

Larth: Dataset and Machine Translation for Etruscan

Gianluca Vico and Gerasimos Spanakis

Department of Advanced Computing Sciences / Paul-Henri Spaaklaan 1

Maastricht University / Maastricht, The Netherlands

g.vico@student.maastrichtuniversity.nl

jerry.spanakis@maastrichtuniversity.nl

Abstract

Etruscan is an ancient language spoken in Italy from the 7th century BC to the 1st century AD. There are no native speakers of the language at the present day, and its resources are scarce, as there exist only around 12,000 known inscriptions. To the best of our knowledge, there are no publicly available Etruscan corpora for natural language processing. Therefore, we propose a dataset for machine translation from Etruscan to English, which contains 2891 translated examples from existing academic sources. Some examples are extracted manually, while others are acquired in an automatic way. Along with the dataset, we benchmark different machine translation models observing that it is possible to achieve a BLEU score of 10.1 with a small transformer model. Releasing the dataset¹ can help enable future research on this language, similar languages or other languages with scarce resources.

1 Introduction

Etruscan (ISO 639-3 code: `ett`) is a language spoken in the Etruria region (modern-day centre Italy) from the 7th century BC to the 1st century AD (Wallace, 2008). It is written right to left using the Etruscan alphabet, derived from the Greek alphabet (Wallace, 2008). The predominant word order in this language is mostly subject-object-verb (Wallace, 2008). This pattern is similar to Latin, but distinguishing it from other languages like English, where the words follows the subject-verb-object order. It has 5 cases (accusative, nominative, genitive, dative and locative), two numbers (singular and plural) and takes into consideration animacy and gender (Wallace, 2008).

Only a small number of inscriptions in this language survived up to the present day: an estimated 12,000 inscriptions have been recovered (Wallace,

¹The data and code are available here: <https://github.com/GianlucaVico/Larth-Etruscan-NLP.git>

2008). However, only a few of them have a significant length to be considered complete. Other ancient languages used in similar areas and periods in history, such as Latin and Ancient Greek, have more resources, thus, making natural language processing techniques and tools easier to develop for these languages.

The contribution of this paper is threefold: First, we build a corpus of Etruscan inscriptions usable for natural language processing. We use as a starting point existing academic resources for this language exist, and we try to create our corpus both by manual and automatic work. Second, we focus on the machine translation task from Etruscan to English. We evaluate whether neural models can be trained with this data and if they can outperform less data-hungry models. Finally, we investigate if it is possible to exploit any similarity between Etruscan and Latin or Ancient Greek to improve the aforementioned model.

In Section 2, we introduce state-of-the-art techniques relevant to this paper. Then, in Sections 3 and 4 we explain the methods used to work on the data and the model used. Section 5 and Section 6 illustrate the experiments and compare the different techniques. Finally, Section 7 concludes the paper.

2 Literature review

The Etruscan Texts Project (ETP) (Wallace et al., 2004) is a digital Etruscan corpus which contains 369 inscriptions. The project is based on Etruskische Texte (Rix and Meiser, 1991) and is used in the book *Zihk Rasna* (Wallace, 2008). Another digital Etruscan work is the Corpus Inscriptionum Etruscarum Plenissimum (CIEP) (Hill, 2018), based on the Corpus Inscriptionum Etruscarum (CIE) (Pauli, 1893).

