

LLM-Assisted Rule Based Machine Translation for Low/No-Resource Languages

Jared Coleman
Bhaskar Krishnamachari
Khalil Iskarous

University of Southern California
{jaredcol,bkrishna,kiskarou}@usc.edu

Ruben Rosales
rubenrosales@ieee.org

Abstract

We propose a new paradigm for machine translation that is particularly useful for no-resource languages (those without any publicly available bilingual or monolingual corpora): LLM-RBMT (LLM-Assisted Rule Based Machine Translation). Using the LLM-RBMT paradigm, we design the first language education/revitalization-oriented machine translator for Owens Valley Paiute (OVP), a critically endangered Indigenous American language for which there is virtually no publicly available data. We present a detailed evaluation of the translator’s components: a rule-based sentence builder, an OVP to English translator, and an English to OVP translator. We also discuss the potential of the paradigm, its limitations, and the many avenues for future research that it opens up.

1 Introduction

Large Language Models like OpenAI’s GPT series (OpenAI, 2023) have shown remarkable capability at an impressively wide range of natural language tasks (Bubeck et al., 2023) including machine translation (Hendy et al., 2023). These models work because they are trained on vast amounts of natural language data, primarily from the internet (OpenAI, 2023). For languages that are low-resource (languages for which there is little publicly available data) or no-resource (languages for which there is *no* publicly available data), models like these do not perform well on their own (Chowdhery et al., 2022; Robinson et al., 2023). There have been many efforts in improving machine translation for low-resource languages (see (Ranathunga et al., 2023) for a comprehensive survey), but no-resource languages have received much less attention in the literature.

Over past decades, researchers and community leaders have led many remarkable efforts in endangered language revitalization (Coronel-Molina and

McCarty, 2016; Taylor and Kochem, 2022) and reclamation (Baird, 2016). In this paper, we begin to explore how the impressive general-purpose language skills of LLMs might be helpful for these kinds of efforts by introducing a new paradigm for low/no-resource machine translation: LLM-RBMT (LLM-Assisted Rule-Based Machine Translation). The intuition behind our approach is that, given the correct tools, humans are good at translating simple sentences even for languages they do not know. For example, a common task for students in a language class is to translate sentences given a set of words and rules. Given a conjugation table and a set of verbs a motivated student could probably translate a list of simple subject-verb (e.g., “I eat”, “you sing”, etc.) sentences with extremely high accuracy. Of course, they are limited to only translating specific kinds of sentences with a limited vocabulary. Still, the idea is interesting: if we provide enough context to an LLM like ChatGPT, which has been shown to exhibit human-level performance on many natural language tasks, we may not need it to actually *know* (exhibit general fluency in) the target language we want to translate to/from. While rule-based translation can likely never achieve the quality of modern ML translators for high-resource languages, this technique has a few potential advantages for low/no-resource languages. First and most importantly, it requires no parallel corpus. Also, it is amenable to *partial translation*. Humans (and, as we’ll show in this paper, LLMs) are capable of giving partial translations when they don’t have all of the necessary vocabulary. For example, a Spanish student that doesn’t know the word for “dog” might still be able to translate the rest of the sentence “the dog ate the apple yesterday” by saying “el [dog] comió la manzana ayer”.

In this paper, we leverage LLMs (specifically, gpt-3.5-turbo and gpt-4 from OpenAI’s GPT series) to break down natural language sentences into

structured simple sentences compatible with hand-crafted rule-based translators. We also use them to turn sentence structured information (in JSON format) into natural language sentences. Using this approach, the LLM never interacts directly with the target language. Rather, we rely on the LLM to tell us how to use the simple, rule-based translators to provide a translation as close as possible to the user’s original input. We use this technique to build and evaluate two machine translation tools for Owens Valley Paiute (OVP, also called Eastern Mono or Monache in linguistics literature, ISO 639-3 mnr (SIL International, 2024)), a critically endangered Indigenous American language in the Uto-Aztecan language family (Moseley, 2010). The first is a selection-based OVP to English simple sentence translator and the second is an English to OVP translator that uses available vocabulary to construct translations of arbitrary user input sentences. The translators are oriented toward language teaching and revitalization. They are not designed to be general-purpose translators, but rather as tools to help no-resource language learners express ideas using simple sentence constructions.

1.1 Contributions

The main contributions of this work are¹:

1. An extensible OVP sentence builder for constructing valid subject-verb and subject-verb-object sentences in Owens Valley Paiute.
2. An LLM-assisted OVP to English translator that translates OVP sentence builder sentences to English with high accuracy.
3. An LLM-assisted English to OVP translator that translates English sentences to Owens Valley Paiute using the sentence-builder and OVP to English translation tools.
4. A novel methodology for the design and evaluation of no-resource language translators.

The rest of this paper is organized as follows. We discuss related work in Section 2. We present the sentence building tool and OVP to English translation system in Section 3. Then, in Section 4, we present the English to OVP translation system and report results on translation quality for different types of input sentences using embeddings models

¹The code behind all contributions is open-source: https://github.com/kubishi/kubishi_sentences

to measure semantic similarity. We conclude the paper with a summary of contributions and discussion of future research directions in Section 5.

2 Related Work

The landscape of low-resource machine translation is vast and constantly growing. A comprehensive survey on this subject is provided by (Ranathunga et al., 2023), which outlines the current techniques, guidance for selecting which techniques to use for a given language, and future directions for research. Of particular interest within this survey is the examination of unsupervised neural machine translation. While there exists promising research on constructing translators from minimal corpora, these methods invariably require *some* natural language data and thus have limited applicability to no-resource languages (such as OVP). The survey also discusses data augmentation strategies, including word or phrase replacement-based augmentation and Back-Translation-based Data Augmentation, both of which could potentially be integrated with some of the solutions presented in this paper (in particular the sentence builder and OVP to English translator to be presented in Section 3). Such an exploration is an interesting topic for future work. Other approaches discussed in the survey such as supervised learning, transfer learning, and semi-supervised learning are inapplicable to our scenario due to the absence of bilingual or monolingual corpora.

