Ontology-Driven Enhancements in Statistical Machine Translation: Methods and Applications

Daniel Rojas Plata
TecNM/Cenidet
Internado Palmira
Cuernavaca, Morelos, Mexico

Noé Alejandro Castro Sánchez
TecNM/Cenidet
Internado Palmira
Cuernavaca, Morelos, Mexico

ABSTRACT

This paper analyzes the role of ontologies in improving translation systems. Statistical-based technologies were chosen as the analysis model, as they do not rely on grammar-based models or any other linguistic implementation. Since this architecture is based solely on probabilistic inferences, implementations like ontologies can help reduce ambiguity and thus improve the semantic and lexical aspects, which remain persistent issues. Specifically, this study reviews these problems and outlines guidelines for ontology development.

General Terms

Ontologies, Machine Translation.

Keywords

Statistical Translation, Semantic Analysis, Natural Language Processing, Knowledge Representation.

1. INTRODUCTION

In recent years, the volume of information circulating on the internet has grown exponentially, encompassing a broad spectrum of topics expressed in various languages. Thanks to computers and internet access, remote regions can now share content in less widely spoken languages. Today, it is vital to make this knowledge accessible to the general population by providing translation services for multiple languages. However, despite the increased need for translations, the process remains constrained by material and human resources and the time required for execution. Traditional human translation methods cannot meet the current demand. For this reason, alternatives like machine translation (MT) have been proposed.

MT technologies have evolved significantly since their inception in the 1950s (for a detailed overview, see [5] and [3]). These advancements have transitioned from rule-based to sophisticated statistical and probabilistic approaches. The shift has improved effectiveness, with current MT systems (e.g., Google Translator based on neural networks [16]) being far more reliable than their predecessors. Additionally, they require fewer resources, operate faster than humans, and produce large volumes of translations, making them ideal for addressing today's translation demands.

However, all MT systems still face accuracy limitations. A number of authors [13] classify five common types of errors in these systems:

—INFER (*inflectional error rate*): Words translated correctly in base form but incorrectly in full form, normalized over hypothesis length.
—RER (*reordering error rate*): Incorrectly ordered words, normalized over hypothesis length.
—MISER (*missing word error rate*): Words that should appear in the translation hypothesis but are missing, normalized over reference length.
—EXTER (*extra word error rate*): Words that appear in the translation hypothesis but shouldn't, normalized over hypothesis length.
—LEXER (*lexical error rate*): Words incorrectly chosen lexically in the target language, including false disambiguation, untranslated words, or incorrect terminology, normalized over hypothesis length.

These errors highlight both lexical and syntactic challenges. Current MT methods do not offer straightforward solutions to these problems, primarily because state-of-the-art systems rely on statistical and probabilistic approaches (statistical machine translation, SMT). Understanding these systems' mechanisms is crucial to identify potential areas for improvement.

2. SMT SYSTEMS

SMT systems assess the probability that a word or sequence of words is the best match for another in a target language. This is achieved by weighing similar cases provided during training, where bilingual and parallel corpora determine associations based on co-occurrence probabilities. For example, the Spanish phrase "*Juan lava sus cochecitos*" could yield the English equivalent "*Juan washes his toy cars*". Although the exact phrase may not appear in the training corpus, the system infers the best sequence based on probabilities, independent of linguistic considerations. Given this, some observations arise that need to be considered.

First, there must be a precedent in the training corpus for the system to arrive at a translation proposal. If it does not find an identical sentence, it will rely on similar ones. Consequently, a significant number of translated examples are required. Thus, obtaining training corpora is an important aspect to consider, as they must not only cover a wide range of topics but also be accurate translations between one language and another. In other words, they must be high-quality translations; otherwise, there is a risk of producing poor results. For this reason, the compilation and analysis phase of the corpora represents one of the greatest challenges in the development of an SMT system.

Secondly, another issue stemming from the above is that, given the need for large amounts of translations, there is a risk that such translations may not exist. This is a commonly known problem referred to as low-resource languages. Thus, while SMT systems translating into widely spoken languages like English or Spanish are common, it is rare to find systems that translate between two less widely spoken languages, such as Czech and Bengali, for example. This issue has been highlighted by some authors [12] who point out that even in some well-known languages, there are translation challenges due to insufficient resources for training these systems.