Similar works exist for Latin and Ancient Greek, like I.PHI (Sommerschild et al., 2021) and Perseus (Crane, 1985). In addition, toolkits like CLTK (Johnson et al., 2021) offer natural language pro-

Etruscan	ETP	Larth
A	a	a
B	b	b
C	c	c
D	d	d
⌘	e	e
⌚	v	v
I	z	z
⌘	h	h
⊗	θ	th
┆	i	i
K	k	k
┐	l	l
⌚	m	m
┆	n	n
⊕	+	s
⊕	s	s
○	o	o
⌚	σ, σ	s, sh
┆	p	p
⊙	q	q
⌘	r	r
⋈	s, ś, ζ, ζ	s, sh, s, sh
T	t	t
V	u	u
X	š	sh
ϕ	ϕ	ph
Υ	χ	kh
8	f	f

Table 1: Texts from ETP are already transliterated, but CIEP transliteration is sometimes ambiguous. We further reduce the number of symbols by using a subset of the Latin alphabet.

inscriptions can be damaged, so parts of the words cannot be read and the translation models have to either discard those words or rely only on the remaining characters. So, we generate more training samples by damaging more words. We assume that the damage occurs at the beginning or end of the words with a set probability. Also, we assume the number of damaged characters follows a geometric distribution. In this way, for instance, the word "clan" can stay unchanged or it might become "-an", "cla-", "-l-".

3.2 Latin and Ancient Greek

Models introduced later in the paper use Latin or Ancient Greek documents. Tatoeba eng-lat (Tiedemann, 2020) is used to train the Latin model. The text is normalised and non-Latin characters are removed. For Ancient Greek, we use Perseus (Crane,

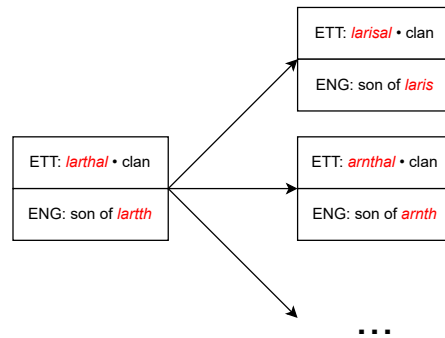


Figure 1: Example of data augmentation by replacing proper names. The name is replaced both in the Etruscan text and the English translation.

1985). In this case, we also remove all diacritical marks and transliterate the text to Latin. In this way, all the languages used share the same alphabet.

4 Machine Translation

We compare different models for machine translation on the BLEU metric but chr-F and TER metrics are also reported. The metrics are computed by SacreBLEU (Post, 2018). Higher BLEU and chr-F and lower TER indicate a better-performing model. Moreover, we evaluate the case where we use only ETP and ETP+CIEP for training and testing the models.

4.1 Random Model

The output of this model does not depend on the Etruscan inputs, but only on the training translations. It assumes that the length of the translations follows a normal distribution whose parameters are estimated from the training data. Then, it samples English tokens from the training distribution. The experiment is repeated 10 times with random splits of the dataset in training and testing data. The resulting metrics are then averaged.

4.2 Dictionary-based Model

The second model is a dictionary-based model based on the vocabulary provided in Zihk Rasna (Wallace, 2008). The model assumes that each word has one meaning and one translation. Moreover, it does not rearrange the word order and it does not consider the grammar of the source language or the target language. This model splits the input text into word tokens. Then, for each token, it searches for the exact match in the dictionary. If a

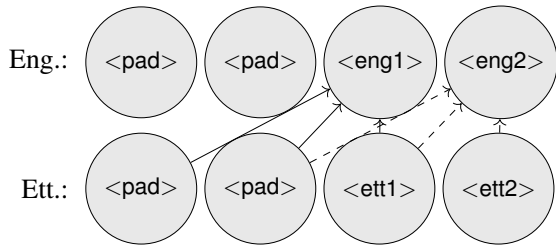


Figure 2: The first approach for the n-gram model. $\langle eng\ n \rangle$ indicates English tokens, while $\langle et\ n \rangle$ are Etruscan tokens; $\langle pad \rangle$ is the padding token. The example shows $P(\langle eng1 \rangle | \langle pad \rangle \langle pad \rangle \langle et1 \rangle)$ and $P(\langle eng2 \rangle | \langle pad \rangle \langle et1 \rangle \langle et2 \rangle)$. The context is made up of Etruscan trigrams.

match is found, it adds the translation to the output; otherwise, the token is ignored.

