



HAL
open science

Computational Pattern Recognition in Linear A

Jajwalya R Karajgikar, Amira Al-Khulaidy, Anamaria Berea

► **To cite this version:**

Jajwalya R Karajgikar, Amira Al-Khulaidy, Anamaria Berea. Computational Pattern Recognition in Linear A. In press. hal-03207615

HAL Id: hal-03207615

<https://hal.science/hal-03207615v1>

Preprint submitted on 26 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Computational Pattern Recognition in Linear A

Jajwalya R. Karajgikar^{1*}, Amira Al-Khulaidy², Anamaria Berea³

1 Computational Sciences, MS, George Mason University

2 Computational Social Sciences, PhD student, George Mason University

3 Computational and Data Sciences, Associate Professor, George Mason University

*Corresponding author: Jajwalya R. Karajgikar¹

Abstract

Linear A is an ancient Mycenaean Greek language that is as of yet undeciphered. Computational analysis of the symbols using natural language processing and data mining techniques can aid to better understand this logogram language, and to help us uncover interesting statistical and information-theoretic patterns even without deciphering the language. Using information theory and data science techniques, we show an exploratory analysis of these symbols. We use n-gram analysis to summarize and predict the possible symbols on lost stone tablets, while symbols are clustered by topic modeling as well as k-means, for comparison. The results show that there are universal features in this language that we can explore further with computational methods.

keywords

Linear A, Information Theory, NLP, Topic modeling, Clustering, Knowledge Mining, Humanities and Social Sciences/Literature, Humanities and Social Sciences/History

INTRODUCTION

The historical context of the ancient languages may be lost to time but some of the archaeological evidence found in the form of stone tablets (as in our case), contain information that can be considered for the possibility of computational study. In the case of undeciphered languages, such as Linear A, many researchers turned to new computational techniques such as machine learning and natural language processing with the aim to decode them. But fewer authors have approached these languages from the data science and information theory perspective, with the aim to discover statistical and informational patterns which, even in the absence of a decipherment, would help us understand better the universality of human language, both ancient and modern.

From the data science perspective, languages are repositories of information that hold quantitative value in addition to the meaning that we are assigning them from the human cultural perspective. Undeciphered languages represent a type of agnostic data that can help us design better, unbiased algorithms, particularly for natural language processing methods. Therefore, in this paper, our main goal is to extract meaningful emergent patterns of vocabulary from an independent primary source, Linear A.

British Archeologist Arthur Evans discovered the hieroglyphic seal-stones with successive types of script around 1893 along the eastern Mediterranean basin, which culminated in the dramatic excavation of the palace hill of Mycenaean Knossos on the Greek island of Crete (Evans, 1909). Dated from around 1400 BCE, the ancient Minoan civilization inscriptions are one of the earliest

¹jkrajgi@gmu.edu / jajwalyak@gmail.com

forms of writing discovered. Evans was the first to define the pictographic Cretan hieroglyphics, and the subsequent linear of classes A and B. The unique structure of the ancient palace site was such that Evans believed the possibility of tracing the whole evolution of the art of writing all the way up to the Neolithic times, thereby establishing its importance in historical linguistics and art history. The Linear A writing system was prevalent in the Bronze Age Minoan civilization from 1800 to 1450 BC. It was the primary script used in palace and religious writings of the Minoan civilization. The etchings on the clay tablets found in the excavation are called logogram symbols² and form the basis of our dataset that will be run through several analysis techniques before they are available for subject matter expertise' discernment and approval. This research work is an attempt at taking a step on the identification, simulation, and study of language families origins and evolution with constellations of computational methodology.

I RELATED WORKS

1.1 Literature Survey

While Linear A is as yet undeciphered, the Linear B syllabary has been computationally considered (Papavassiliou et al., 2020) via smoothed n-gram analysis, Hidden Markov Models, Bayesian classifiers, and Conditional Random Fields (CRF). This paper proposes to conduct an exploratory approach to conducting data mining on an undeciphered language. Linear B has been challenging to decipher due to reasons similar to Linear A's. There is a limited amount of data that can be used for processing since there are only so many archaeological remains from the historical evidence of the Linear A language, which makes it a greater challenge to decipher than Linear B. Furthermore, there is no real way of knowing if the output will be accurate without validation data or material.

A machine-learning system capable of deciphering lost languages (Luo et al., 2020) was demonstrated by having it decipher Linear B—the first time this has been done automatically. The approach used was very different from the standard machine translation techniques by using a novel neural decipherment approach for automatic decipherment. The outcome accuracy can be verified and validated against existing Linear B literary decipherment.

In the below Linear B tablet facsimile³ (see figure 1, right), the following transliterations⁴ are known: men (ideogram⁵ VIR), women (ideogram MUL), girls (ko-wa) and boys (ko-wo). This is available because the women and children of Knossos were identified by the occupational, geographical, or ethnic names. The DĀMOS or Database of Mycenaean at Oslo (Aurora, et al., 2013) is a corpus of complex architecture just for Linear B syllabary that is standardized by the TEI epigraphic standard and Leiden convention that contains 5900 documents of metadata, textual markup, annotation, and modularity in a relational database MySQL workbench. The purpose is to translate each lemma, disseminate knowledge of the language, and its irregularities.

