

Characterizing the Effects of Translation on Intertextuality using Multilingual Embedding Spaces

Hope McGovern¹ Hale Sirin² Tom Lippincott²

¹ Department of Computer Science & Technology, University of Cambridge, U.K.

² Center for Digital Humanities, Johns Hopkins University, Baltimore, U.S.A.

¹ hope.mcgovern@cl.cam.ac.uk ² {hsirin1, tom.lippincott}@jhu.edu

Abstract

Rhetorical devices are difficult to translate, but they are crucial to the translation of literary documents. We investigate the use of multilingual embedding spaces to characterize the preservation of intertextuality, one common rhetorical device, across human and machine translation. To do so, we use Biblical texts, which are both full of intertextual references and are highly translated works. We provide a metric to characterize intertextuality at the corpus level and provide a quantitative analysis of the preservation of this rhetorical device across extant human translations and machine-generated counterparts. We go on to provide qualitative analysis of cases wherein human translations over- or underemphasize the intertextuality present in the text, whereas machine translations provide a neutral baseline. This provides support for established scholarship proposing that human translators have a propensity to amplify certain literary characteristics of the original manuscripts.¹

1 Introduction

Coined by the semiotician and literary critic Julia Kristeva in 1969, *intertextuality* is a term that encompasses the ways in which one piece of text can refer to another (Kristeva, 1986 [1969]). It can range from direct quotation to semantic resemblance, both within and between works, highlighting that “no text is an island,” and that a text can only be understood as part of a matrix of other texts, impacting both literary theory and translation theory that followed (Moyise, 2002). For example, intertextual allusions can be seen throughout James Joyce’s retelling of Homer’s *Odyssey* in his 1922 novel, *Ulysses*, realized through a broad range of linguistic and narrative correspondences (Currie, 2019), such as the pairing between characters from

¹All code and data available at <https://github.com/comp-int-hum/literary-translation>

each book: Molly/Penelope, Stephen/Telemachus and Leopold/Odysseus.

As earlier scholarship on computational detection of intertextuality points out, intertextual references have two main functions: to express similarity between two passages, “so that the latter can be interpreted in light of the former”; but also to highlight their differences, in that the earlier context they reference can be revised (Bamman and Crane, 2008). For example, in the film *The Matrix* (1999) the white rabbit serves as an intertextual reference to *Alice’s Adventures in Wonderland* (1865) by Lewis Carroll. However, inverting the original context in which Alice was falling into a dreamland, Neo is now waking up from one. Intertextual references of this type set up a link akin to a two-way “traffic”—inviting both similarities and differences (Hays, 1989). They are a prominent feature of Classical texts, notably the New Testament and its references to the Hebrew Bible (Bamman and Crane, 2008). Table 1 shows one such example. Biblical intertextuality can range from explicit quotation to echoes of formulaic language, and many examples have since been cataloged (Hays, 1989).

Detecting intertextual references contributes toward a contextualized understanding of the “full semiotic density” of a any given text (Broderick, 2017) and therefore identifying intertextuality and the degree to which it is preserved in translation is crucial for the interpretation and appreciation of literary and historical texts. Due to its significance, computational methods for identifying intertextuality have become an expanding field of research, and it is closely connected to other NLP tasks that are grouped under narrative reasoning and comprehension (Sang et al., 2022; Piper et al., 2021). Significant attention has been devoted to identifying text reuse (implicit intertextuality) in Biblical text (Lee, 2007; Moritz et al., 2016), classical Latin poetry (Burns et al., 2021; Bamman and Crane, 2008), Latin prose (Dexter et al., 2017), and Romantic po-

Exodus 14:21 <i>Hebrew Bible</i>	Revelation 16:12 <i>New Testament</i>
20 [...] and it was a cloud and darkness to them, but it gave light by night to these: so that the one came not near the other all the night.	11 And blasphemed the God of heaven because of their pains and their sores, and repented not of their deeds.
21 And Moses stretched out his hand over the sea; and the LORD caused the sea to go back by a strong east wind all that night, and made the sea dry land, and the waters were divided.	12 And the sixth angel poured his vial upon the great river Euphrates; and the water thereof was dried up, that the way of the kings of the east might be prepared.
22 And the children of Israel went into the midst of the sea upon the dry ground [...]	13 And I saw three unclean spirits like frogs come out of the mouth of the dragon [...]