Contrary to the prevailing assumption in the literature that rule-based machine translation (RBMT) is a relic of the past, there remains active research and development in RBMT systems tailored for the most under resourced of languages (Khanna et al., 2021; Pirinen, 2019; Torregrossa et al., 2019). Recent work has also explored the utilization of Large Language Models (LLMs) for enhancing translation capabilities in low-resource languages through fine-tuning techniques (Lankford et al., 2023). Although this approach has shown promise in improving LLM translation quality for low-resource languages like (e.g., Irish), its reliance on bilingual corpora make it infeasible for no-resource languages like OVP. Recently, semantic similarity has been used to evaluate the quality of Machine Translation systems (Cao et al., 2022; Song et al., 2021). We observe that this technique is particularly useful for evaluating the quality of the English to OVP translator presented in this paper, due to the

lack of parallel corpora. See (Muennighoff et al., 2023) for a comprehensive benchmark of different embeddings models used for computing semantic similarities.

3 OVP to English Translation

In this section, we present an LLM-assisted selection-based OVP to English translator. The first piece of this translation system is a selection-based OVP sentence builder. The sentence builder allows users to select each of the following parts of speech from a list of choices until they form a valid sentence:

- Subject: The subject of the sentence.
- Subject Suffix: In OVP, noun subjects are always suffixed with either **-ii** (if the subject is proximal to the speaker) or **-uu** (if the subject is distant to the speaker).
- Verb: the verb of the sentence.
- Verb Suffix: In OVP, verb suffixes express the tense/aspect of the verb.
- Object: The object of the sentence (disallowed if the selected verb is intransitive, optional otherwise).
- Object Suffix: In OVP, noun objects are always suffixed with either **-(n)eika** (if the subject is proximal to the speaker) or **-(n)oka** (if the subject is distant to the speaker).
- Verb Object Pronoun Prefix: In OVP, object pronouns are prefixed to the verbs they are the object of. Object pronouns are required for all transitive verbs. Even when a noun object is specified, the pronoun is still required and should match the object suffix (**-(n)eika** matches **a-** or **ma-** and **-(n)oka** matches **u-**).

Not all of these parts of speech are required to create a valid sentence in OVP, though. In fact, some of them are incompatible. For example, if an intransitive verb is chosen, then it cannot have an object or object pronoun. In other words, the valid choices a user can make is a function of the choices they have already made. In our python implementation for this translator, we process each user selection and change the list of valid options for each part-of-speech based on their current selections to ensure they always create a valid OVP sentence. This is the *rule* part of LLM-Assisted *Rule*-Based Machine Translation that requires expert knowledge in the target language to implement. The entire vocabulary available for the users to select among fits on

a single page (and can be found in Appendix C).

After the user creates a valid OVP sentence, we translate it by first encoding the following sentence information into an English-only (using vocabulary definitions) *structured* simple sentence:

- Subject: noun or pronoun subject in English
- Subject Proximity: **proximal** if user selected subject suffix **-ii** or **distant** if the user selected subject suffix **-uu**
- Object: noun or pronoun object in English
- Object Proximity: **proximal** if user selected object suffix **-(n)eika** or **distant** if the user selected object suffix **-(n)oka**
- Verb: verb in English
- Verb Tense/Aspect: one of **past**, **present**, **future**, **past-continuous**, **present-continuous**, **present-perfect**

Then, we use few-shot prompting to encourage an LLM to transform the structured English data into a natural language sentence. Consider the example in Figure 1 where few-shot training examples (colored black) tell the LLM how to respond to the actual structured data for the randomly generated sentence for “Wo’ada-ii pagwi-noka u-zawa-dü.” (colored blue). Observe, the LLM is prompted to translate using only the English, structured version of the selected sentence.

To evaluate the accuracy of the translator, we generated 100 random valid OVP sentences by iteratively selecting a random choice among available choices for each of the parts of speech until the sentence is valid. Of the 100 random sentences generated, 98 were translated into English accurately using gpt-3.5-turbo model from OpenAI’s GPT-series. Translations and their accuracy labels can be found in Appendix A. While impressively accurate, this translator has many disadvantages. It only works for simple subject-verb and subject-verb-object sentences that use the nouns and verbs available in the system. Also, since pronouns and suffixes in OVP encode temporal/spatial information, translations don’t always capture full meaning of the sentence. The English translations are correct, but may be missing useful information. For example “kamü-uu wo’abi-neika a-düka-ti” translates to “the jackrabbit is eating the worm”, which is technically correct, but -uu also indicates that the jackrabbit is not present and the -neika indicates the worm *is* present. Then, since -ti can be used for both the present *and* past continuous tenses (*is*

Wo'ada-ii pagwi-noka u-zawa-dü.

system

You are an assistant for translating structured sentences into simple natural English sentences.

user

```
[[{'part_of_speech': 'subject', 'positional': 'proximal', 'word': 'wood'}, {'part_of_speech': 'object', 'positional': 'proximal', 'word': 'dog'}, {'part_of_speech': 'verb', 'tense': 'present continuous (-ing)', 'word': 'see'}]]
```

assistant

This wood is seeing this dog.

user

```
[[{'part_of_speech': 'subject', 'positional': 'distal', 'word': 'pinenuts'}, {'part_of_speech': 'object', 'positional': 'distal', 'word': 'horse'}, {'part_of_speech': 'verb', 'tense': 'future (will)', 'word': 'see'}]]
```

assistant

Those pinenuts will see that horse.

user

```
[[{'part_of_speech': 'subject', 'word': 'mosquito', 'positional': 'proximal'}, {'part_of_speech': 'object', 'word': 'fish', 'positional': 'distal'}, {'part_of_speech': 'verb', 'word': 'cook', 'tense': 'present'}]]
```

assistant

This mosquito is cooking that fish.

Figure 1: Few-shot examples for translating “Wo’ada-ii pagwi-noka u-zawa-dü.” using gpt-3.5-turbo.

[*x*]-ing or was [*x*]-ing), a better translation would be “the jackrabbit was eating this worm”. More advanced rules and better prompt-engineering may help mitigate this issue.