²In a written language, a logogram or logograph is a written character that represents a word or morpheme.

³A facsimile is a copy or reproduction of an old book, manuscript, map, art print, or other item of historical value that is as true to the original source as possible.

⁴A transliteration is a transcription from one alphabet to another transcription, written text - something written, especially copied from one medium to another, as a typewritten version of dictation.

⁵An ideogram is a character or symbol representing an idea or a thing without expressing the pronunciation of a particular word or words for it, as in the traffic sign

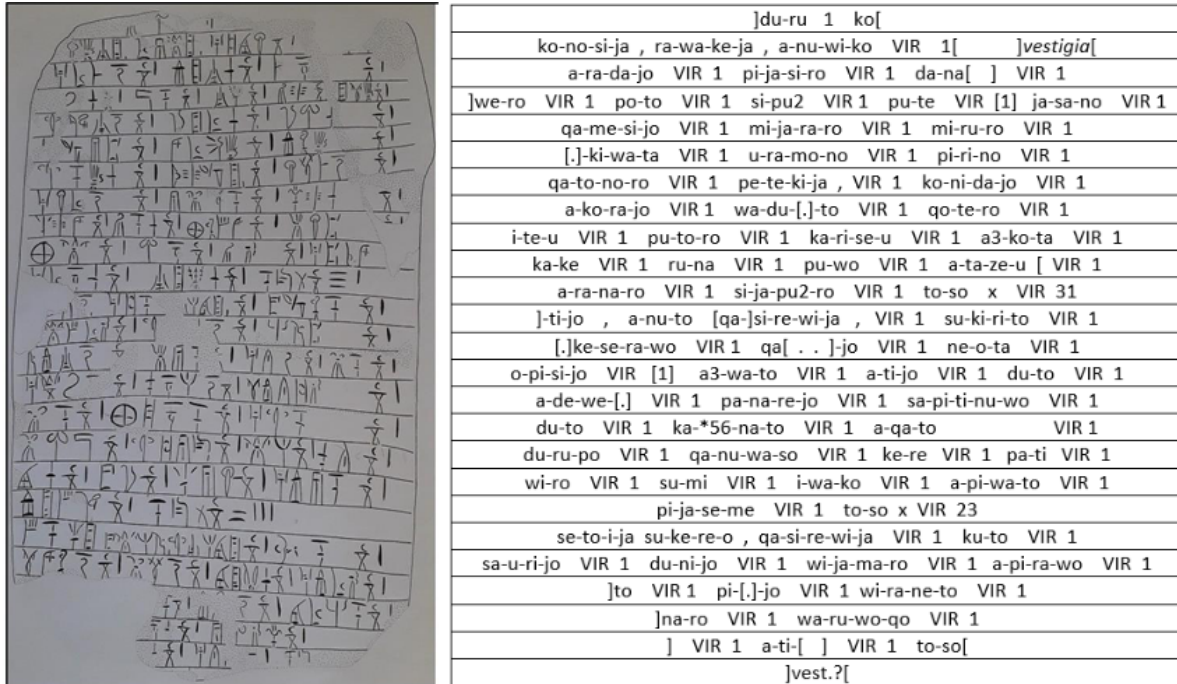


Figure 1: Left, the Mycenaean Tablet KN As 1516 (Chadwick, et al., 1987). Right, the transliterated text of the KN As 1516 (Chadwick, et al., 1987)

There are rich interdependencies between various unknown languages and their applicability in unraveling mysteries of modern evolutionary biology studies (Daggumati et al., 2019). Genomic data availability lies in the gray area of abundance and sparseness depending on the kind of algorithm at work. The construction of phylogenetic trees uses algorithms like Neighbor joining method, Maximum Likelihood, Common Mutations Similarity Matrix, which are comparatively objectively free of human bias (Daggumati et al., 2019). The ClustalW2 cladogram with 2 phylogenetic tree generation algorithms using both Neighbour Joining method and UPGMA method (unweighted pair group method with arithmetic mean) give the similarity index between various comparable language models.




Figure 2: Evolutionary Tree of ancient scripts reconstructed using Neighbor Joining algorithm in ClustalW2 (Daggumati et al., 2019)

As such, languages with the same proto-language structures having similar distribution profiles is not unknown to linguists. The sequence-to-sequence model capturing character-level correspondence between cognates using minimum cost flow problem, that is Long Short Term Memory algorithm. There are many more research papers to consider going forward but these provide a guideline to previous methodology and are indicative of the expected results.

1.2 Research Question(s)

Can we use natural language processing and data mining methods to create a meaningful taxonomy of vocabulary from a sparse dataset showing patterns of a language that is yet to be deciphered? To specify, the question is: What are the frequently occurring patterns in the Linear A corpus that are also similar in form?

1.2.1 Objectives

- Computational linguistic preprocessing and analysis via vector embeddings and similarity analysis via K-Means Clustering.
- Methods for the detection of intertexts and text patterns using Zipf Law and Entropy possibilities.