4.3 N-gram and Naïve Bayes Models

Then, we try to translate Etruscan taking into consideration the previous n tokens. The model estimates the probability distribution $\mathbb{P}(eng_i | ett_i, ett_{i-1}, \dots, ett_{i-n})$, where eng_i and ett_i are tokens at position i . This is done either directly from the training data or as a Naïve Bayes model with the following expression:

$$\begin{aligned} \mathbb{P}(eng_i | ett_i, ett_{i-1}, \dots, ett_{i-n}) &\propto \\ &\propto \mathbb{P}(eng_i) \prod_{j=0}^n \mathbb{P}(ett_{i-j} | eng_i) \quad (1) \end{aligned}$$

The model assumes that one n^{th} Etruscan token is translated into the single n^{th} English token. Figure 2 shows how the sequences are aligned and which Etruscan context is used for each English token.

A second N-gram model also includes the previous English tokens in the context by computing $\mathbb{P}(eng_i | ett_i, \dots, ett_{i-n}, eng_{i-1}, \dots, eng_{i-n-1})$ as shown in Figure 3. When the probability distribution is estimated directly, we consider the case when the word order is taken into account and when it is not. We use beam search to generate the output.

4.4 IBM Models

Next, we compare our models to existing ones. To do so, we consider the IBM models (Koehn, 2010) from the NLTK package (Bird and Loper, 2004). They are a series of 5 models with increasing complexity. These models consider the alignment between the source strings and the target strings, however, the Etruscan-English pairs we are using do

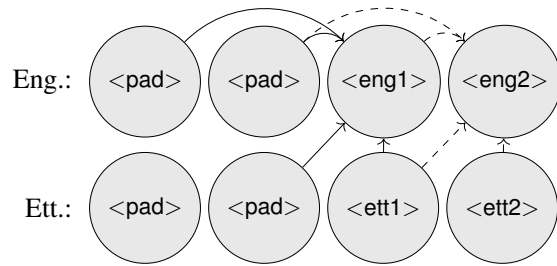


Figure 3: the second approach for the n-gram model. $\langle eng\ n \rangle$ indicates English tokens, while $\langle ett\ n \rangle$ are Etruscan tokens; $\langle pad \rangle$ is the padding token. The example shows $P(\langle eng1 \rangle | \langle pad \rangle \langle ett1 \rangle, \langle pad \rangle \langle pad \rangle)$ and $P(\langle eng2 \rangle | \langle ett1 \rangle \langle ett2 \rangle, \langle pad \rangle \langle eng1 \rangle)$. The context is made up of Etruscan and English bigrams.

not contain this information. Therefore, we test the models as if the sequences were aligned.

IBM1 does not consider the word order. IBM2 introduces the word order, while IBM3 takes also into consideration that a word can be translated into zero or more words. IBM4 and IBM5 can also reorder the output words. Moreover, IBM4 and IBM5 also need the part-of-speech (POS) tags of both the source and target sequences. POS tags are inferred from the grammatical features listed in the dictionary. For Etruscan, these are obtained by a manually annotated list of words, while the English sequences are tagged by NLTK perceptron tagger.

4.5 Transformer Models - Larth

Finally, we propose a transformer model, **Larth**. The encoder is based on Ithaca (Assael et al., 2022). It takes both the characters and the words as input and concatenates their embeddings. Then, the sequence is encoded with a BigBird attention block (Zaheer et al., 2021). The character and word sequences are aligned so that they have the same length. To do so, we test two approaches: we either extend the word sequence by repeating the word tokens or by adding space tokens as shown in Figure 4.

The decoder uses the encoded and the translated word sequences as input. First, it applies self-attention to the translated sequence, and then it computes the cross-attention between the translation and the encoded inputs. A feed-forward layer generates the output. Figure 5 illustrates this architecture.