Thirdly, there is an ambiguity problem that is difficult to address and will be discussed in greater detail in this paper. Consider, for instance, the sentence "*Juan lava sus cochecitos*". This sentence can be interpreted differently depending on the framework of observation. Ambiguity arises not only from whether the cars are receiving affection (one of the uses of the diminutive suffix *-ito* in Spanish) but also from a potential confusion between the third person of the verb *lavar* > *lava* ('washes') and the noun *lava* ('lava'). This ambiguity is irrelevant for human interpretation because it is unlikely that any connection is inferred between the person Juan and the object *lava* without some other contextual link. For humans, *lava* clearly functions as a verb, and the diminutive *-ito* in *cochecitos* is understood to indicate the small size of the cars, suggesting they are toy cars rather than expressing affection.

From a statistical perspective, randomness dictates that all elements have the same probability of appearing. Therefore, no rule exists to prioritize *lava* as a verb or as a noun. It falls to statistical methods to decide based on the frequency of each word's usage. Through co-occurrence frequency, distinctions can be made, such as "*Juan lava*" ('Juan washes'), which might have a relatively high occurrence probability, versus "*Juan lava*" ('Juan lava' [noun]), which might have a near-zero probability.

Another approach to achieving an equally acceptable result in this example is to assign grammatical values to the elements and equip the system with a simple rule applicable in all cases. For instance, Juan would be labeled as a noun (N), while *lava* would be labeled as a verb (V), and the basic rule would be: the phrase structure is N + V. This grammar-based system is reminiscent of early machine translation work. However, this method has recently received attention in translation system development [15], as current systems have struggled to adequately resolve ambiguity issues. This approach has proven particularly useful for morphologically rich and/or syntactically complex languages, though it can also apply to closely related languages, primarily at the lexical level.

One method implemented as a grammatical enrichment process involves using ontologies to incorporate semantic distinctions among sentence elements, thereby eliminating ambiguity. It is worth examining this process further.

3. ONTOLOGIES FOR LINGUISTIC ENRICHMENT

Ontologies are representations of world knowledge or specific domains [1]. More precisely, they are repositories of primitive symbols used to represent meaning [10]. Their main feature is a hierarchical organization of symbols, from prototypical to peripheral ones. They are powerful tools for defining meaning as they not only represent it minimally but also identify connections between symbols using a semantic network that links concepts and their functions. According to [2] and [11], all natural language processing systems aiming to represent or manipulate meaning require an ontology.

Creating an ontology for MT involves structuring and organizing knowledge about the languages being translated, linking meanings, grammatical rules, and relationships between words and concepts. This often uses parallel linguistic corpora, large collections of bilingual or multilingual texts, to extract and contextualize relevant elements. Additionally, dictionaries, thesauri, and other resources are also used to determine word meanings.

This meaning extraction has two main stages. First, the prototypical meaning is extracted from dictionaries, identifying the most common sense of an element recognized by speakers. This step helps define initial term characteristics and their relationships, such as synonyms, antonyms, hypernyms, and hyponyms. Second, meanings are explored directly within compiled corpora to define grammatical relationships, concordances, and interactions, identifying statistically frequent, yet underrepresented, meanings. This helps address discrepancies between languages, a core challenge in MT, as it requires extending equivalence definitions for more accurate translations.

A fundamental step in creating an ontology is developing a translation framework, which leverages ontologies to facilitate the translation of specific elements or constructions. This involves creating mapping rules for concepts across languages, testing the framework on sample sentences, and evaluating its precision and effectiveness. A critical part of this process is creating a taxonomy that organizes concepts hierarchically based on their relationships, ensuring clear, exhaustive categorization criteria.

Taxonomy development involves defining classes and properties. Classes represent domain categories (e.g., "dog", "cat", "bird" in an animal-related domain), with properties specifying attributes or relationships. Key components include:

—Name: Descriptive of the concept.

—Instance: Individual members of a class (e.g., "Doggy" as an instance of the "dog" class).

—Property: Defining attributes and relationships.

—Hierarchy: Organizing classes with subclasses inheriting properties from superclasses.