1.2.2 Motivation

We are interested in this particular problem for its applications in historical Computational Linguistics research. The multidisciplinary nature of the project is exciting for gaining knowledge in several domains in Digital Humanities. Furthermore, algorithms used with Natural Language Processing are applicable in a variety of fields like Evolutionary Biology Genomics. This is part of the wider project on studying the origins of socio-linguistic complexity combining natural language processing, information theory, data science techniques, and eventually even agent-based modeling.

II Model

2.1 Dataset

2.1.1 Preparing the Data

The Linear A data was available in JSON format in the Github repository ([mwenge / lineara.xyz](https://github.com/mwenge/lineara.xyz), 2019). The original .ods file was converted into a series of .csv extension files for easy machine readability, based on each attribute in the ODS extension file⁶. To read the inscriptions in Linear A and Linear B, special unicode font character installations are required: NotoEmojiRegular.ttf⁷, NotoMono-Regular.ttf⁸, NotoSansLinearA-LinearB.ttf⁹. When text is

⁶A file with the .ODS file extension is most likely an OpenDocument Spreadsheet file that contains spreadsheet information like text, charts, pictures, formulas and numbers, all placed within the confines of a sheet full of cells.

⁷ <https://github.com/googlefonts/noto-emoji>

⁸ <https://github.com/googlefonts/japanese/tree/main/build/fonts/noto>

⁹ <https://github.com/googlefonts/noto-fonts>

rendered by a computer, the characters in the text can not all be displayed, because no font that supports them is available to the computer. When this occurs, small boxes called “tofu” are shown to represent the characters. The Python package Matplotlib also does not render the font as of the current version 3.3.4.

2.1.2 Dataset Schema

Now the files can be opened to show the attribute features:
Name, Parsed Inscription, Transcription, Translated Words

#	source	name	normal copy	line	sign groups	ideograms	integers	fractions	probable “words”
0	1.0	8.052	AN Zb 1	None	None	None	None	None	None
1	2.0	4.006	AP Za <3>	None	None	None	None	None	None
2	3.0	4.002	AP Za 1	None	None	None	None	None	None
3	4.0	4.004	AP Za 2	1	None	None	None	None	None
4	5.0	4.004	AP Za 2	2	None	None	None	None	None
5	6.0	4.142	AR Zf 1	None	None	None	None	None	None
6	7.0	4.143	AR Zf 2	None	None	None	None	None	None
7	8.0	8.048	ARG Zg 1	None	None	None	None	None	None

Figure 3: Linear A Dataset from Github repository [mwenge](#) / [lineara.xyz](#)

2.1.3 Preprocessing the Data

The first step in the preprocessing pipeline is the fragmentation of the text blocks into sentences. However, in the case of our dataset, the values are in dataframe format and not in larger blocks of text. It is possible to consider the possible sentence boundaries depending on the transliterated pieces.

During the initial steps of the pre-processing pipeline, the content is fragmented into sentences. In numerous dialects and languages -- for example, English -- accentuation, particularly the full stop or period character, outcry, and question marks can be utilized to distinguish the finish of the sentence. However in the case of our dataset, we are unsure of the actual sentence boundaries since tablets are not indicative of word tokenization although it has been considered within the *probable “words”* column of the dataset. This is a major challenge in both processing the data and deciphering the language in general.

2.2 System Framework

2.2.1 Exploratory N-Gram Analysis

The purpose of this module is to explore the Linear A dataset for better understanding of the distribution of symbols in the corpus. We implemented different NLP techniques such as generating bigrams, trigrams, predicting symbols based on weighted frequency and generating a summary of the dataset. The diagram (see Figure 4) shows the process flow of this module. The process starts with importing required Python libraries nltk. The bigrams and trigrams collections

are collocations that help us learn of the summary and predicted symbols in the missing or lost stone tablets. Calculating the word weighted frequency generates a summary of the dataset and predicting symbols.




Figure 4: Natural Language Processing N-gram analysis process flow

The frequency distribution of the symbols gave us an interesting Lexical Dispersion plot¹⁰ (see Figure 5). The frequency of a word across the corresponding parts of the corpus highlights the word offsets.