First, we train the model from scratch on Etruscan \rightarrow English. Then, the model is initially trained for Latin \rightarrow English or Ancient Greek \rightarrow

Repeated word tokens										
Char:	<v>	<i>	<n>	<u>	<m>	<_>	<t>	<h>	<i>	<c>
Word:	<vinum>	<vinum>	<vinum>	<vinum>	<vinum>	<_>	<thic>	<thic>	<thic>	<thic>

Space tokens										
Char:	<v>	<i>	<n>	<u>	<m>	<_>	<t>	<h>	<i>	<c>
Word:	<vinum>	<_>	<_>	<_>	<_>	<_>	<_>	<_>	<_>	<_>

Figure 4: Example of how the character and word sequence are aligned. The string *vinum thic* means *wine and water*.

Dataset	BLEU	chr-F	TER
ETP+	0.059	9.263	194.977
CIEP	(0.0174)	(0.295)	(10.676)
ETP	0.324	13.970	133.878
	(0.064)	(1.150)	(11.877)

Table 2: Performance of the random model on the different Etruscan datasets. The table reports the mean value and the standard deviation of the metrics.

English and later fine-tuned on the original task Etruscan \rightarrow English.

Moreover, we investigate the effect of using both the character and the word sequence by training with only one of the sequences and the effect of data augmentation. The model uses beam search when generating the output sequences, but we use one beam when evaluating during the training for efficiency. Sequences are truncated at 256 tokens due to memory and computational resources.

5 Experiments

In this section, we compare different machine translation models trained on Etruscan data. The models are compared on the BLEU score.

5.1 Random Model

First, we run the random model on the Etruscan-English data. The dataset is split into 80 % for training and 20 % for testing. Only English labels are used for the training. Each experiment is repeated 10 times with random dataset splits. Table 2 reports the mean scores and the standard deviation of the models with different combinations of the datasets.

5.2 Dictionary-based Models

From the book Zikh Rasna it is possible to build a dictionary containing 821 vocables and their trans-

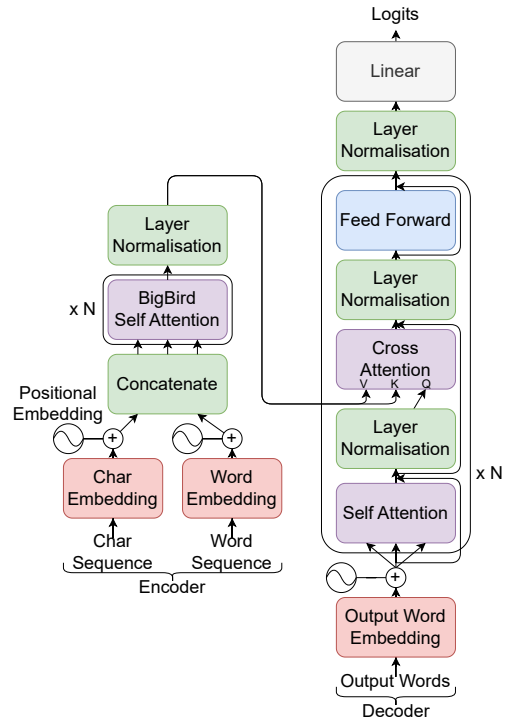


Figure 5: Transformer architecture used to translate Etruscan to English. The encoder imitates Ithaca’s torso. For both the encoder and the decoder, one attention block is used.

Dataset	BLEU	chr-F	TER
ETP+CIEP	0.167	9.120	89.799
ETP	4.505	40.771	68.135
CIEP	0.000	1.896	98.672
ETP (Suffix)	1.605	37.669	82.666

Table 3: Results of the dictionary-based model when tested on the different sets. *ETP (Suffix)* is the model tested on ETP with the suffix tokenizer.

lations.

We compare two tokenizers for Etruscan: the first uses white spaces to split the tokens, and the second also separates the suffixes from the root. The list of suffixes is also obtained from Zikh Rasna and the tokenizer recognises 178 suffixes. Table 3 shows the results of this model when translating Etruscan.