Table 1: **Biblical intertextuality.** The highlighted middle verse shows the intertextual reference from the New Testament to the Hebrew Bible, establishing a connection between the drying up of the Euphrates River and Moses parting the Red Sea. Both are instances of divine intervention in the context of a body of water. However, intertextuality here not only establishes a semantic parallel between two events, but it also emphasizes the difference. The passage from Exodus is moment of the divine judgment that leads to safety, whereas the drying up of the Euphrates is a preparation for the final judgment of the world.

etry (Forstall and Scheirer, 2019). Several of these works, Burns et al. (2021) in particular, highlight that neural embeddings can be used effectively to capture intertextuality. However, not much attention has been paid to the *effects of translation* on intertextual references.

In this work, we look at translation effects on intertextuality in the Bible through neural embedding spaces. While the Bible is often treated as one text, it is in fact a library of texts written by an estimated 60 different authors over the course of 4,000 years, and therefore offers a unique test bed for the detection of intertextuality and the effects of translation. This is especially true given the multilingual nature of the intertextuality between the New Testament and the Hebrew Bible in their original Greek and Hebrew.

Our main contributions are as follows:

1. We show that multilingual embedding spaces may be effectively used to characterize intertextuality in original documents as well as their translations.
2. We provide a new method for characterizing intertextuality within and across translations.
3. We conduct a comparative study of human- and machine-generated translations of the same corpus into different languages of varying resource levels.
4. We contribute to Classical and Biblical scholarship that qualitatively explores whether human translations have, purposefully or not,

amplified intertextuality between the old and new testaments for the sake of continuity².

2 Characterizing intertextuality between Corpora

Our intertextuality measure is simply the cosine similarity of a pair of verse embeddings from a multilingual embedding model. For a given set of ground-truth references, we can also compute *baseline* similarities by randomly swapping one of the verses with another from the same chapter³. The ratio of the average intertextuality similarity to the average baseline similarity can be used to compare the degree of intertextuality across different sets of translations.

Intuitively, a ratio much larger than one (1) indicates strong intertextuality, whereas anything less than one indicates that supposedly intertextual verses are not more similar than random pairings. When comparing changes in intertextuality ratio across translation, we compute the 95% confidence interval via bootstrapping. Specifically, we resample the original data with replacement 10,000 times, recalculating the ratio for each resample.

Note that this method relies upon having access

²For instance, Erich Auerbach underlines that Paul’s historical mission among the Gentiles needed to separate Christianity from Judaism by conveying the idea that “the old Law is suspended and replaced” through references that both alluded to and recontextualized the Hebrew Bible (Auerbach, 1959) (Sirin, 2022).

³Maintaining the same chapter ensures that false pairs likely remain upon the same topic, as opposed to choosing a random verse from anywhere in the Bible.

Language	Family	Bitext pairs
English	West Germanic	> 10M
Finnish	Uralic	> 1M
Turkish	Turkic	> 100K
Swedish	North Germanic	> 10K
Marathi	Indo-Aryan	Small

Table 2: **Languages by family.** Summary of languages used in this study: each has a full, aligned human translation of both the Jewish and Christian texts. The sizes are reported from Tang et al. (2020) training data and reflect the variety of resource-levels.

to ground-truth references — or suspected references — and would likely be too crude a method to discover novel instances of intertextuality without extensive threshold tuning. Instead, we use this measure to ascertain the *degree* of intertextuality within a set of texts known to be intertextual. We can then use this measurement to characterize changes in intertextuality across the same set of texts in translation.

We compute intertextuality ratios for all original, human, and machine-translation texts, distinguishing the sets of references that are internal to a testament (*within*) and that cross between them (*across*). This distinction allows us to consider whether Christian writing is particularly referential to the Jewish Testament, or if it became so through the effects of translation. Christian theologians throughout history have often underscored the continuity of the Christian and Jewish testaments (van der Waal, 1980), and human translators may have sought to emphasize this continuity in their translations. The full tables of these ratios can be found in Table 4.