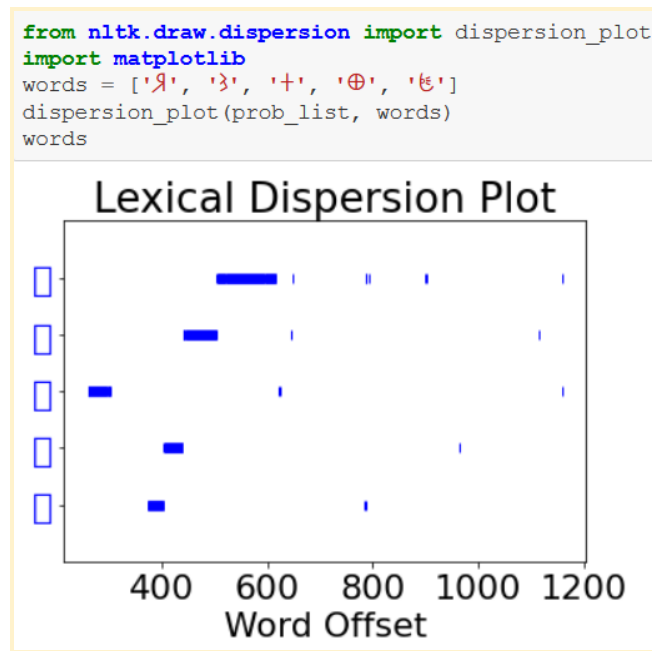



Figure 5: Lexical Dispersion plot for the most frequently occurring symbols

The plot of frequency of the terms is a short tail distribution, that is universal with Zipf's Law¹¹ parameters (see Figure 6). The terms considered for summary using n-gram analysis where (n-1) terms give us the nth terms, are thus the more highly frequent symbols through the relative frequency table. We find that, interestingly, the predicted symbols are also the ones found at the beginning of other tablets.

¹⁰Lexical dispersion is a measure of how frequently a word appears across the parts of a corpus. This plot notes the occurrences of a word and how many words from the beginning of the corpus it appears (word offsets).

¹¹Zipf's law is an empirical law formulated using mathematical statistics with related discrete power law probability distributions. The underlying principle in Zipf's law accepting models is that short sequences have a high probability of occurring while long sequences have a low probability and hence a high energy (Zipf, 1936).



FreqDist({'𐎠': 119, '𐎡': 66, '𐎢': 46, '𐎣': 38, '𐎤': 33, '𐎥': 31, '𐎦': 22, '𐎧': 15, '𐎨': 14, '𐎩': 12, ...})

Figure 6: Short-tailed distribution plot of the most frequently occurring symbols

Then we test the Zipf law by plotting first the frequency of the symbol ranked in descending order, then taking the logarithm of these values. Given that this dataset is allegedly 90% administrative records with billing records, the summary of the “goat” symbol along with fractions is an apt summary, we believe, for the dataset in question.

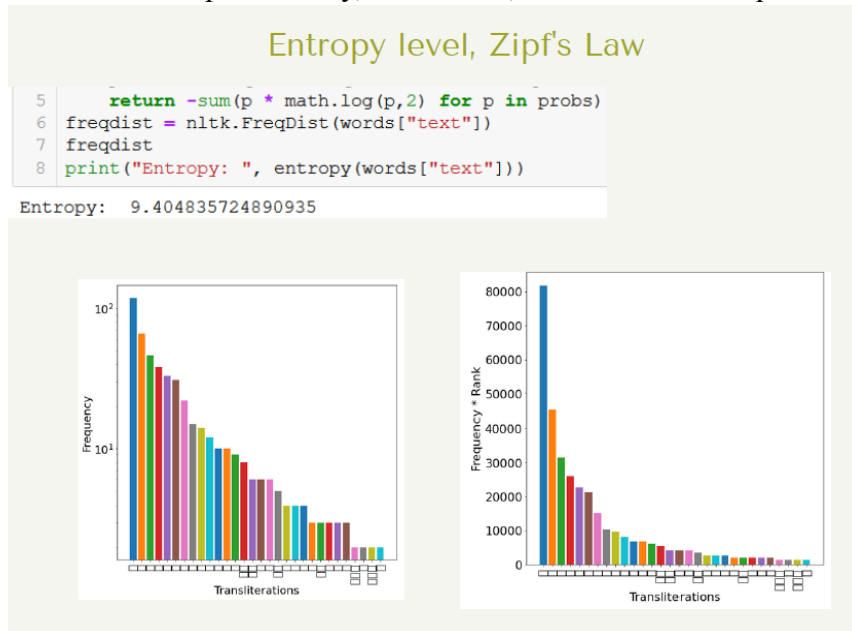


Figure 7: Plotting Zipf Law short-tailed distribution of probable words in the data set

For Entropy analysis, we expected the output value to be comparable to other known languages. The stone tablets containing only 96 unique symbols or transliterations transfers to the entropy¹² analysis where we gain the high entropy level of slightly above 9.4. The short tail distribution of the log of frequency and rank versus the transliterations is also very evident at each iteration.

¹²Entropy is the measure of uncertainty, randomness, or a disorder of a system (Clausius, 1865)

In this case, our analysis of entropy reveals a larger level of disorder in the Linear A language. If the disorder of the symbols is high, that could mean that it is not a language as we understand human languages to be, since most human languages have a similar entropy value (Shannon, 1949). The higher the Shannon entropy¹³ value, the greater the disorder and, subsequently, the less likely we are to think of that dataset as a language.

In this case, we obtained two results. The first result was obtained using only the “probable words” from the corpus, instead of the entire corpus. The output revealed a greater entropy of 9.404, than that of languages we know, which was inconclusive in terms of defining whether what we were looking at was a human language. However, when including more of the symbols and ideograms, the entropy decreased to 3, showing that the disorder had been greatly reduced (Matricciani, 1994).