Despite some expected disadvantages, this translator has many advantages. First, it is the first machine translator for OVP. It is also easy to extend the tool with new nouns and verbs. Also, while implementing the rules required expert knowledge of what makes an OVP sentence valid, no expert knowledge of how the rules map to English was required (or needed to be programmed), thanks to the LLM. Finally, we believe this kind of translation system might be a useful educational tool that helps students learn how to build simple sentences. It also may be useful as a data augmentation technique for training neural machine translation models for low-resource languages.

4 English to OVP Translation

In this section, we present the first English to Owens Valley Paiute machine translator. Unlike the OVP to English translator presented in the previous section, the user can input any sentence in natural language. The translator works by first using an LLM to break the input sentence into a set of simple

structured subject-verb and subject-verb-object sentences, discarding any adjectives, adverbs, prepositions, objects (except for direct objects), etc. We encourage (through few-shot prompt engineering) the LLM to preserve as much semantic meaning as possible between the original sentence and the set of simple sentences. Consider the example in Figure 2 where few-shot examples (colored black) tell the LLM how to respond to a given input sentence “We are playing and laughing.” (colored blue)². Then, we use these structured sentences and available vocabulary to build valid OVP sentences with the sentence-building tool described in Section 3. Once the sentence is built, we use the translator described in Section 3 to translate the OVP sentences back into English. The idea is that, while some meaning may have been lost between the original input sentence and the final output English sentences, the user can be fairly confident (given the accuracy of the OVP to English translator) that the final translations are correct. The entire English to OVP translation process is depicted in Figure 3.

²We also leverage OpenAI’s function calling capability to ensure that responses are consistently formatted. We refer interested readers to the open-source implementation’s documentation at https://github.com/kubishi/kubishi_sentences for full details.

We are playing and laughing.

system

You are an assistant that splits user input sentences into a set of simple SVO or SV sentences. The set of simple sentences should be as semantically equivalent as possible to the user input sentence. No adjectives, adverbs, prepositions, or conjunctions should be added to the simple sentences. Indirect objects and objects of prepositions should not be included in the simple sentences.

user

I am sitting in a chair.

assistant

[[{'subject': 'I', 'verb': 'sit', 'verb_tense': 'present_continuous', 'object': null}]]

user

I saw two men walking their dogs yesterday at Starbucks while drinking a cup of coffee

assistant

[[{'subject': 'I', 'verb': 'see', 'verb_tense': 'past', 'object': 'man'}, {'subject': 'man', 'verb': 'walk', 'verb_tense': 'past_continuous', 'object': 'dog'}, {'subject': 'man', 'verb': 'drink', 'verb_tense': 'past_continuous', 'object': 'coffee'}]]

user

We are playing and laughing.

assistant

[[{'subject': 'we', 'verb': 'play', 'verb_tense': 'present_continuous', 'object': null}, {'subject': 'we', 'verb': 'laugh', 'verb_tense': 'present_continuous', 'object': null}]]

Figure 2: Few-shot training examples for the English to OVP using gpt-3.5-turbo.




Figure 3: The entire English to OVP translation process. The box with a red, dashed border indicates the set of sentences in Owens Valley Paiute (the target language) and the box with a blue, dashed border indicates the set of English sentences they translate to. Ideally, the input sentence, simple sentences, and English output sentences will have equivalent or very similar semantic meaning.

We evaluate the system by translating a set of 125 sentences. There are five types of sentences in the dataset (25 of each):

- subject-verb (e.g., “I read” or “she sings”)
- subject-verb-object (e.g., “Mom made dinner” or “John read a book”)
- two-verb (e.g., “She sings and dances.” or “I ate while watching TV.”)
- two-clause (e.g., “My brother drove and I waited.” or “Harry wrote and Ron read.”)

- complex (e.g., “Rachel and Monica share an apartment.” or “Romeo and Juliet loved deeply.”)

We translated all 125 sentences using both the gpt-3.5-turbo and gpt-4 models from OpenAI’s GPT-series, resulting in a total of 250 translations.

To measure the quality of each translation, we compute the semantic similarity between the input sentence and:

- the set of simple sentences generated by the LLM-powered segmenter (denoted *simple*).

- the set of simple sentences with unknown vocabulary removed (denoted *comparator*). Intuitively, this represents the best the translator can achieve given the vocabulary it has access to. For example, suppose the verb “wash” is not in our available vocabulary. Then the comparator sentence for the simple sentence “The woman is washing” would be “The woman is [VERB]-ing”
- the “round-trip” English translation (denoted *backwards*). This is the sentence produced using by translating the translated sentence (in OVP) to English using the method described in Section 3.

The usefulness of these measurements depends greatly on the function used to compute semantic similarity. We compute semantic similarity by generating embeddings for sentences (using some embeddings model) and computing the normalized cosine similarity between these embeddings. For our application, we want semantically similar sentences to have small normalized cosine similarity, independent of other linguistic features like syntax. For example, an ideal semantic similarity function would rank “an apple is eaten by a man” more similar to “a man eats apples” than the sentence “a woman drinks coffee”, despite the latter sentence being essentially grammatically equivalent to the target sentence.

We evaluated seven different embeddings models for this purpose and measured the semantic similarity between twelve target sentences and a ranked list of 10 sentences for each ranging from most to least semantically similar (sentences can be found in Appendix B). For each target sentence, we compare the ground-truth ranking of the 10 sentences to the ranking determined by the semantic similarity scores yielded by a particular embeddings model. We measure the similarity between the two rankings using two metrics: average displacement (average distance between a sentence’s position in the computed ranking and its position in the target ranking) and RBO (Rank-biased Overlap (Webber et al., 2010)). Table 1 tabulates the results of this evaluation. Results indicate that the all-MiniLM-L6-v2 embeddings model perform well with respect to both Average Displacement and RBO. For this reason, we run the rest of our experiments using this embeddings model for computing the semantic similarity between sentences.

We computed the semantic similarity between

all pairs of sentences in the dataset to establish a baseline for comparison. The mean semantic similarity between a pair of unrelated sentences was $\mu \approx 0.574$ with a standard deviation of $\sigma \approx 0.061$. Furthermore, the distribution appears to be relatively Gaussian (a histogram can be found in Appendix E). Intuitively, this suggests that semantically unrelated sentences are very unlikely to have a semantic similarity score of above $\mu + 3\sigma$ (i.e., greater than 0.757).