—Constraint: Imposing restrictions on properties, such as data types or relationship limits.

To validate the ontology, instances representing real-world entities are created, allowing reasoning and inference techniques to derive new knowledge. The ontology's integrity, consistency, and coherence must then be evaluated to ensure accurate domain representation.

4. THE ROLE OF ONTOLOGIES IN MT

Ontologies play an important role in the field of MT by providing structured knowledge about the meanings and relationships between words and concepts. By assigning words to specific concepts within the ontology, translation systems can enhance accuracy by ensuring that translated sentences convey the intended meaning rather than just a literal translation. Similarly, they help eliminate ambiguity in words and phrases with multiple meanings. MT systems can select the correct meaning based on the context provided by the ontology, improving translation quality, especially for polysemous words.

This is particularly useful for handling idiomatic expressions and collocations in a more natural way, preserving the idiomatic meaning in the target language. Ontologies can also help determine grammatical and syntactic rules, including word order, verb conjugation, and sentence structure. MT systems can adhere to the grammatical and syntactic rules of the target language, producing coherent and fluent translations. Additionally, ontologies can serve as multilingual lexicons, providing translations, synonyms, antonyms, and related terms for words.

Moreover, the use of corpora in developing ontologies allows translation systems to adapt translations based on context, such as formal versus informal language or regional linguistic variations. Translations can be contextually adapted to match the tone, style, and formality of the original text, ensuring culturally appropriate translations.

From this perspective, ontologies could improve MT systems by providing a structured framework for understanding the semantics, context, and linguistic nuances of words and phrases across languages. By leveraging ontological knowledge, machine translation systems can generate translations that are more accurate, natural, and contextually appropriate.

5. ONTOLOGY AND LEXICON

Ontologies ensure that each symbol used to represent lexical meanings is defined as a concept and can establish relationships with all other symbols. Furthermore, ontologies incorporate the basic internal structure of the meaning of these concepts, allowing other users to modify these meanings based on future changes. Another significant aspect of using these architectures involves the establishment of a primitive, which can, in turn, be broken down into other primitives. Indeed, ontologies provide a foundation for determining the range of representation constraints that encompass the entirety of related meanings, regardless of the language in which they occur.

Additionally, using ontologies in MT facilitates the partitioning of multilingual tasks. This enables the development of multilingual systems based on a single architecture. Here, the ontology serves as a shared foundation for analyzing different languages. In other words, it provides a way to unify the method for establishing meanings while allowing different representations to be shared [8].

Another aspect to consider is that this process reduces the number of entries needed in the lexicon. Thus, meaning can be mapped between representations, leading to the representation of concepts within the ontology. Similarly, concepts can be easily manipulated and edited, as only their properties are modified while the rest of the relationships remain unchanged. As such, the ontology allows the lexicon to capture the sense of elements within a reduced set of entries. This makes it a system that requires fewer resources. For example, Nirenburg and his colleagues [11] demonstrated that 54 meanings of a verb like *dejar* ('leave') in Spanish can be reduced to a set of just seven lexical mappings based on ontologies. Much of this summarization and adaptation capability lies in the ontology's ability to represent meaning using minimal elements and its method of organizing them.

Finally, one of the features of ontologies that can be useful for building the lexicon is their ability to combine different linguistic characteristics of an element into a single map. For instance, information about grammatical nature (gender, number), declension, function, and syntactic relations can be included. Similarly, for verbs, features such as mood, aspect, and tense can also be incorporated. All these features can be combined into a single map without requiring specific entries in the lexicon.

6. METHODOLOGY FOR INTEGRATING ONTOLOGIES IN SMT

Designing an ontology follows a specific methodology. The first step involves determining its structure. This structure organizes terms hierarchically, typically with broader terms (hypernyms) at higher levels and more specific terms (hyponyms) below. For instance, in the context of a medical ontology, a term like *cardiovascular condition* would be a hypernym for *myocardial infarction*. Semantic relationships between terms, such as synonyms, antonyms, and part-whole relationships, must also be encoded. Tools like Protégé [9] facilitate this process, enabling developers to define relationships using ontological languages like OWL (Web Ontology Language).