One possible conclusion is that the introduction of more symbols in the dataset allowed the algorithm to recognize more patterns, which reinforces that we are looking at a human language from the information theory perspective (see Figure 8).




Figure 8: Entropy of natural languages and other sequences. (Rao, 2018)

We trained the Word2Vec model for one epoch and we fit the same model to look at the distribution of the symbols in the corpus, visualized with a scatter plot.




Figure 9: Word Embedding Process Flow

¹³Shannon Entropy formula where it is defined as “In information theory, the entropy of a random variable is the average level of "information", "surprise", or "uncertainty" inherent in the variable's possible outcomes.” (Shannon, 1949)

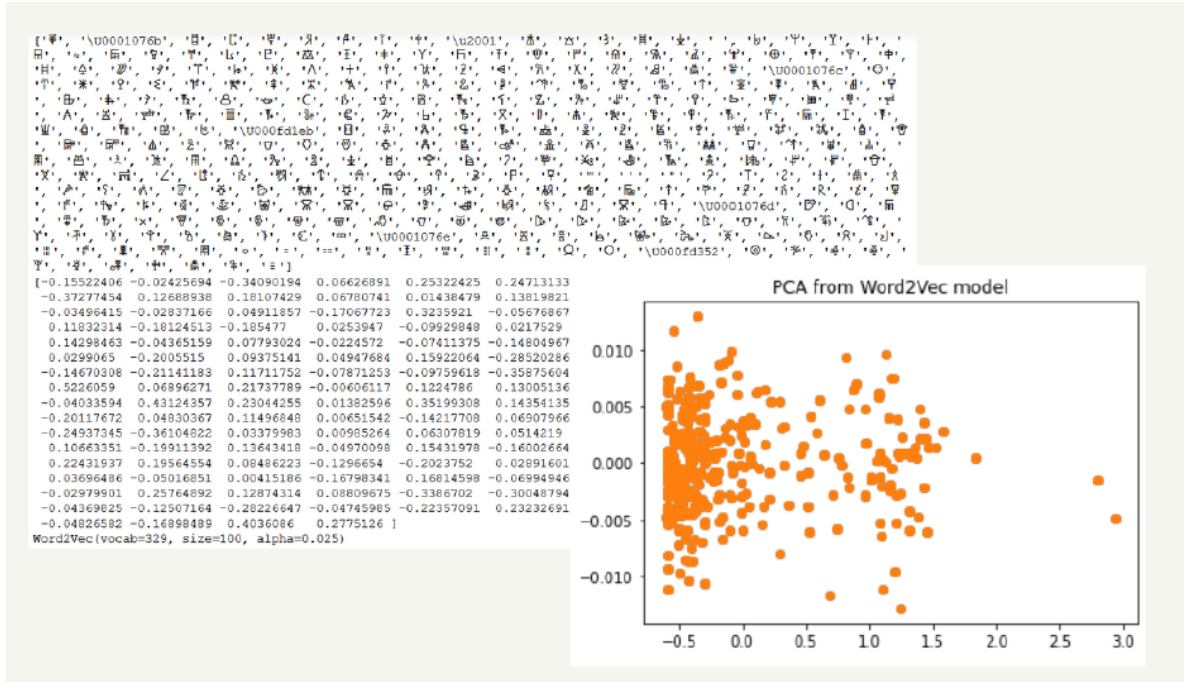


Figure 10: Word Embedding and Principal Component Analysis of Linear A data set

Word embedding of the transliterations show us the entropy value to be of less structured format, as expected with such a sparse dataset. The pictographic representation also leaves much to the randomness. The stochastic gradient descent representation of the covariance and correlation matrix for the PCA shows the dimensionality for a critical value of the alpha at a default number of 0.025. There are certain outliers which are further explored with Topic Modeling and k-Means clustering.

```

1 print (model.similarity('𐌆', '𐌆'))
2 print (model.similarity('𐌆', '𐌆'))
3 print (model.most_similar(positive=['𐌆'], negative=[], topn=4))
4 print (model.most_similar(positive=[], negative=['𐌆'], topn=4))

1.0
0.9994925
[('𐌆', 0.999587893486023), ('𐌆', 0.9995784759521484), ('𐌆', 0.9995618462562561), ('𐌆', 0.9995594620704651)]
[('𐌆', -0.9992014169692993), ('𐌆', -0.9994107484817505), ('𐌆', -0.9994924664497375), ('𐌆', -0.9995063543319702)]

```

Figure 11: Cosine Similarity analysis of the symbols

Word2Vec embedding also gives us the similarity analysis of the symbols by their vector values. As seen in the output, the first set of symbols are the same, so their similarity comes to 1.0, while the similar but not same symbols values are 0.99. The list of symbols with similar positive values as the first set of symbols and then the negative associations of the second symbol in the second set is also apparent.

```


Score (Opposite of the value of X on the K-means objective which
is Sum of distances of samples to their closest cluster center):
-0.22897068
Silhouette_score:
0.48115668

```

Figure 12: K-means clustering scores

K-Means Clustering gives us the groups of the transliterations based on their cosine similarity distance. The silhouette score is then grouped by the simpler Euclidean distance metric.