A good translation, then, should score high on all three semantic similarity metrics. For example, the translation

I am swimming.	
Simple	Semantic Similarity
I am swimming.	1
Comparator	
I am swimming.	1
Target	
Nüü pahabi-ti.	
Backwards	
I am swimming.	1
(model: gpt-3.5-turbo)	

is perfect. There are other interesting cases too, though. For example, when the comparator score is low but simple and backwards scores are high, the translator appears to do well with the vocabulary available, but is only able to give a partial translation. For example, the translation

Birds will migrate and return.	
Simple	Semantic Similarity
A bird will migrate.	
The bird will return.	0.955
Comparator	
A bird will [VERB].	
The bird will [VERB].	0.778
Target	
[migrate]-wei tsiipa-uu.	
[return]-wei tiipa-uu.	
Backwards	
That bird will migrate.	
That bird will return.	0.944
(model: gpt-4)	

has a high simple score, since the “Birds will migrate and return” is quite similar (semantically) to “A bird will migrate. The bird will return.”. The comparator score is much lower, however, since

Embeddings Model		Average Displacement		RBO	
		mean	std	mean	std
text-embedding-ada-002	(OpenAI, 2024a)	0.967	0.442	0.885	0.053
all-MiniLM-L6-v2	(Reimers and Gurevych, 2020)	0.933	0.323	0.884	0.050
text-embedding-3-small	(OpenAI, 2024b)	1.000	0.362	0.882	0.051
text-embedding-3-large	(OpenAI, 2024b)	0.917	0.463	0.882	0.054
paraphrase-MiniLM-L6-v2	(Reimers and Gurevych, 2020)	1.150	0.410	0.870	0.054
bert-base-uncased	(Reimers and Gurevych, 2019)	1.600	0.703	0.777	0.100
spacy/en_core_web_md	(Explosion, 2024)	1.833	0.466	0.760	0.090

Table 1: Quality of different embeddings models in measuring semantic similarity between sentences. A lower average displacement and higher RBO indicate a better embeddings model for this purpose.

the words for “migrate” and “return” are not available in the vocabulary. The backwards score is good because the sentence “[migrate]-wei tsiipa-uu. [return]-wei tiipa-uu.”, when translated using the OVP to English translator described in Section 3, becomes “That bird will migrate. That bird will return”, which is almost equivalent to the simple sentence. This example also highlights an advantage to our approach. No-resource language learners can use this system to understand how sentences should be structured even if it doesn’t have all of the vocabulary necessary to generate the entire translation. Furthermore, users can be fairly confident that the OVP target sentence is accurately translated into the English backwards sentence, due to the accuracy of the OVP to English translator.

The simple score being low indicates that a significant portion of the input sentence’s meaning is lost during segmentation (due, for example, to the input sentence containing adverbs, adjectives, prepositions, etc. that don’t fit anywhere in a simple subject-verb or subject-verb-object sentence). For example, the translation

My brother and I went hiking.	
Simple	Semantic Similarity
Brother went. I went.	0.794
Comparator	
[SUBJECT] went. I went.	0.608
Target	
mia-ku [brother]-uu. nüü mia-ku.	
Backwards	
The brother went. I went.	0.806
(model: gpt-3.5-turbo)	

lost meaning in the first step segmenting the input sentence into simple sentences because it chose to

use the verb "to go" instead of "hike" which is the main topic of the sentence. Perhaps a better way to have segmented this sentence would be: “Brother hiked. I hiked”. It may be possible to encourage the LLM to prefer “topic” verbs through prompt engineering.

Another interesting case is when the simple and comparator scores are high and only the backwards score is lower. This is observed in cases where there is ambiguity in OVP where there is not in English. For example, in the translation

She is cooking.	
Simple	Semantic Similarity
She is cooking.	1
Comparator	
She is cooking.	1
Target	
Uhu sawa-ti.	
Backwards	
He is cooking.	0.836
(model: gpt-4)	

“she” turns to “he” in the backwards translation because OVP does not have gendered pronouns. Despite the lower backwards score, this translation is accurate.

In general, both gpt-3.5-turbo and gpt-4 models do well with respect to the simple and backwards semantic similarity scores. Table 2 summarizes the mean semantic similarity scores for each model and type of sentence. Figure 4 depicts results for subject-verb sentences. Plots for the rest of the results can be found in Appendix D. That the simple and backwards scores are consistently higher than the comparator scores suggests that the translator can be greatly improved simply by expanding its




Figure 4: Results for subject-verb sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) \pm one, two, and three standard deviations, respectively.

Model	Type	Mean Sim.
gpt-3.5-turbo	subject-verb	0.941
	two-verb	0.906
	subject-verb-object	0.869
	two-clause	0.879
gpt-4	complex	0.829
	subject-verb	0.941
	subject-verb-object	0.866
	two-verb	0.905
	two-clause	0.877
	complex	0.830

Table 2: Translation qualities by model and sentence type: mean semantic similarities between input sentence and the simple, comparator, and backwards sentences produced during translation.

vocabulary. It is also interesting to note that the cheaper, weaker model performs quite well³.

5 Conclusion

In this paper, we present the first translation tool for Owens Valley Paiute (a critically endangered Indigenous American language) and, in doing so, propose a new methodology for low/no-resource machine translation: LLM-RBMT (LLM-Assisted Rule-Based Machine Translation). Due to a lack of bilingual or monolingual corpora, we use semantic

³The cost to run generate all translations was \$0.09 using gpt-3.5-turbo and \$5.48 using gpt-4.

similarity to measure the quality of translations, and observe interesting cases where the translator performs well, where it performs well given the vocabulary (primarily verbs and nouns) available to it, and where it performs poorly due to current implementation limitations.