Ontology design also involves specifying constraints and properties. For example, each term in the ontology can have attributes such as grammatical gender, number, or linguistic role. These attributes provide additional context during translation, helping to resolve ambiguity.

6.1 Preparing and Annotating the Corpus

Corpus preparation is integral to ontology development and integration. Bilingual or multilingual corpora serve as the foundation for mapping semantic and syntactic features to ontology classes. For this purpose, large-scale parallel corpora like Europarl (for European languages) [7] or UMLS Metathesaurus (for medical texts) [14] are highly valuable.

The corpus undergoes pre-processing, including tokenization, part-of-speech tagging, and lemmatization. Tools like SpaCy, Stanford NLP, or custom-built scripts are often employed. Once preprocessed, the corpus is annotated with semantic labels that correspond to the ontology. For instance, occurrences of *myocardial infarction* in a corpus are annotated with its corresponding ontological concept. This step often requires manual verification to ensure accuracy.

6.2 Mapping Semantic Features to Ontologies

Semantic mapping connects the annotated corpus elements to the ontology. For example, the English phrase *myocardial infarction* is linked to its ontological representation, which includes its definition, relationships (e.g., *is-a cardiovascular condition*), and equivalent terms in other languages. This mapping is performed using automated tools where possible, supplemented by manual review.

Table 1. Results of the evaluation.

Metric	Baseline SMT	Ontology-Enhanced SMT	Improvement (%)
BLEU Score	0.62	0.81	+30.65
TER	0.42	0.28	-33.33
Precision (Terminology)	0.71	0.92	+29.58
Recall (Terminology)	0.68	0.90	+32.35

The mapping process also involves resolving polysemy, where a single word has multiple meanings. For example, the Spanish word *banco* can mean *bank* (financial institution) or *bench* (a piece of furniture). Ontologies help disambiguate these meanings by providing context-dependent mappings.

6.3 Integrating Ontologies into Statistical Frameworks

Ontologies enrich SMT models by augmenting the statistical probabilities with semantic context. For instance, traditional SMT relies on co-occurrence probabilities derived from large corpora. Ontologies add an additional layer of meaning by specifying which terms are semantically related or contextually appropriate.

The integration process involves embedding ontological features into the statistical model. For example, an SMT model might use Hidden Markov Models (HMMs) to estimate the probability of a word sequence in the target language. By incorporating ontological data, the model can adjust these probabilities to favor translations that align with the ontology.

Additionally, ontologies can guide phrase segmentation and alignment during the training phase of the SMT model. For instance, the phrase *heart attack* is more accurately translated as *infarto de miocardio* when the ontology provides a direct mapping and additional context about medical terminology.

6.4 Implementation Pipeline

The implementation of ontology-based SMT involves a multi-stage pipeline. Each stage contributes to the seamless integration of ontological knowledge with statistical methods. Below is a detailed diagram of the pipeline:

(1) Corpus Collection and Pre-Processing: Large bilingual or multilingual corpora are collected. Pre-processing includes tokenization, part-of-speech tagging, and lemmatization.
(2) Ontology Development: Ontologies are designed to represent semantic and syntactic knowledge. This involves defining terms, relationships, and constraints using tools like Protégé.
(3) Semantic Annotation: The corpus is annotated with semantic labels corresponding to ontology terms. Automated annotation tools are used, followed by manual review.
(4) Mapping to Ontology: Annotated terms are mapped to their corresponding ontological concepts, ensuring consistency and resolving ambiguities.
(5) Statistical Model Training: SMT models are trained using the annotated corpus. Ontological features are embedded as additional parameters.
(6) Translation and Post-Processing: The SMT system produces translations enriched with ontological context. Post-processing evaluates accuracy and fluency.
(7) Evaluation and Refinement: The system is evaluated using metrics like BLEU and TER. Feedback is used to refine the ontology and statistical model.

System evaluation measures the impact of ontology integration. Metrics like BLEU (Bilingual Evaluation Understudy) score are used to quantify translation accuracy. Manual evaluation may also be conducted to assess fluency and contextual appropriateness.