3 Method

3.1 Data

We use three primary sources for our textual analysis: the Translator’s Amalgamated Hebrew Old Testament (TAHOT) and Greek New Testament (TAGNT)⁴, as well as a digitized copy of the Septuagint (LXX)⁵. The TAHOT is based on the Leningrad Codex, the oldest complete extant version of the Hebrew Old Testament. The TAGNT consolidates the Greek New Testament text from multiple early extant editions, and these are both compiled by Bible scholars at Tyndale House in Cambridge, UK, and released as part of the STEP

⁴www.STEPBible.org

⁵<https://sourceforge.net/projects/zefania-sharp/files/Bibles/GRC>

Target	Source Manuscript		
	Hebrew OT	Greek OT	Greek NT
English	69.5	61.2	72.6
Finnish	47.6	43.9	48.8
Turkish	66.7	65.4	68.2
Swedish	54.0	53.7	56.0
Marathi	27.6	26.5	29.8

Table 3: **COMET scores.** Top-scoring translation for each source manuscript is in bold text. Second top-scoring translation is in italics.

Bible project⁶. The Septuagint is the earliest Greek translation of the Hebrew Old Testament, completed by Jewish scribes in the few centuries preceding the events of the New Testament.⁷

For modern human translations, we use the Johns Hopkins University Bible Corpus (McCarthy et al., 2020) for the five languages in Table 2, each of which include both testaments.

To independently evaluate our method, we use a benchmark corpus for intertextuality provided by Burns et al. (2021) detailing intertextual references in Classical Latin literature. Specifically, it contains 945 references curated by subject matter experts connecting Valerius Flaccus’ *Argonautica I* to earlier and contemporary Roman authors.

3.2 Translation

To compare the effects of human and machine translation, we employ Cohere’s multilingual model Aya23⁸ (Aryabumi et al., 2024) to translate all of the original manuscripts into the five languages of varying resource levels from Table 2. Aya23 is chosen for this task as it has been shown to outperform other multilingual models of similar, and sometimes larger, sizes for machine translation (Aryabumi et al., 2024), but is small enough to be practical for academic research settings with limited compute power. Full pre-processing, prompting, and post-processing details may be found in § 6. We report translation quality scores using the COMET metric (Rei et al., 2020) in Table 3, providing references (human-translated text in the target

⁶We release a formatted version of STEP Bible’s data on the Huggingface Hub. DOIs: [10.57967/hf/4174](https://doi.org/10.57967/hf/4174), [10.57967/hf/4184](https://doi.org/10.57967/hf/4184).

⁷The complex history of Biblical scribal tradition means that almost all modern English translations use a versification system which at many points differs from the versification in the original Hebrew (cf. Genesis 31:55 in English translations is considered Genesis 32:1 in the Leningrad Codex). For consistency, we align all documents to use the English versification system across all experiments.

⁸We use the model version with 8B parameters.

language), predictions (machine-generated text in the target language), and sources (original text in the original language).

3.3 Gold standard for intertextuality

For ground-truth information about which passages are truly interlinked, we use a dataset of Bible cross-references (Owens, 2023). According to the dataset’s documentation, the initial data was seeded largely from the Treasury of Scripture Knowledge (Torrey and Canne, 1982), an authoritative compilation of cross-references from prominent Biblical scholars over many centuries, which was then cleaned to remove duplicates and concatenate separate entries for adjacent references. Finally, the references were opened to crowd-sourcing annotation for voting on relevant connections.

We limit consideration to verse-to-verse links that connect passages from different books and can be resolved in all manuscripts. We disregard ordering by summing the votes for both directions between a pair of verses, and use a vote threshold of 50 to consider a reference valid.⁹ This produces a total of 2183 references: 548 are entirely within the Jewish testament, 961 within the Christian, and 674 that span them. We differentiate these two cases with the qualifiers *within*, meaning within the same testament, and *across*, meaning across the two testaments.