This research opens up many directions for future work. First, expanding the translation tool (e.g., by adding more vocabulary, building versions for other languages, etc.) will make it more useful for the communities it is meant to serve. We are currently working to add more complex sentence structures to test the limits of this methodology. Building this system required expert knowledge of the target language. It may be interesting to explore how large language models might be able to help with this task as well (e.g., in finding patterns and implementing RBMT logic) to make it easier for non-experts (and even experts in the language who are less comfortable writing code) to build these kinds of translators. Finally, other methodologies for leveraging LLMs for low/no-resource machine translation should be explored. For example, a RAG (retrieval augmented generation) (Lewis et al., 2020) approach that searches a knowledge-base for known sentence structures, vocabulary, grammar rules, etc. and uses them to perform zero-shot translation might work well. The remarkable general-purpose language skills that LLMs exhibit make them a promising tool in helping revitalize critically endangered languages.

References

- Jessie Little Doe Baird. 2016. [Wopanaak language reclamation program: bringing the language home](#). *Journal of Global Indigeneity*, 2(2).
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of Artificial General Intelligence: Early experiments with GPT-4](#).
- Jialun Cao, Meiziniu Li, Yeting Li, Ming Wen, Shing-Chi Cheung, and Haiming Chen. 2022. [Semmt: A semantic-based testing approach for machine translation systems](#). *ACM Trans. Softw. Eng. Methodol.*, 31(2).
- Aakanksha Chowdhery et al. 2022. [PaLM: Scaling Language Modeling with Pathways](#).
- Serafin M. Coronel-Molina and Teresa L. McCarty. 2016. [Indigenous language revitalization in the americas](#).
- Explosion. 2024. [Industrial-Strength Natural Language Processing](#). <https://spacy.io/>. Accessed: 3 Mar 2024.
- Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. [How Good Are GPT Models at Machine Translation? A Comprehensive Evaluation](#).
- Tanmai Khanna, Jonathan N. Washington, Francis M. Tyers, Sevilay Bayath, Daniel G. Swanson, Tommi A. Pirinen, Irene Tang, and Hector Alòs i Font. 2021. [Recent advances in Apertium, a free/open-source rule-based machine translation platform for low-resource languages](#). *Machine Translation*, 35(4):475–502.
- Séamus Lankford, Haithem Afli, and Andy Way. 2023. [adaptmllm: Fine-tuning multilingual language models on low-resource languages with integrated LLM playgrounds](#). *Inf.*, 14(12):638.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Christopher Moseley. 2010. *Atlas of the World's Languages in Danger*. Unesco.
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2023. [MTEB: massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 2006–2029. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 Technical Report](#).
- OpenAI. 2024a. [New and Improved Embedding Model](#). <https://openai.com/blog/new-and-improved-embedding-model>. Accessed: 3 Mar 2024.
- OpenAI. 2024b. [New embedding models and API updates](#). <https://openai.com/blog/new-embedding-models-and-api-updates>. Accessed: 3 Mar 2024.
- Tommi A Pirinen. 2019. [Workflows for kickstarting RBMT in virtually no-resource situation](#). In *Proceedings of the 2nd Workshop on Technologies for MT of Low Resource Languages*, pages 11–16, Dublin, Ireland. European Association for Machine Translation.
- Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2023. [Neural machine translation for low-resource languages: A survey](#). *ACM Comput. Surv.*, 55(11):229:1–229:37.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Nathaniel R. Robinson, Perez Ogayo, David R. Mortensen, and Graham Neubig. 2023. [ChatGPT MT: Competitive for High- \(but not Low-\) Resource Languages](#).
- SIL International. 2024. [639 Identifier Documentation: mnr](#). Accessed: 11 Mar 2024.
- Yurun Song, Junchen Zhao, and Lucia Specia. 2021. [Sentsim: Crosslingual semantic evaluation of machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 3143–3156. Association for Computational Linguistics.
- Joshua Taylor and Timothy Kochem. 2022. [Access and empowerment in digital language learning, maintenance, and revival: a critical literature review](#). *Diaspora, Indigenous, and Minority Education*, 16(4):234–245.
- Daniel Torregrosa, Nivranshu Pasricha, Maraim Masoud, Bharathi Raja Chakravarthi, Juan A. Alonso, Noe Casas, and Mihael Arcan. 2019. [Leveraging](#)

rule-based machine translation knowledge for under-resourced neural machine translation models. In *Proceedings of Machine Translation Summit XVII Volume 2: Translator, Project and User Tracks, MT-Summit 2019, Dublin, Ireland, August 19-23, 2019*, pages 125–133. European Association for Machine Translation.

William Webber, Alistair Moffat, and Justin Zobel. 2010. [A similarity measure for indefinite rankings](#). *ACM Trans. Inf. Syst.*, 28(4):20:1–20:38.