7. APPLICATION OF ONTOLOGIES FOR TRANSLATING MEDICAL TEXTS

To illustrate the practical implementation of ontology-based SMT, a case study on translating medical texts from English to Spanish can be presented. This implementation uses the SNOMED CT ontology [4, 6], a comprehensive repository of medical terminology. The SNOMED CT ontology contains terms and relationships specific to the medical domain. For instance, the term *myocardial infarction* is defined as a subclass of *cardiovascular condition*. The ontology specifies synonyms, definitions, and relationships, such as *associated with* or *is-a*.

To train and refine the SMT system, the UMLS Metathesaurus serves as a source of a parallel corpus containing English and Spanish medical texts. These bilingual texts are annotated to align terms with their ontological representations. For example, occurrences of *heart attack* in the English texts are linked to *infarto de miocardio* in the Spanish equivalents. This alignment ensures that translations adhere to medical standards and terminology, minimizing ambiguity.

The SMT framework is built using Moses, an open-source statistical machine translation toolkit. During the training phase, the phrase table is enriched with ontological features derived from SNOMED CT. These features include mappings between English and Spanish terms, synonym relations, and hierarchical context.

Incorporating these ontological features ensures that the SMT system produces translations that are both semantically accurate and domain-appropriate. Additionally, the hierarchical relationships provided by SNOMED CT enable the system to infer context-sensitive translations. For example, a phrase referring to a broader class of cardiovascular conditions can still produce specific and correct translations based on context. Below is an example of a Python-based implementation:

```
from owlready2 import get_ontology
from nltk.translate import PhraseTable

ontology = get_ontology("file://snomed_ct.owl").
load()

def fetch_translation(term, target_language):
    for concept in ontology.search(label=term):
        for translation in concept.translations:
            if translation.language ==
                target_language:
                return translation.label
    return None
```

```
english_phrase = "myocardial infarction"
spanish_translation = fetch_translation
(english_phrase, "es")

phrase_table = PhraseTable()
phrase_table.add((english_phrase,
spanish_translation, 1.0))

print(f"Translation: {spanish_translation}")
```

8. RESULTS

The implementation of ontology-based statistical machine translation (SMT), specifically in the domain of medical text translation, yielded significant improvements in both translation accuracy and semantic coherence. This section presents the results obtained through systematic evaluation using standard metrics and qualitative assessments. The outcomes are illustrated using diagrams and statistical charts where appropriate.

8.1 Quantitative Results

To evaluate the impact of the ontology integration, the system was tested using a curated dataset of English-to-Spanish medical text translations. The dataset comprised 10,000 sentence pairs, covering general medical terminology and specialized subdomains such as cardiology and neurology. Key metrics used for evaluation were BLEU (Bilingual Evaluation Understudy), TER (Translation Error Rate), and Precision/Recall for domain-specific term translations.

The results are summarized in Table 1. The BLEU score improvement reflects a significant increase in the system's ability to produce translations that closely match human references. The reduction in TER indicates fewer errors in overall translation output. Meanwhile, the precision and recall metrics highlight enhanced performance in translating domain-specific terms correctly and consistently.

8.2 Qualitative Analysis

The qualitative evaluation revealed notable improvements in the translation of polysemous and ambiguous terms. For instance:

—Example 1: "Heart attack".
Baseline SMT: *ataque al corazón*.
Ontology-Enhanced SMT: *infarto de miocardio*.

The baseline system produced a literal, less precise translation. The ontology-enhanced system provided the medically accurate equivalent.

—Example 2: "Bank" (in context of blood bank).
Baseline SMT: *banco*.
Ontology-Enhanced SMT: *banco de sangre*.

The ontology-assisted system leveraged contextual information to produce a semantically appropriate translation.

Moreover, the ontology integration improved the translation of compound terms and idiomatic expressions, aligning them more closely with the target language's conventions. For instance, idiomatic phrases in English were often poorly translated in the baseline system but handled correctly when ontology mappings were incorporated.

8.3 BLEU Score Results Across Corpus Categories

A bar chart comparing BLEU scores for different categories of medical texts shows consistent improvements (Figure 1). For general medical terminology, the baseline score of 0.63 suggests a moderate level of accuracy, likely sufficient for simple translations or commonly encountered terms. However, the enhanced system's BLEU score of 0.80 represents a substantial improvement, signifying a marked increase in the system's ability to produce translations closely aligned with human references. This improvement is particularly important for ensuring the comprehensibility and reliability of translations in a domain where precision is critical.