		Within		Across
		Jewish (OT)	Christian (NT)	
Orig.	Ancient Hebrew	0.98 ± 0.14	–	–
	Ancient Greek	1.27 ± 0.21	1.30 ± 0.19	1.31 ± 0.20
Human	English	1.66 ± 0.21	1.70 ± 0.30	1.69 ± 0.27
	Finnish	1.42 ± 0.41	1.36 ± 0.53	1.48 ± 0.22
	Turkish	1.50 ± 0.18	1.43 ± 0.29	1.51 ± 0.12
	Swedish	1.33 ± 0.15	1.39 ± 0.12	1.37 ± 0.09
	Marathi	1.35 ± 0.12	1.42 ± 0.12	1.44 ± 0.10
NMT	English	1.32 ± 0.20	1.50 ± 0.17	1.48 ± 0.20
	Finnish	1.24 ± 0.22	1.28 ± 0.18	1.26 ± 0.11
	Turkish	1.60 ± 0.15	1.71 ± 0.12	1.52 ± 0.32
	Swedish	1.31 ± 0.22	1.29 ± 0.31	1.36 ± 0.30
	Marathi	1.02 ± 0.10	1.22 ± 0.25	1.30 ± 0.18

Table 4: Intertextuality ratios for source manuscripts and their human translations. Ratios within, and where possible between, testaments, for the Septuagint and TAGNT (Greek), TAHOT (Hebrew), and five human translations with 95% CI.

4 Experiments

Benchmark Corpus: First, we evaluate our method on a benchmark corpus for intertextuality between Valerius Flaccus’ *Argonautica I* to earlier

⁹We independently verify that 96.0% of the cross-references in our dataset with at least 50 votes are attested in an online version of the Treasury of Scripture Knowledge <https://www.tsk-online.com/>.

a	ὅτι ἴλεως ἔσομαι ταῖς ἀδικίαις αὐτῶν, καὶ τῶν ἁμαρτιῶν αὐτῶν οὐ μὴ μνησθῶ ἔτι.
b	ἐγὼ εἰμι ἐγὼ εἰμι ὁ ἐξολείψων τὰς ἀνομίας σου καὶ οὐ μὴ μνησθῆσομαι.
a	For I will be merciful to their unrighteousness, and their sins and their iniquities will I remember no more.
b	I, even I, am he that blotteth out thy transgressions for mine own sake, and will not remember thy sins.
a'	For I will beware of their iniquity, and their sinner’s iniquity; for I will not abhor them:
b'	I am the last of thy iniquitous acts, and I hate not myself.

Table 5: Overemphasized Intertextuality by Human Translation. The intertextuality from Hebrews 8:12 to Isaiah 43:25 is amplified by the human translator’s decision to render different words as "sin". The machine translation abstains from this and restores the original distance, but loses coherence.

and contemporary Roman authors. We calculate an intertextuality ratio of 1.55, 95% CI [1.53,1.56], indicating that our method succeeds at characterizing known intertextuality at the corpus level.

Translation Quality: Table 3 shows translation scores from the Hebrew Old Testament, Greek Old Testament, and Greek New Testament into five target languages. English and Turkish consistently achieve the highest scores across all manuscripts, with English translations ranging from 61.2 to 72.6, and Turkish from 65.4 to 68.2, suggesting strong translation quality for these language pairs. In contrast, translations into Marathi show the lowest scores, ranging from 26.5 to 29.8, likely due to the complexity of translating between less common language pairs. These results establish a valuable benchmark for evaluating translation quality for underrepresented languages in historical texts.

5 Analyzing Intertextuality in Translation

Table 4 shows that there is a higher degree of intertextuality *across* the New Testament and the Hebrew Bible compared to intertextual references *within* each book.

The degree to which intertextuality is preserved is highest for the English translation and lowest for Marathi. Human translations consistently show higher levels of intertextuality.