If we consider the example *itun turuce venel atelinas tinas dlniiaras* with the reference translation *venel atelinas dedicated this vase to the sons of tinia*, this model predicts *this dedicated venel atelina tinia*. If we use the suffix tokenizer the prediction is *this for him dedicated three this venel laris atelina shows*.

Context: ETT - Word order: No			
N-gram	BLEU	chr-F	TER
1	0.406 (0.163)	7.727 (0.867)	92.605 (0.960)
2	0.006 (0.001)	3.249 (0.752)	98.035 (0.821)
3	0.001 (0.001)	2.523 (0.753)	98.553 (0.821)
Context: ETT - Word order: Yes			
N-gram	BLEU	chr-F	TER
1	0.405 (0.163)	7.727 (0.867)	92.605 (0.960)
2	0.005 (0.005)	3.211 (1.004)	98.013 (1.089)
3	0.001 (0.001)	2.523 (0.748)	98.531 (0.870)

Table 4: Mean scores and their standard deviation (in parenthesis) of the n-gram models that use only the Etruscan texts.

5.3 N-gram and Naive Bayes models

Similarly to the random models, 80% of the data is used for training, while the remaining 20% is for testing. The dataset is ETP. Each experiment is repeated 10 times with different random splits.

With the N-gram models, we compare models with a context size of 1, 2 and 3 that use only Etruscan or both Etruscan and English as context and whether they consider the word order. Out-of-vocabulary (OOV) tokens are handled with additive smoothing. We use 8 beams when generating the output sequence, however, this is equivalent to greedy search when the context uses only Etruscan. Table 4 shows the results of the models that use only the Etruscan sequence, while Table 5 shows the models that also use the English translations.

For the Naive Bayes models, we only use a context size of 2 and 3, and the models always consider the word order. Table 6 reports the results.

5.4 IBM models

We split 80 % of the data for training and 20 % for testing. Moreover, we use the previously built dictionary as training data. No alignment information is given to the model, but IBM4 and IBM5 receive a dictionary that maps words to POS tags. We assume that words can only have one tag.

IBM3, IBM4, and IBM5 are trained only with the dictionary data. Models trained on ETP+CIEP are tested on ETP+CIEP, while models trained on

Context: ETT-ENG - Word order: No			
N-gram	BLEU	chr-F	TER
1	0.218 (0.018)	3.059 (0.301)	92.902 (1.160)
2	0 (0)	0 (0)	100 (0)
3	0 (0)	0 (0)	100 (0)
Context: ETT-ENG - Word order: Yes			
N-gram	BLEU	chr-F	TER
1	0.447 (0.211)	5.360 (0.856)	92.105 (1.117)
2	0.000 (0.000)	0.370 (0.167)	99.705 (0.346)
3	0.000 (0.000)	0.357 (0.097)	99.690 (0.297)

Table 5: Mean scores and their standard deviation (in parenthesis) of the n-gram models that use the Etruscan texts and the English translations. When the scores are zero is because the models immediately predict the end-of-sequence (EOS) token.

N	Context	BLEU	chr-F	TER
2	Ett.	0.160 (0.023)	12.609 (1.009)	101.482 (1.251)
3	Ett.	0.146 (0.030)	12.708 (0.921)	103.867 (1.220)
2	Ett.-Eng.	0.055 (0.048)	9.547 (1.821)	101.522 (0.851)
3	Ett.-Eng.	0.055 (0.048)	9.954 (2.103)	103.038 (1.005)

Table 6: Mean scores and their standard deviation (in parenthesis) of the Naïve Bayes models.

ETP are tested on ETP as shown in Tables 7 and 8.

As an example, IBM3 translates "*eca shuthic velus ezipus clensi cerine*" as "*this funerary vel et-spus son constructed*", while the reference translation is "*this funerary monument belongs to vel etspu it is constructed by his son*".