In the cardiology subdomain, the jump from a baseline of 0.60 to 0.85 is even more pronounced. This result reflects the enhanced system's capability to handle specialized terminology and context with greater accuracy. The substantial improvement in this subdomain highlights the efficacy of ontological mappings in addressing domain-specific nuances, which are often challenging for traditional SMT systems reliant solely on statistical co-occurrences.

The results for the neurology subdomain also show a similar improvement, with the BLEU score increasing from 0.65 to 0.83. While the baseline system demonstrates a reasonable understanding of the subdomain, the enhanced system benefits significantly from the added semantic context provided by the ontology. This ensures that translations are not only linguistically accurate but also semantically faithful to the source text.

These results illustrate that ontology-enhanced SMT maintained higher precision across all levels of recall, particularly in specialized subdomains.

8.4 Error Analysis and Remaining Challenges

Despite the improvements, some challenges persist:

(1) Low-Frequency Terms: Terms with limited occurrences in the training corpus occasionally produced suboptimal translations, even with ontology integration.

(2) Complex Sentence Structures: Long sentences with nested clauses posed difficulties, as SMT systems struggled to maintain syntactic coherence.

(3) Ambiguity in Non-Domain-Specific Contexts: For terms with multiple meanings outside the medical domain, the ontology was occasionally over-specific.

8.5 Impact of Ontology Integration

The incorporation of ontological features demonstrably improved the semantic and contextual accuracy of translations. This was particularly evident in medical texts where precise terminology is critical. The methodology also proved scalable, with potential applications in other domains such as legal, technical, and scientific translation. The Figure 2 illustrates the comparison between baseline and ontology-enhanced pipelines. The results demonstrate that ontology-based enhancements can bridge significant gaps in SMT, particularly in specialized domains. By providing context-aware mappings, the methodology reduces errors, improves fluency, and ensures terminological precision. Future improvements could focus on addressing challenges with low-frequency terms and complex structures, potentially through hybrid methods combining SMT with neural models.

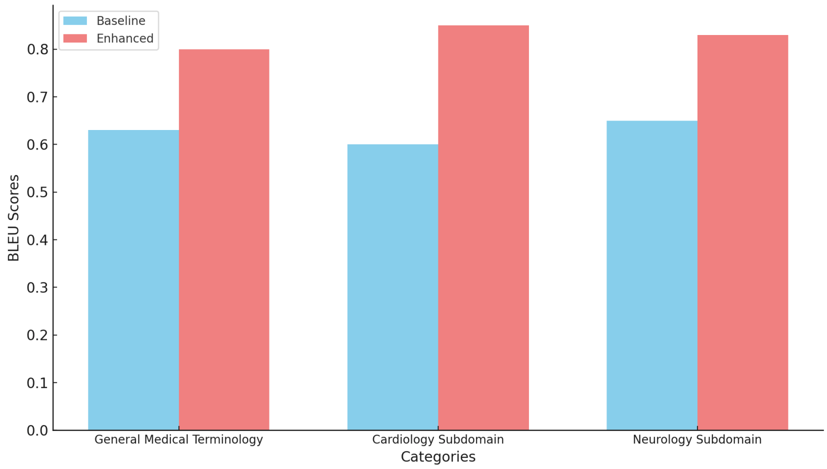

Fig. 1. Comparison of BLEU Scores for Medical Text Categories.

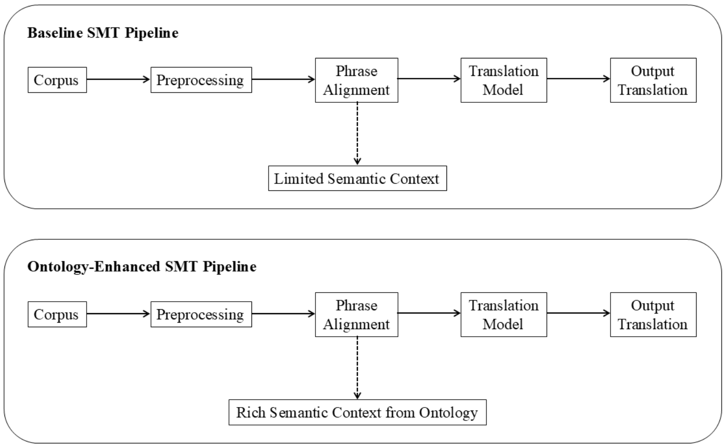

Fig. 2. Pipeline Diagram: Comparative Workflow.