As suggested by McGovern et al. (2024), we indeed see that human translations over or under-emphasize the intertextuality present in the text, whereas machine translations provide a neutral baseline, based on these results. We can look closer at the translation effects by sorting intertextual

pairs according to the absolute shift in similarity. Table 5 shows the original Greek, the human English translation, and the unconstrained machine translation for one such pair, between the Epistle to the Hebrews and the Book of Isaiah. The pair of verses has strong similarity in the original Greek (0.332), but this is nearly doubled by the human English translation (0.656). The highlighted Greek word, *hamartion*, typically translated as *sin*, occurs in Hebrews but not Isaiah, yet the latter’s human translation makes a point of using the term. Surface-level lexical decisions like this, and presumably many less direct choices, lead to uncalibrated translations that reinforce the received interpretation.

6 Future work

We plan to address the persistent issue of misalignment in parallel Bible corpora. Even in scholarly editions of digitized texts, misalignment is persistent. However, by applying the alignment methodology proposed by (Craig et al., 2023), we could unify alignment for research purposes. Finally, we leave to future work exploring larger narrative contexts by examining narrative episodes instead of verse-level intertextuality.

Limitations

In this work, we generate machine translations working from the oldest extant manuscripts of the Biblical texts. However, most translations present in the JHUBC were not translated directly from ancient manuscripts but instead work from English translations, which themselves were often translations of the Greek texts. So direct comparisons of the human translations and machine translations in this work should be treated with caution.

Acknowledgments

We would like to thank Matt Post for very helpful discussions about machine translation evaluation, as well as scholars at Tyndale House, especially Ellie Wiener and Caleb J. Howard, for useful discussions about Biblical Hebrew. Thanks to Andrew Caines for his helpful comments on previous drafts. Hope McGovern’s work is supported by the Woolf Institute for Interfaith Relations and the Cambridge Trust.

References

- Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Kelly Marchisio, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick Frosst, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. [Aya 23: Open weight releases to further multilingual progress.](#)
- Erich Auerbach. 1959. *Scenes from the Drama of European Literature: six essays.* Meridian Books.
- David Bamman and Gregory R. Crane. 2008. [The logic and discovery of textual allusion.](#) In *In Proceedings of the 2008 LREC Workshop on Language Technology for Cultural Heritage Data.*
- Damien Broderick. 2017. Reading sf as a mega-text. *Science fiction criticism: an anthology of essential writings*, pages 139–48.
- Patrick J Burns, James A Brofos, Kyle Li, Pramit Chaudhuri, and Joseph P Dexter. 2021. Profiling of intertextuality in latin literature using word embeddings. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4900–4907.
- C Craig, K Goyal, G Crane, F Shamsian, and DA Smith. 2023. Testing the limits of neural sentence alignment models on classical greek and latin texts and translations. In *Proceedings http://ceur-ws.org ISSN 1613 0073.*
- Bruno Currie. 2019. [The iliad, the odyssey, and narratological intertextuality*.](#) *Symbolae Osloenses*, 93(1):157–188.
- Joseph P. Dexter, Theodore Katz, Nilesh Tripuraneni, Tathagata Dasgupta, Ajay Kannan, James A. Brofos, Jorge A. Bonilla Lopez, Lea A. Schroeder, Adriana Casarez, Maxim Rabinovich, Ayelet Haimson Lushkov, and Pramit Chaudhuri. 2017. [Quantitative criticism of literary relationships.](#) *Proceedings of the National Academy of Sciences of the United States of America*, 114(16):E3195–E3204.
- Christopher W. Forstall and Walter J. Scheirer. 2019. [Quantitative Intertextuality - Analyzing the Markers of Information Reuse.](#) Springer.
- R.B. Hays. 1989. *Echoes of Scripture in the Letters of Paul.* Yale University Press.
- Julia Kristeva. 1986 [1969]. *Word, dialogue and novel, in: T Moi (ed), The Kristeva Reader.* Columbia University Press, New York.
- John Lee. 2007. A Computational Model of Text Reuse in Ancient Literary Texts. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 472–479, Prague, Czech Republic. Association for Computational Linguistics.

