relative	P1038: relative	Human (Q5)	Human (Q5)	significant person (P3342)
birth_date	P569: date of birth	Human (Q5)		inception (P571)
death_date	P570: date of death	humann, group of humans		end time (P582)
death_date	P570: date of death	humann, group of humans		dissolved, abolished or demolished date (P576)
birth_place	P19: place of birth	Human (Q5)	geographic location (Q2221906)	location (P276)
has_occupation	P106: occupation	Human (Q5), person (Q215627)	occupation (Q12737077)	root
has_conflict_with	P607: conflict	Human (Q5), group of humans (Q16334295), fictional military organization (Q18011141)	Conflict (Q180684)	participant in (P1344)
has_competitor	league or competition (P118)	Organization (Q43229)	Organization (Q43229)	participant in (P1344)
partner_with	P2652: partnership with	Organization (Q43229), administrative territorial entity (Q56061)	Organization (Q43229), administrative territorial entity (Q56061)	root
manager_of	P1037: manager/director	Human (Q5)	Organization (Q43229)	significant person (P3342)
president_of	P488:chairperson union	administrative territorial entity (Q56061)	Human (Q5)	significant person (P3342)
president_of	head of government (P6)	administrative territorial entity (Q56061), Organization (Q43229)	Human (Q5)	director / manager (P1037)
leader_of	general secretary (P3975)	Organization (Q43229)	Human (Q5)	significant person (P3342)
leader_of	general secretary (P3975)	Organization (Q43229)	Human (Q5)	director / manager (P1037)
geopolitical_division	P150: contains administrative territorial entity	Administrative Entity (Q56061), administrative territorial entity (Q56061)	Administrative Entity (Q56061), administrative territorial entity (Q56061)	has part(s) (P527)
subsidiary	P355: has subsidiary	Organization (Q43229)	Organization (Q43229)	owner of (P1830)
subsidiary	P355: has subsidiary	Organization (Q43229)	Organization (Q43229)	has part(s) (P527)
member_of	P463: member of	Any entity	Organization (Q43229)	part of (P361)
employee_of	P108: employer	Human (Q5), Organization (Q43229), group of humans (Q16334295)	Organization (Q43229)	affiliation (P1416)
student_at	P69: educated at	Human (Q5)	Educational Institution (Q2385804)	affiliation (P1416)
owner_of	P1830: owner of	Human (Q5), Organization (Q43229), group of humans (Q16334295)	Human (Q5), Organization (Q43229)	root
inventor_of	P61: discoverer or inventor	none	Human (Q5), facility (Q13226383), organization (Q43229), group of humans (Q16334295)	root
manufacturer_of	P176: manufacturer	Organization (Q43229), Human (Q5)	Product (Q2424752)	root
builder_of	main building contractor (P193)	Organization (Q43229)	Organization (Q43229), Human (Q5)	manufacturer (P176)
founder_of	P112: founded by	organization (Q43229), group of humans (Q16334295), website	Human (Q5), organization (Q43229), group of humans (Q16334295)	creator (P170)
lives_in	P551: residence	Human (Q5), group of humans (Q16334295)	Location (Q17334923)	location (P276)
located_in	P276: location	Entity	Location (Q17334923), facility (Q13226383), administrative territorial entity (Q56061)	root
headquartered_in	P159: headquarters location	Organization (Q43229)	Location (Q17334923), administrative territorial entity (Q56061)	significant place (P7153)
has_border_with	P47: shares border with	Geopolitical Entity (Q15642541)	Geopolitical Entity (Q15642541)	root
nearby				
has_property				
branch_count	P8368: number of branches	Organization (Q43229)	Quantity	root
org_has_revenue	P2139: total revenue	Organization (Q43229)	Monetary Value (Q13624636)	root
number_of_employees	P1128: employees	Organization (Q43229), facility	Quantity	root
org_found_date	P571: inception	root	-	start time (P580)
has_alternate_name	P4970: alternate names	-	-	root
geopolitical_entity_has_area	P2046: area			-
official_language	P37: official language	org, gpe, norp	-	language used (P2936)
has_currency	P38: currency	gpe, human	Currency (Q8142)	uses (P2283)
has_population	P1082: population	gpe, norp	Quantity	root
capital_of	P1376: capital of	Geopolitical Entity (Q15642541)	administrative territorial entity	located in the administrative territorial entity (P131)

Table 6: Mapping Wojood relations with Wikidata properties.

	Schema.org			
Wojood Relations	Property name	Property URI	Domain	Range
has_parent	parent	https://schema.org/parent	person	person
has_spouse	spouse	https://schema.org/spouse	person	person
has_sibling	sibling	https://schema.org/sibling	person	person
has_relative	relatedTo	https://schema.org/relatedTo	person	person
birth_date	birthDate	https://schema.org/birthDate	person	Date
death_date	deathDate	https://schema.org/deathDate	person	Date
birth_place	birthPlace	https://schema.org/birthPlace	person	Place
has_occupation	hasOccupation	https://schema.org/hasOccupation	person	occupation
has_conflict_with	-			
has_competitor	competitor	https://schema.org/competitor	sport event	person, sport team
partner_with	-			
manager_of	-			
president_of	-			
leader_of	-			
geopolitical_division	containedInPlace	https://schema.org/containedInPlace	place	place
subsidiary	subOrganization	https://schema.org/subOrganization	organization	organization
member_of	memberOf	https://schema.org/memberOf	person, organization	organization
employee_of	employee	https://schema.org/employee	organization	person
student_at	alumniOf	https://schema.org/alumniOf	person	organization
owner_of	owns	https://schema.org/owns	person, organization	product
inventor_of	creator	https://schema.org/creator	person, organization	creativework
manufacturer_of	manufacturer	https://schema.org/manufacturer	organization	product
builder_of	-			
founder_of	founder	https://schema.org/founder	organization	person, organization
lives_in	homeLocation	https://schema.org/homeLocation	person	place
located_in	location	https://schema.org/location	organization	place
headquartered_in	-			
has_border_with	-			
nearby	-			
has_property	-			
branch_count	-			
org_has_revenue	-			
number_of_employees	numberOfEmployees	https://schema.org/numberOfEmployees	organization	quantitative vlaues
org_found_date	foundingDate	https://schema.org/foundingDate	organization	Date
has_alternate_name	alternateName	https://schema.org/alternateName	thing	text
geopolitical_entity_has_area	-			
official_language	-			
has_currency	-			
has_population	-			
capital_of	-			

Table 7: Mapping *Wojood*^{Relations} with Schema.org properties.