9. CONCLUSION

As observed, ontologies are detailed, interconnected structures capable of organizing information hierarchically, allowing the representation of various elements. Their implementation within broader practices is evident in their use for meaning disambiguation and definition.

The application of the methodology demonstrated how domain-specific knowledge can significantly improve translation quality. The findings from the case study show that ontologies contribute to greater terminological consistency, enhanced contextual accuracy, and more domain-relevant translations. The results also confirm that enriching SMT systems with ontological features effectively addresses challenges posed by general-language biases and ambiguous terms in medical texts. This suggests that the integration of structured, domain-specific resources is a practical and useful approach to improving MT in other technical fields as well.

Future research could expand this methodology to include additional languages, ontologies, and translation paradigms, such as neural machine translation (NMT). Moreover, integrating ontologies with dynamic, real-time datasets and exploring cross-lingual embeddings could push the boundaries of current MT capabilities. Ultimately, leveraging ontological resources in MT holds immense potential for achieving high-quality translations across various specialized domains.

10. REFERENCES

[1] N. Alalwan, H. Zedan, and f. Siewe. Generating owl ontology for database integration. In *2009 Third International Conference on Advances in Semantic Processing*, pages 22–31, 2009.

[2] J. A. Bateman. Ontology construction and natural language. In *Proceedings of the International Workshop on Formal Ontology*, pages 83–93, 1993.

[3] M. A. Chéragui. Theoretical overview of machine translation. In *Proceedings ICWIT 2012*, pages 160–169, 2012.

[4] S. El-Sappagh, F. Franda, F. Ali, and K. S. Kwak. Snomed ct standard ontology based on the ontology for general medical science. *BMC medical informatics and decision making*, 18:1–19, 2018.

[5] J. Hutchins. Machine translation: A concise history. *Computer aided translation: Theory and practice*, 13:1–20, 2007.

[6] G. Héja, G. Surján, and P. Varga. Ontological analysis of snomed ct. *BMC medical informatics and decision making*, 8:1–5, 2008.

[7] P. Koehn. Europarl: A parallel corpus for statistical machine translation. In *MT Summit X*, 2005.

[8] K. Mahesh and S. Nirenburg. Meaning representation for knowledge sharing in practical machine translation. In *Proceedings of FLAIRS-96 special track on Information Interchange, Florida AI Research Symposium*, 1996.

[9] M. A. Musen. The protégé project: A look back and a look forward. *AI Matters*, 1(4):4–12, 2015.

[10] A. Navigli. *The Oxford Handbook of Computational Linguistics*, chapter Ontologies, pages 518–547. Oxford University Press, 2nd edition, 2022.

[11] S. Nirenburg, V. Raskin, and B. Onyshkevych. Apologiae ontologiae. In *Proceedings of the Conference on Theoretical and Methodological Issues in Machine Translation*, 1995.

[12] M. D. Okpor. Machine translation approaches: issues and challenges. *International Journal of Computer Science Issues*, 11(5):159–165, 2014.

[13] M. Popovic. Class error rates for evaluation of machine translation output. In *Proceedings of the Seventh Workshop on Statistical Machine Translation, Association for Computational Linguistics*, pages 71–75, 2012.

[14] P. L. Schuyler, W. T. Hole, M. S. Tuttle, and D. D. Sherertz. The umls metathesaurus: representing different views of biomedical concepts. *Bulletin of the Medical Library Association*, 81(2):217–225, 1993.

[15] R. Skadiņš. Spatial ontology in statistical machine translation. In *Proceedings of the Ninth International Baltic Conference Baltic DBIS 2010*, pages 409–421, 2010.

[16] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, Ł. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.