5.5 Transformer Models - Larth

The model is trained for Etruscan \rightarrow English translation with ETP+CIEP and with ETP only. The models are tested on the same split of the dataset. Due to the small size of the dataset, 95 % of the data is used for training.

The optimizer is RAdam (Liu et al., 2019), with an initial learning rate of 0.002 and 250 warmup steps. We use a reverse square root learning sched-

ETP+CIEP			
Model	BLEU	chr-F	TER
IBM1	0.402	19.744	89.213
	(0.183)	(1.178)	(0.693)
IBM2	0.392	19.450	89.551
	(0.183)	(1.383)	(0.487)
IBM3(*)	0.105	8.629	91.052
	(0.046)	(1.148)	(1.507)
IBM4(*)	0.105	8.627	91.052
	(0.046)	(1.148)	(1.507)
IBM5(*)	0.105	8.631	91.063
	(0.046)	(1.147)	(1.516)

Table 7: Performance of the IBM models on the ETP+CIEP dataset. (*): IBM3, IBM4 and IBM5 are trained only with the dictionary.

ETP			
Model	BLEU	chr-F	TER
IBM1	2.187	37.363	73.917
	(0.596)	(2.011)	(2.163)
IBM2	2.104	36.721	74.334
	(0.449)	(2.098)	(2.090)
IBM3(*)	2.482	39.393	71.270
	(0.513)	(2.229)	(2.456)
IBM4(*)	2.482	39.391	71.270
	(0.514)	(2.228)	(2.456)
IBM5(*)	2.481	39.416	71.331
	(0.513)	(2.235)	(2.415)

Table 8: Performance of the IBM models on the ETP dataset. (*): IBM3, IBM4 and IBM5 are trained only with the dictionary.

ule. The loss function is cross-entropy, and the batch size is 32. We set the label smoothing to 0.1.

We first try to train from scratch and with different alignment techniques. The BLEU, chr-F and TER scores are shown in Table 9. We use data augmentation with ETP+CIEP with the sequences aligned by repeating the word tokens, however, we do not use it on ETP due to the decrease in performance.

Next, we train the same architecture with only the word sequence or only the character sequence. The results are shown in Table 10.

When training the same model with the Latin and Greek data, it achieved, respectively, BLEU/chr-F/TER of 0.4968/5.01/151.4 and 0.12/6.186/107.3. Then, we fine-tune those models with Etruscan as shown in Table 11.

Larth trained on ETP translates "*mi aveles me-*

ETP+CIEP			
Model	BLEU	chr-F	TER
repeat	10.1	15.11	144.5
space	5.201	16.9	274.8
repeat+unk	2.8	14.8	189.1
repeat+name	1.004	12.2	615.9
ETP			
Model	BLEU	chr-F	TER
repeat	9.053	17.24	137
space	5.784	15.88	124.7

Table 9: Larth trained from scratch for Etruscan \rightarrow English. *Repeat* and *space* indicate how the character and the word sequence are aligned. *+name* is trained with data augmented by changing names, while *+unk* is augmented by deleting characters.

ETP+CIEP			
Inputs	BLEU	chr-F	TER
char	0.9694	14.42	254.8
word	2.776	13.49	99.88
char+word	10.1	15.11	144.5
ETP			
Inputs	BLEU	chr-F	TER
char	0.1431	11.22	528.1
word	7.679	18.48	131.6
char+word	9.053	17.24	137

Table 10: Larth trained from scratch for Etruscan \rightarrow English with only the character or the word sequence or both as input.

tienas" as "*i am the tomb*" while the reference translation is "*i am the tomb of avele metienas*". Note that in this example "*the tomb*" is implied and not mentioned explicitly.

When trained on ETP+CIEP, we have "*e ca shuthi anes cuclnies*" translated as "*this tomb*" but the reference is "*this is the tomb of ane cuclnies*". In this case "*the tomb*" is mentioned, but the model misses the name of the owner, which is also mentioned.