Appendices

A OVP to English Translator Results

sentence	translation	label
tei-zawa-ku ihi	This cooked us.	1
isha'-uu tüba-neika ai-hibi-pü	That coyote has drunk these pinenuts.	1
tüba-uu tsibui-dü	Those pinenuts climb.	1
ta-ŵui-gaa-wei mahuŵa	"They are going to write to us, you and I."	1
wo'ada-uu aponu'-oka u-zawa-ti	The mosquito is cooking the apple.	1
tübbi-uu tüwoobü-neika a-yadohi-pü	The rock has talked to the earth.	1
payahuupü-uu katü-ti	That river is sitting.	1
kwisha'i-wei üü	You will sneeze.	1
isha'-ii tübbi-neika mai-ŵui-gaa-wei	This coyote is going to write these rocks.	1
toni-uu wünü-ti	The wickiup is standing.	1
isha'pugu-neika ihi mai-dama'i-ku	This found these dogs.	1
wo'ada-neika ihi mai-dama'i-gaa-wei	This will find these mosquitoes.	1
tabuutsi'-uu tüba-noka u-buni-ku	The cottontail saw those pinenuts.	1
maishibü-neika uhu ai-nia-ti	He/she/it is reading these corn.	1
koopü'-ii wükada-noka ui-nia-ku	This coffee read those bird snakes.	1
tuunapi-uu waakü-pü	That food has worked.	1
katünu-ii koopü'-oka ui-nobini-wei	This chair will visit those coffees.	1
aingwü-neika mahu ma-nia-dü	He/she/it reads the squirrel.	1
maishibü-uu wükihaa-gaa-wei	That corn is going to smile.	1
isha'-uu isha'-oka ui-zawa-wei	That coyote will cook those coyotes.	1
aingwü-ii tsibui-dü	This squirrel climbs.	1
katünu-ii tübbi-neika ma-buni-wei	This chair will see this rock.	1
wükada-uu tsibui-ku	The bird snake climbed.	1
wo'ada-uu paya-neika ma-hibi-ti	The mosquito is drinking the water.	1
pagwi-neika mahu ma-ŵui-pü	He/she/it has written/is writing this fish.	1
tabuutsi'-uu isha'pugu-neika mai-nobini-gaa-wei	That cottontail is going to visit those dogs.	1
paya-neika mahu mai-hibi-gaa-wei	He/she/it is going to drink this water.	1
nishua'i-pü nüü	I am laughing.	1
aaponu'-ii küna-neika a-düka-pü	This apple has eaten this wood.	1
katü-dü uhu	He/she/it sits.	1
mukita-uu isha'pugu-noka u-naki-ti	The lizard is chasing the dog.	1
isha'-oka üü ui-dama'i-ku	You found those coyotes.	1
pahabichi-uu wükihaa-dü	That bear smiles.	1
pahabichi-ii wo'abi-noka ui-naka-dü	This bear hears those worms.	1
habi-ku ihi	This lay down.	1
tümüi-ku taa	You and I wrote.	1
üwi-ku ihiŵa	These slept.	1
tübbi-uu ta-naka-ku	"That rock heard us, you and I."	1
wo'ada-uu tei-gwana-dü	That mosquito smells us.	1
tümüi-dü mahuŵa	They write.	1
üwi-dü nüügwa	We are sleeping.	1
tsibui-ti mahu	He/she/it is climbing.	1
mukita-uu tsibui-pü	That lizard has climbed.	1
payahuupü-uu toyabi-neika ma-zawa-gaa-wei	The river is going to cook the mountain.	1

nobi-uu kwisha'i-ku	That house sneezed.	1
kamü-uu wükihaa-dü	That jackrabbit smiles.	1
toni-uu katü-wei	That wickiup will sit.	1
aingwü-uu katünu-noka u-zawa-gaa-wei	The squirrel is going to cook that chair.	1
paya-uu pasohobü-neika ai-buni-wei	That water will see those trees.	1
toyabi-ii tsibui-ku	The mountain climbed.	1
tsibui-wei taa	You and I will climb.	1
pugu-uu wo'abi-neika ai-naki-ku	That horse chased those worms.	1
mukita-uu wai-noka u-nobini-gaa-wei	The lizard is going to visit the rice.	1
wükihaa-ti mahu	He/she/it is smiling.	1
tüsüga-ii tüwoobü-neika ma-naki-gaa-wei	This weasel is going to chase this earth.	1
yadoha-ku uhuŵa	They talked.	1
pahabichi-ii pugu-noka ui-nia-ku	The bear read those horses.	1
paya-uu katünu-noka ui-yadohi-gaa-wei	Water is going to talk to those chairs.	1
pagwi-ii wo'abi-noka ui-düka-ti	This fish is eating those worms.	1
tabuutsi'-uu tübinohi-ku	That cottontail played.	1
tünia-ku nüü	I read.	1
poyoha-pü ihiŵa	These are running.	0
mukita-uu yotsi-wei	That lizard will fly.	1
tabuutsi'-uu tübbi-neika ma-buni-pü	That cottontail has seen this rock.	1
isha'pugu-ii tüba-neika ai-nobini-ku	This dog visited these pinenuts.	1
isha'-uu katü-dü	That coyote sits.	1
pasohobü-ii tsiipa-noka ui-naka-ku	The tree heard those birds.	1
kamü-uu mukita-neika ma-zawa-wei	That jackrabbit will cook this lizard.	1
tuunapi-uu tümui-gaa-wei	That food is going to write.	1
wai-uu aingwü-neika ai-naka-pü	That rice has heard those squirrels.	1
tsiipa-uu pugu-noka ui-naka-pü	The bird has heard those horses.	1
pagwi-ii wükihaa-gaa-wei	These fish are going to smile.	1
tei-nobini-dü uhuŵa	They visit us.	1
paya-neika ihi ma-dama'i-dü	This finds this water.	1
isha'pugu-neika nüügwa ma-düka-ku	We ate this dog.	1
tabuutsi'-uu hubiadu-dü	The cottontail sings.	1
kwadzi-ii yadoha-wei	This tail will talk.	1
isha'pugu-uu mukita-neika ai-naka-ti	That dog is hearing these lizards.	1
pasohobü-noka üü u-zawa-gaa-wei	You are going to cook that tree.	1
tsiipa-noka uhu u-buni-pü	He/she/it has seen or is seeing that bird.	1
tuunapi-neika mahu mai-gwati-ku	He/she/it hit the foods here.	0
isha'-eika nüügwa ai-ŵui-gaa-wei	We are going to write coyotes.	1
pasohobü-uu toyabi-neika ma-hibi-pü	The tree has drunk the mountain.	1
aingwü-uu aponu'-eika ai-naka-ti	The squirrel is hearing these apples.	1
tabuutsi'-uu wai-noka u-yadohi-pü	The cottontail has talked to the rice.	1
katünu-noka mahuŵa ui-gwati-dü	They are hitting those chairs.	1
tüwoobü-neika uhuŵa mai-buni-ti	They are seeing these earths.	1
koopii'-uu hukaŵia-ti	Coffee is walking.	1
küna-ii ni-naka-ti	This wood is hearing us.	1
toyabi-neika taagwa ma-buni-gaa-wei	We are going to see this mountain.	1
isha'pugu-ii üwi-dü	This dog sleeps.	1
pagwi-neika ihi mai-naka-gaa-wei	This will hear these fish.	1
kwadzi-ii toni-neika ai-gwati-gaa-wei	This tail is going to hit those wickiups.	1

paya-ii tuunapi-noka u-düka-ti	This water is eating that food.	1
mukita-uu tümui-gaa-wei	That lizard is going to write.	1
pahabichi-uu küna-neika ma-zawa-ku	The bear cooked the wood.	1
isha`pugu-uu tabuutsi`-eika a-zawa-dü	That dog is cooking this cottontail.	1
katünu-uu pahabichi-noka u-naki-ku	The chair chased the bear.	1
küna-uu waakü-gaa-wei	That wood is going to work.	1
pugu-neika mahu ai-naka-ku	He/she/it heard these horses.	1

Table 3: One hundred randomly generated sentences (using the OVP sentence building tool described in Section 3) and translated (using the OVP to English translator described in Section 3) labeled 1 if the translation is accurate and 0 otherwise.