When the number of clusters =2,



When the number of clusters =5,

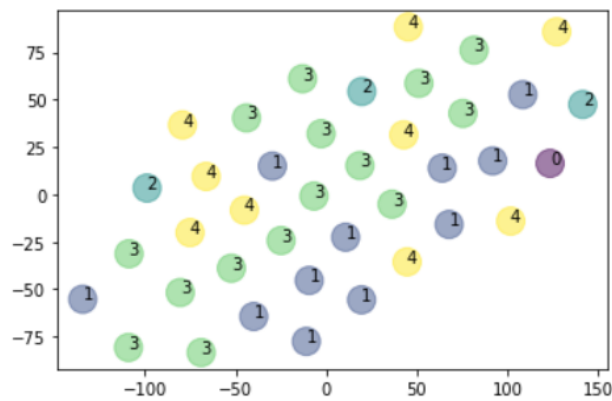


Figure 13: K-means clustering by groups

The within-cluster sum of squares is a measure of the different types of transliterations within each cluster. The average distance from observations in this graph is seemingly random which can be attributed to the high entropy given that it is a measure for the variability of the transliterations within each cluster. Centroid helps interpret each cluster as a focus for the cluster.

As a generative statistical model, Latent Dirichlet Allocation (LDA)¹⁴ uses known observations to group similar unobserved groups of terms, which is extremely useful for undeciphered languages.

¹⁴ The Latent Dirichlet Allocation (LDA) is used to generate a probabilistic model of a corpus with a variety of topics using Jensen–Shannon divergence between topics computation and term frequency relevance. LDA is a useful technique within the machine learning toolbox and thus, the artificial intelligence toolkit.




Figure 14: Topic Modeling process flow

LDA was originally developed for Evolutionary biology, bio-medicine, and engineering problems, and in computational linguistics as well.

A word cloud gives the visual representation of most common stone tablets in the text corpus (see Figure 15)



Figure 15: Word cloud of locations the Linear A tablets were found by the number of symbols gathered from each

The stone tablet “HT” seems to be the most popular in the text corpus which corresponds to the known archaeological study of stone tablets being found mainly in Hagia Triada (HT).




Figure 16: Interactive Topic Modeling graph

[pyLDAvis](#) helps to interpret individual topics and their relationships. The circles represent the multidimensional divergence, while the bar chart gives the relevance measure. The bar charts are for both corpus-wide frequency and topic-wise.

The terms are ranked by their relevance in the corpus data with the λ slider. While the terms are ranked by the probability of the topic, since the pyLDAvis is an interactive tool, the topics can be selected and adjusted accordingly. The λ slider adjusts the critical value of the relevant term rank for each specific topic. The value that is usually used is 0.6. The Topic Modeling results go on to show that there are five major groupings of the transcriptions in Linear A language as the Intertopic Distance Map (via multidimensional scaling) shows. While group 2, 3, and 4 are overlapping, which seems to indicate closer similarity, 1 and 5 are on the opposite side of the graph in the modeling graph.

The LDA proved to be an improvement in similarity analysis by topics. While this is one methodology to understand the patterns in the undeciphered language, and the visual gave a better idea of token relationships, by itself LDA was an improvement over the base NLTK entropy function.

2.2.2 Exploratory Knowledge Mining

In addition to NLP, we also used some machine learning methods to detect most common symbols. In order to use machine learning, the data needed to be pre-processed so that it could work best within the constraints of the two following methods we tried: Naïve Bayes¹⁵ (NB) and K-nearest-neighbor¹⁶ (kNN).

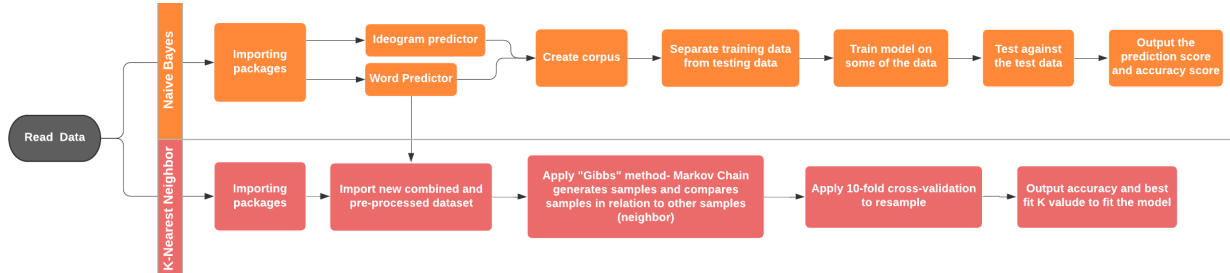


Figure 17: Knowledge Mining process flow

NB uses a probabilistic method for breaking down the components of the symbols. From a data pre-processing perspective, we needed to understand which symbols and representation to use (Bunescu and Mooney, 2004). We opted for using the symbols which were likely words, as well as symbols that were likely ideograms. The difference between symbols and words were great enough so that assigning a true or false to whether the symbol represented a word or ideograms, could allow for NB to make a determination on future symbols or ideograms, or if a symbol fell outside of these categories.