6 Results & Discussion

Figure 6 and Figure 7 compare the scores of the models presented in the previous Section. Compared to the random model, the dictionary-based model shows higher BLEU and chr-F scores and lower TER scores except when tested only on CIEP. This suggests that CIEP is noisier than ETP and that the dictionary is not suited for CIEP.

The N-gram models perform better than random

Data	BLEU	chr-F	TER
Lat+ETP+CIEP	0.1965	2.195	351.6
Grc+ETP+CIEP	1.011	8.148	215.3
Lat+ETP	0.293	3.784	654.4
Grc+ETP	2.037	6.04	164

Table 11: Larth trained with Latin (Lat) or Ancient Greek (Grc) and then fine-tuned on Etruscan.

only when using unigrams as context. With longer n-grams, the performance decrease until the model only predicts the EOS token. We can make similar observations for Naïve Bayes models.

IBM models are able to perform better than random. When trained on ETP+CIEP, simpler models work better. This, again, might depend on the noise in CIEP. IBM3 works better on ETP despite being trained only with the dictionary. Adding POS information (IBM4 and IBM5) does not improve the results. However, on ETP the dictionary-based model still performs better than the IBM models.

Larth is able to achieve a better BLEU score than the previous models on both ETP and ETP+CIEP. However, it needs to use both the character and word sequences and the word tokens are repeated to align the two sequences, whereas the other models only use the word tokens. Using the space token to align the sequences decrease the performance, but the BLEU score is still higher than the dictionary-based model. A similar observation can be made for the model using only the word sequence. Using data augmentation or only the character sequences reduces the performance that is still higher than random.

Fine-tuning from Latin and Ancient Greek always performs worse than the dictionary-based model. This may depend on the small size of the model that is not able to adapt.

As for chr-F and TER, the dictionary model and IBM models perform better than Larth. These two models can only output tokens from the training set and ignore unknown tokens. Thus, they can generate longer sequences of correct characters (high chr-F) and the errors are mainly for unknown tokens or from English tokens that are not directly present in the Etruscan texts like articles (low TER). Whereas, Larth uses tokens that can be word pieces and it still generates a translation for unknown tokens.

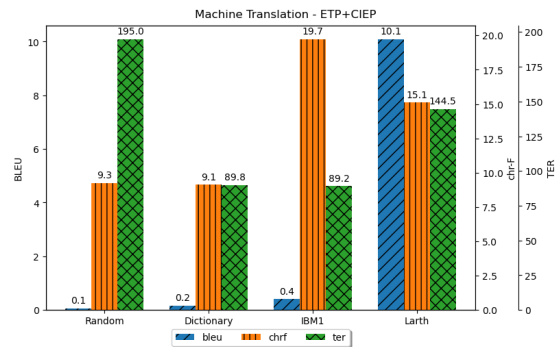


Figure 6: Comparison of the models with the best BLEU scores on ETP+CIEP. One model from each type is selected.

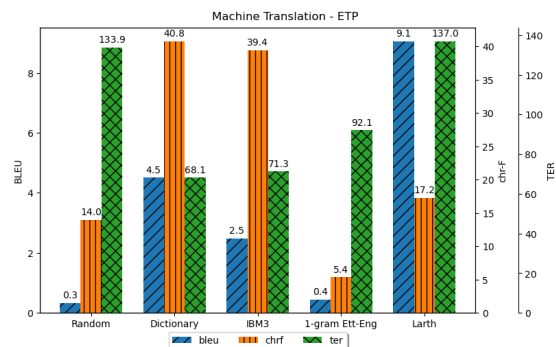


Figure 7: Comparison of the models with the best BLEU scores on ETP. One model from each type is selected.

7 Discussion

In this paper, we present a dataset for Etruscan \rightarrow English machine translation. Although the dataset is not very big, we show that it is possible to train statistical and transformer models. Given the unexplored nature of Etruscan language, the fact that trained models perform better than random is an important first step for this language. Moreover, we demonstrated that Larth performs better than the IBM models when trained on the available data.