B Semantic Similarity Sentences

Base Sentence	Other Sentences (in order of most to least semantically similar to Base Sentence)
She sings.	He sings. He/she/it sings. She performs a song. A song is being sung by her. She hums a tune. She listens to music. She dances. She eats. The cat sleeps. Mountains echo silently.
The dog fell.	The dog fell yesterday. A dog stumbled. The puppy tripped over. The cat is running. An animal is in motion. The bird flies. Leaves fall in autumn. He reads a book. Clouds cover the sky. Apples on the moon are hungry.
The man ate an apple.	The apple was eaten by the man. A man consumes a fruit. The boy nibbles on an apple. Someone is eating. He drinks water. The woman ate a pie. A cat chases a mouse. Trees grow in the forest. The car is red. Stars twinkle at night.
The sun rises in the east.	The east welcomes the sunrise. Sunrise occurs in the east. Day breaks in the east. The moon sets in the west. The stars shine at night.

	<p>Clouds gather before rain. The wind changes direction. Leaves fall in autumn. Snow covers the mountains. A book rests on the table.</p>
Birds fly south for the winter.	<p>For winter, birds head south. Migratory birds travel south when it gets cold. Birds migrate to warmer climates during winter. Fish swim upstream. Bears hibernate in winter. Flowers bloom in spring. The earth orbits the sun. Trees lose their leaves in fall. The sky is blue. A cat sleeps on the couch.</p>
I read a book yesterday.	<p>Yesterday, I finished reading a book. A book was read by me yesterday. I watched a movie last night. I'll visit the library tomorrow. She writes a letter. He cooks dinner. They are painting a house. The sun sets in the evening. A dog barks at night. The car needs fuel.</p>
The cake was delicious.	<p>Delicious was the cake. The dessert tasted great. We enjoyed the tasty cake. The pie is sour. Coffee complements breakfast. Leaves rustle in the wind. A bird sings outside. Children play in the park. Traffic is heavy today. The phone is ringing.</p>
Lightning precedes thunder.	<p>Thunder follows lightning. First comes lightning, then comes thunder. The storm brings lightning and thunder. Rain refreshes the earth. The sun warms the ground. A river flows to the sea. Mountains reach towards the sky. A cat chases a mouse. Books fill the shelf. The clock ticks steadily.</p>
She painted a beautiful picture.	<p>A beautiful picture was painted by her. The painting she created is beautiful. She sketches a portrait. He writes a poem. They are filming a movie. Birds nest in spring.</p>

	<p>Flowers wilt in the heat. Kids play video games. Cars fill the parking lot. The sun sets late in summer.</p>
The computer is broken.	<p>A broken state afflicts the computer. The machine isn't working. We need to repair the computer. The phone's battery is dead. Lights flicker during a power outage. A book lies open on the desk. Water boils at 100 degrees Celsius. A cat purrs contentedly. The door creaks when opened. Birds migrate in autumn.</p>
He solved the puzzle quickly.	<p>The puzzle was quickly solved by him. Quickly, he found the solution to the puzzle. She completes the crossword. The mystery remains unsolved. A race against time. Flowers are sold at the market. The river cuts through the valley. A key unlocks the door. Leaves turn red in autumn. The train arrives at noon.</p>
The stars twinkle at night.	<p>At night, the stars shimmer. Twinkling stars fill the night sky. Night unveils a sky full of stars. The moon glows brightly. Clouds mask the moon. The sun sets, stars appear. A comet streaks through the sky. Fireflies glow in the dark. Crickets chirp in the evening. A candle flickers in the window.</p>

Table 4: Base sentences and other sentences ordered by their semantic similarity to the base sentence (as determined by authors).

C Vocabulary

		<i>isha'</i>	coyote		
		<i>isha'pugu</i>	dog		
		<i>kidi'</i>	cat		
<i>tüka</i>	eat	<i>pugu</i>	horse		
<i>puni</i>	see	<i>wai</i>	rice		
<i>hibi</i>	drink	<i>tüba</i>	pinenuts		
<i>naka</i>	hear	<i>maishibü</i>	corn	<i>nüü</i>	I
<i>kwana</i>	smell	<i>paya</i>	water	<i>uhu</i>	he/she/it
<i>kwati</i>	hit	<i>payahuupü</i>	river	<i>uhuwa</i>	they
<i>yadohi</i>	talk to	<i>katünu</i>	chair	<i>mahu</i>	he/she/it
<i>naki</i>	chase	<i>toyabi</i>	mountain	<i>mahuwa</i>	they
<i>tsibui</i>	climb	<i>tuunapi</i>	food	<i>ihi</i>	this
<i>sawa</i>	cook	<i>pasohobü</i>	tree	<i>ihüwa</i>	these
<i>tama'i</i>	find	<i>nobi</i>	house	<i>taa</i>	you and I
<i>nia</i>	read	<i>toni</i>	wickiup	<i>nüügwa</i>	we (exclusive)
<i>mui</i>	write	<i>apo</i>	cup	<i>taagwa</i>	we (inclusive)
<i>nobini</i>	visit	<i>küna</i>	wood	<i>üü</i>	you
		<i>tübbi</i>	rock	<i>üügwa</i>	you (plural)
		<i>tabuutsi'</i>	cottontail		
		<i>kamü</i>	jackrabbit		
		<i>aaponu'</i>	apple		
		<i>tüsüga</i>	weasle	<i>ii</i>	(proximal)
		<i>mukita</i>	lizard	<i>uu</i>	(distal)
		<i>wo'ada</i>	mosquito		
		<i>wükada</i>	bird snake		
		<i>wo'abi</i>	worm	<i>i</i>	me
		<i>aingwü</i>	squirrel	<i>u</i>	him/her/it (distal)
		<i>tsiipa</i>	bird	<i>ui</i>	them (distal)
		<i>tüwoobü</i>	earth	<i>ma</i>	him/her/it (proximal)
		<i>koopü'</i>	coffee	<i>mai</i>	them (proximal)
		<i>pahabichi</i>	bear	<i>a</i>	him/her/it (proximal)
		<i>pagwi</i>	fish	<i>ai</i>	them (proximal)
		<i>kwadzi</i>	tail	<i>ni</i>	us (plural, exclusive)
				<i>tei</i>	us (plural, inclusive)
				<i>ta</i>	us (dual), you and I
				<i>ü</i>	you (singular)
				<i>üü</i>	you (plural), you all
		<i>ku</i>	completive (past)		
		<i>ti</i>	present ongoing (-ing)		
		<i>dü</i>	present		
		<i>wei</i>	future (will)		
		<i>gaa-wei</i>	future (going to)	<i>eika</i>	(proximal)
		<i>pü</i>	have x-ed, am x-ed	<i>oka</i>	(distal)