¹⁵In statistics, Naive Bayes classifiers are a family of simple "probabilistic classifiers" based on applying Bayes' theorem with strong independence assumptions between the features.

¹⁶ In instance-based classification, each new instance (of symbol, in our case) is compared with existing ones using a distance metric, and the closest existing instance is used to assign the class to the new one. This is called the nearest-neighbor classification method. (Witten et al. 2017)

of the text is present, it does so with greater error. With the “Gibbs” method, along with 10-fold cross validation, the accuracy of prediction was around 64%, with a kNN of 7 being used for the final model.

III Conclusion

3.1 Discussion

The NLP analysis gave us the frequency distribution of popular symbols in the Linear A stone tablets. n-gram summary of the tablets’ contents and predicted the possible subsequent symbols in other tablets. With the inclusion of symbols beyond probable words, the entropy reduced from 9.4 to 3 thereby verifying the consistent unstructured nature of the language via word embedding process using Word2Vec and Principal Component Analysis. Three methods of grouping the symbols were explored from Cosine Similarity Analysis, K-Means Clustering, Topic Modeling, and K-Nearest Neighbor verified by 10-fold Cross Validation, all of which were successful in various measures. Majority of the symbols were grouped in similar buckets except a few interesting outliers. Attempting to classify the symbols was an interesting exercise for ascertaining the accuracy of ideograms particularly.

Deciphering Linear A, the ancient Mycenaean language has been an ongoing effort for many decades now. Through our exploratory computational analysis, we hope to add to the discourse via introducing various new methodologies to this study of this language. To the best of our knowledge, this is the first study to discuss and show computational analysis of Linear A. The purpose of the current study was to examine and apply the possibilities of a series of Natural Language Processing and Machine Learning techniques to undeciphered languages. The paper also proposes to broaden the toolkit for these languages and bridging the gap between a linguistic approach and a data-mining approach. Utilizing computational techniques on the undeciphered scripts in Linear A, we yield some preliminary results. Our future work includes a comparison and analysis of Linear A with other ancient languages, both deciphered (Linear B) and undeciphered (Rongorongo).

3.2 Acknowledgments

We are incredibly grateful for Dr. William Kennedy, and Dr. Duoduo Liao’s support and guidance in this endeavour. Further specialized inferences will be gleaned from the results by Dr. Frederico Aurora, in conjunction with our team.

References

- "Using Zipf’s Law To Improve Neural Language Models", Medium, 2020. [Online]. Available: https://medium.com/@_init_/using-zipfs-law-to-improve-neural-language-models-4c3d66e6d2f6. [Accessed: 18-Oct- 2020].
- A. Bouchard-Côté, D. Hall, T. L. Griffiths, and D. Klein, “Automated reconstruction of ancient languages using probabilistic models of sound change,” Proc Natl Acad Sci USA, vol. 110, no. 11, p. 4224, Mar. 2013, doi: 10.1073/pnas.1204678110.

- Aurora, F. (2015). DĀMOS (Database of Mycenaean at Oslo). Annotating a Fragmentarily Attested Language. *Procedia - Social and Behavioral Sciences*, 198, 21-31. Doi:10.1016/j.sbspro.2015.07.415
- Buitinck et al., (2013). API design for machine learning software: experiences from the scikit-learn project.
- Bunescu, R. and Mooney, R.J. (2004). Collective information extraction with relational Markov networks. In *Proceedings of the 42nd ACL*, pages 439–446.
- C. E. Shannon., A mathematical theory of communication, *Bell System Tech. J.*, 27 (July 1948), 379–423; (October 1948), 623–656. Reprinted in: C.E. SHANNON, W. WEAVER, *The Mathematical Theory Of Communication*, The University Of Illinois Press, Urbana, 1949.
- Classen, M., & Safrany, L. (1975). Endoscopic papillotomy and removal of gall stones. *British Medical Journal*, 4(5993), 371–374. <https://doi.org/10.1136/bmj.4.5993.371>
- Daggumati, S., & Revesz, P. Z. (2019). Data mining ancient scripts to investigate their relationships and origins. *Proceedings of the 23rd International Database Applications & Engineering Symposium*, 1–10. <https://doi.org/10.1145/3331076.3331116>
- E. Matricciani. Shannon's entropy as a measure of the “life” of the literature of a discipline. *Scientometrics* 30, 129–145 (1994). <https://doi.org/10.1007/BF02017218>
- Emerging Technology. “Machine learning has been used to automatically translate long-lost languages,” *MIT Technology Review*. <https://www.technologyreview.com/2019/07/01/65601/machine-learning-has-been-used-to-automatically-translate-long-lost-languages/> (accessed Oct. 18, 2020).
- Evans, A. (1909). *Scripta Minoa: The written documents of Minoan Crete, with special reference to the archives of Knossos*. Oxford; Clarendon Press. <http://archive.org/details/scriptaminoawrit01evanuoft>
- F.R.S, K. P. (1901). LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11), 559–572. <https://doi.org/10.1080/14786440109462720>
- Harris, C.R., Millman, K.J., van der Walt, S.J. et al. Array programming with NumPy. *Nature* 585, 357–362 (2020). DOI: 0.1038/s41586-020-2649-2. (Publisher link).
- J. D. Hunter, "Matplotlib: A 2D Graphics Environment", *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90-95, 2007.
- J. Luo, Y. Cao, and R. Barzilay, “Neural Decipherment via Minimum-Cost Flow: from Ugaritic to Linear B,” *arXiv:1906.06718 [cs]*, Jun. 2019, Accessed: Oct. 18, 2020. [Online]. Available: <http://arxiv.org/abs/1906.06718>.