However, our model does not provide any explanation about the generated translation neither it guarantees whether it is correct. Our model’s performance also depends on the dataset itself, which does not contain any bibliographic information or the reasoning that the original authors used to translate the inscriptions. Future work includes delivering a cleaner and more complete version of the dataset and the inclusion of additional metadata, such as bibliographic information, more accurate location, or interesting graphical information (e.g. the direction of the inscription).

References

- Yannis Assael, Thea Sommerschild, Brendan Shillingford, Mahyar Bordbar, John Pavlopoulos, Marita Chatzipanagiotou, Ion Androutsopoulos, Jonathan Prag, and Nando de Freitas. 2022. [Restoring and attributing ancient texts using deep neural networks](#). *Nature*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- David Bamman and Patrick J. Burns. 2020. [Latin BERT: A contextual language model for classical philology](#). *CoRR*, abs/2009.10053.
- Mani Bansal and D.K. Lobiya. 2020. [Word-character hybrid machine translation model](#). In *2020 8th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)*, pages 270–274.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Gregory R. Crane. 1985. [Perseus digital library](#).
- Jeff Hill. 2018. [Corpus inscriptionum etruscarum plenissimum](#).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation*, 9(8):1735–1780.
- Kyle P. Johnson, Patrick J. Burns, John Stewart, Todd Cook, Clément Besnier, and William J. B. Mattingly. 2021. [The Classical Language Toolkit: An NLP framework for pre-modern languages](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 20–29, Online. Association for Computational Linguistics.
- Philipp Koehn. 2010. *Statistical Machine Translation*, 1st edition. Cambridge University Press, United States.
- Wang Ling, Isabel Trancoso, Chris Dyer, and Alan W. Black. 2015. [Character-based neural machine translation](#). *CoRR*, abs/1511.04586.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. [On the variance of the adaptive learning rate and beyond](#). *CoRR*, abs/1908.03265.
- Jorj X. McKie and Ruikai Liu. 2016. [PyMuPDF](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: A method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA. Association for Computational Linguistics.
- Carl Pauli. 1893. *Corpus Inscriptionum Etruscarum*. J. A. Barth.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Helmut Rix and Gerhard Meiser. 1991. *Etruskische Texte*. A. G. Narr.
- Khaidzir Muhammad Shahih and Ayu Purwarianti. 2019. [Combining word and character vector representation on neural machine translation](#). In *2019 Fourth International Conference on Informatics and Computing (ICIC)*, pages 1–4.
- Matthew Snover, Bonnier Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of Association for Machine Translation in the Americas*.
- Thea Sommerschild, Yannis Assael, Brendan Shillingford, Mahyar Bordbar, John Pavlopoulos, Marita Chatzipanagiotou, Ion Androutsopoulos, Jonathan Prag, and Nando de Freitas. 2021. [I.PHI dataset: ancient greek inscriptions](#). <https://github.com/sommerschild/iphii>.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014a. [Sequence to sequence learning with neural networks](#). *CoRR*, abs/1409.3215.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014b. [Sequence to sequence learning with neural networks](#). *CoRR*, abs/1409.3215.
- Jörg Tiedemann. 2020. [The Tatoeba Translation Challenge – Realistic data sets for low resource and multilingual MT](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1174–1182, Online. Association for Computational Linguistics.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT — Building open translation services for the World](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Rex Wallace. 2008. *Zikh Rasna: A Manual of the Etruscan Language and Inscriptions*. Beech Stave Press.
- Rex Wallace, Michael Shamgochian, and James Patterson. 2004. [Etruscan texts project](#).
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2016. [A character-aware encoder for neural machine translation](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3063–3070, Osaka, Japan. The COLING 2016 Organizing Committee.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2021. [Big bird: Transformers for longer sequences](#).