(a) Transitive Verbs

<i>katü</i>	sit
<i>üwi</i>	sleep
<i>kwisha'i</i>	sneeze
<i>poyoha</i>	run
<i>mia</i>	go
<i>hukawa</i>	walk
<i>wünü</i>	stand
<i>habi</i>	lie down
<i>yadoha</i>	talk
<i>kwatsa'i</i>	fall
<i>waakü</i>	work
<i>wükihaa</i>	smile
<i>hubiadu</i>	sing
<i>nishua'i</i>	laugh
<i>tsibui</i>	climb
<i>tübinohi</i>	play
<i>yotsi</i>	fly
<i>nüga</i>	dance
<i>pahabi</i>	swim
<i>tünia</i>	read
<i>tümui</i>	write
<i>tsiipe'i</i>	chirp

(b) Intransitive Verbs

(c) Nouns

<i>ku</i>	completive (past)
<i>ti</i>	present ongoing (-ing)
<i>dü</i>	present
<i>wei</i>	future (will)
<i>gaa-wei</i>	future (going to)
<i>pü</i>	have x-ed, am x-ed

(d) Object Suffixes

(e) Subject Pronouns

<i>ii</i>	(proximal)
<i>uu</i>	(distal)

(f) Subject Suffixes

<i>i</i>	me
<i>u</i>	him/her/it (distal)
<i>ui</i>	them (distal)
<i>ma</i>	him/her/it (proximal)
<i>mai</i>	them (proximal)
<i>a</i>	him/her/it (proximal)
<i>ai</i>	them (proximal)
<i>ni</i>	us (plural, exclusive)
<i>tei</i>	us (plural, inclusive)
<i>ta</i>	us (dual), you and I
<i>ü</i>	you (singular)
<i>üü</i>	you (plural), you all

(g) Object Pronouns

<i>eika</i>	(proximal)
<i>oka</i>	(distal)

(h) Object Suffixes

Table 5: Vocabulary available in sentence building system.

D English to OVP Translation Results




Figure 5: Results for subject-verb sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) +/- one, two, and three standard deviations, respectively.

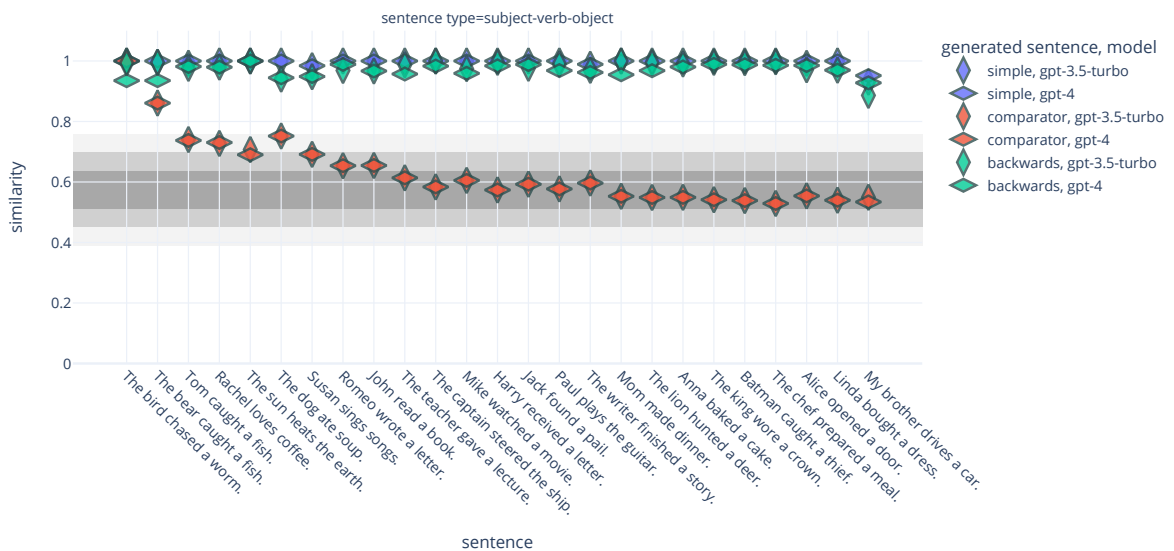


Figure 6: Results for subject-verb-object sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) +/- one, two, and three standard deviations, respectively.




Figure 7: Results for two-verb sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) +/- one, two, and three standard deviations, respectively.




Figure 8: Results for two-clause sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) +/- one, two, and three standard deviations, respectively.




Figure 9: Results for complex sentences. The dark, medium, and light gray bands represent the baseline similarity (between unrelated sentences in the dataset) +/- one, two, and three standard deviations, respectively.

E Semantic Similarity Baseline Score




Figure 10: Distribution of semantic similarity scores between all pairs of sentences in the dataset.