- Arya D. McCarthy, Rachel Wicks, Dylan Lewis, Aaron Mueller, Winston Wu, Oliver Adams, Garrett Nicolai, Matt Post, and David Yarowsky. 2020. The Johns Hopkins University Bible Corpus: 1600+ Tongues for Typological Exploration. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2884–2892, Marseille, France. European Language Resources Association.
- Hope McGovern, Hale Sirin, Tom Lippincott, and Andrew Caines. 2024. [Detecting narrative patterns in biblical Hebrew and Greek](#). In *Proceedings of the 1st Workshop on Machine Learning for Ancient Languages (ML4AL 2024)*, pages 269–279, Hybrid in Bangkok, Thailand and online. Association for Computational Linguistics.
- Maria Moritz, Andreas Wiederhold, Barbara Pavlek, Yuri Bizzoni, and Marco Büchler. 2016. [Non-Literal Text Reuse in Historical Texts: An Approach to Identify Reuse Transformations and its Application to Bible Reuse](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1849–1859, Austin, Texas. Association for Computational Linguistics.
- Steve Moyise. 2002. [Intertextuality and biblical studies: A review](#). *Verbum et Ecclesia*, 23.
- Conley Owens. 2023. [Bible cross references](#).
- Andrew Piper, Richard Jean So, and David Bamman. 2021. [Narrative theory for computational narrative understanding](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Yisi Sang, Xiangyang Mou, Jing Li, Jeffrey Stanton, and Mo Yu. 2022. [A survey of machine narrative reading comprehension assessments](#). In *International Joint Conference on Artificial Intelligence*.
- Hale Sirin. 2022. *The Art of Scholarship: Auerbach, Tanpinar, and the Idea of Literary Knowledge*. Ph.D. thesis, Johns Hopkins University.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#).
- R.A. Torrey and J. Canne. 1982. *The Treasury of Scripture Knowledge: Five-hundred Thousand Scripture References and Parallel Passages*. Hendrickson Publishers Marketing, LLC.
- C van der Waal. 1980. [The continuity between the old and new testaments](#). *Neotestamentica*, 14:1–20.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).

A Additional Implementation Details

Preprocessing We use CohereForAI’s Aya-23 8B model to generate all machine translations. We do not remove any accents or diacritics as preprocessing.

Prompting We use few-shot prompting to obtain our translations. an example prompt can be seen below:

“Translate the following Ancient Greek phrases into English:

1. Ancient Greek: “εἰ δέ τις ἐποικοδομεῖ ἐπὶ τὸν θεμέλιον τοῦτον χρυσόν, ἄργυρον, λίθους τιμίους, ξύλα, χόρτον, κα λάμην,”

English: “Now if any man build upon this foundation gold, silver, precious stones, wood, hay, stubble;”

2. Ancient Greek: “καὶ οὐθὲν διέκρινεν μεταξύ ἡμῶν τε καὶ αὐτῶν τῇ πίστει καθαρίσας τὰς καρδίας αὐτῶν.”

English: “And put no difference between us and them, purifying their hearts by faith.”

3. Ancient Greek: “εἰ δὲ Χριστὸς οὐκ ἐγήγερται, κενὸν ἄρα καὶ τὸ κήρυγμα ἡμῶν, κενὴ δὲ καὶ ἡ πίστις ὑμῶν.”

English: “And if Christ be not risen, then is our preaching vain, and your faith is also vain.”

4. Ancient Greek: “καὶ ἐν τούτῳ γνωσόμεθα ὅτι ἐκ τῆς ἀληθείας ἐσμέν καὶ ἔμπροσθεν αὐτοῦ πείσομεν τὴν καρδίαν ἡμῶν”

English: “And hereby we know that we are of the truth, and shall assure our hearts before him.”

Now, translate this Ancient Greek phrase:

5. Ancient Greek: “INPUT_TEXT”

English:”

For the prompt, we draw four (4) examples of translations from the source texts. Ideally, these translations would be drawn from other parallel sources, but for most of the translation pairs (e.g. Ancient Hebrew → Marathi), the Biblical texts are the only parallel data available.

Generation At inference time, we use a maximum output length of 100 new tokens. We use the default BPE tokenizer with all of the default settings.

Post-Processing We find that we need to post-process the outputs: we grab what is in the first set of quotation marks after our prompt and exclude the rest. We find this is necessary to prevent nonsensical continued generations.

N. B. Models were access through the Huggingface Transformers library (Wolf et al., 2020).