- J. Novembre, “Pritchard, Stephens, and Donnelly on Population Structure,” *Genetics*, vol. 204, no. 2, pp. 391–393, Oct. 2016, doi: 10.1534/genetics.116.195164.
- Jurafsky, D., & Martin, J. H. (2008). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*. Pearson Education.
- Lloyd, S. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2), 129–137. <https://doi.org/10.1109/TIT.1982.1056489>
- Luo, J., Hartmann, F., Santus, E., Cao, Y., & Barzilay, R. (2020). Deciphering Undersegmented Ancient Scripts Using Phonetic Prior. *ArXiv:2010.11054 [Cs]*. <http://arxiv.org/abs/2010.11054>
- M. Waskom, “mwaskom/seaborn: v0.11.0 (Sepetmber 2020)”. Zenodo, 08-Sep-2020.
- MacQueen, J. (1967). *Some methods for classification and analysis of multivariate observations*. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics. <https://projecteuclid.org/euclid.bsmsp/1200512992>
- Malouf, R. (2002) Markov models for language-independent named entity recognition. In Proceedings of the 6th CoNLL, pages 187–190.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. *ArXiv:1310.4546 [Cs, Stat]*. <http://arxiv.org/abs/1310.4546>
- N. M. Dinesh, S. Narasimhan, P. Prashanth, and S. E. P. Pushpa, “Entropy of Tamil Language and Prioritized Coding Algorithm for Encoding of Tamil Letters,” in 2018 International Conference on Smart Systems and Inventive Technology (ICSSIT), Dec. 2018, pp. 395–397, doi: 10.1109/ICSSIT.2018.8748705.
- R. Lee, P. Jonathan, and P. Ziman, “Pictish symbols revealed as a written language through application of Shannon entropy,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 466, no. 2121, pp. 2545–2560, Sep. 2010, doi: 10.1098/rspa.2010.0041.
- Rudolf Clausius. (1867). *The Mechanical Theory of Heat: With Its Applications to the Steam-engine and ...* J. Van Voorst. <http://archive.org/details/mechanicaltheor04claugoog>
- S. Bird, E. Loper, E. Klein (2009), *Natural Language Processing with Python*. O’Reilly Media Inc.
- Sievert, Carson & Shirley, Kenneth. (2014). LDAvis: A method for visualizing and interpreting topics. 10.13140/2.1.1394.3043.
- T. Petrolito, R. Petrolito, F. Perono Cacciafoco, and G. Winterstein, “Minoan linguistic resources: The Linear A Digital Corpus,” in Proceedings of the 9th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH), Beijing, China, Jul. 2015, pp. 95–104, doi: 10.18653/v1/W15-3715.
- Witten, I. H., Frank, E., Hall, M. A., & Pal, C. J. (2017). *Data mining: Practical machine learning tools and techniques*.


```

"QE-RA2-U",
".",
"\n",
"KI-RO",
"197",
"\n",
"□-SU",
"70",
"\n",
"DI-DI-ZA-KE",
"52",
"\n",
"KU□-NU",
"109",
"\n",
"A-RA-NA-RE",
"105"
],
"words": [
"□□□",
".",
"\n",
"□□",
"○≡≡≡≡",
"\n",
"□□",
"≡≡",
"\n",
"□□□□",
"≡≡",
"\n",
"□□□",
"○≡≡≡",
"\n",
"□□□□",
"○≡≡"
]
]

```

}},
The inscription that we need to visualize is encoded in the code above. The image files, the raw transcription, and arrays representing the parsed words of the inscription both in Linear A, transliterated Linear A syllabograms, and proposed translations where applicable are included in this.

ANNEX 2

Cras tristique vel nisi at aliquet. Proin egestas erat sit amet velit lobortis imperdiet. Integer et arcu sapien. Etiam id blandit sapien. Nam tempus lacus ac massa semper, vel laoreet turpis rutrum. Mauris eget nibh vitae justo porta imperdiet sed vel ligula. In imperdiet, augue vel condimentum convallis, neque augue imperdiet neque, eget dapibus nunc mauris ultricies tortor. Nam eget nunc egestas, blandit lectus non, aliquam nunc. Cras sed quam vitae arcu ornare lobortis. Ut ut lacus hendrerit, convallis orci sit amet, commodo nunc. Pellentesque eget tincidunt tortor. Nunc ornare molestie mauris id vehicula. Suspendisse pharetra tortor metus, sit amet fermentum tellus vehicula ut.