

An Independent Unsupervised Examination of the Distinction Between Texts of Priestly and Non-priestly Origins in the Books of Genesis and Exodus

GIDEON YOFFE ^{1,2} AXEL BÜHLER ^{3,4} THOMAS RÖMER ³ NACHUM DERSHOWITZ ⁵ ELIEZER PIASETZKY ⁶
ISRAEL FINKELSTEIN ² AND BARAK SOBER ¹

¹*Dept. of Statistics and Data Science, the Hebrew University of Jerusalem, Mount Scopus, Jerusalem 91905, Israel*

²*Dept. of Archaeology and Ancient Near Eastern Cultures, Tel Aviv University, Ramat Aviv 6997801, Israel*

³*Collège de France, 11 Pl. Marcelin Berthelot, 75231 Paris, France*

⁴*Faculté de théologie, Université de Genève, 1205 Geneva, Switzerland*

⁵*School of Computer Science, Tel Aviv University, Ramat Aviv 6997801, Israel*

⁶*School of Physics and Astronomy, Tel Aviv University, Ramat Aviv 6997801, Israel*

1. INTRODUCTION

It is agreed among scholars that the extant version(s) of the Hebrew Bible is a result of all sorts of various editorial actions (additions, redaction, etc.). As such, it may be viewed as a literary patchwork quilt—whose patches differ, for example, by genre, date, and origin—and the distinction between biblical texts as well as their relation to other ancient texts from the Near East is at the heart of biblical scholarship, theology, ancient Israel studies, and philology. Many debates in this field have been ongoing for decades, if not centuries, with no verdict (e.g., the division of the biblical text into its “original” constituents (8; 19; 20)). While several paradigms gained prevalence throughout the evolution of this discipline (20; 18; 21), the jury is still out on many related hypotheses, such that these paradigms are prone to drastic (and occasionally abrupt) changes over time (e.g., 14).

When scrutinized, the interrelation between various literary features within the text may shed light on the historical context of its composition (e.g., 20; 18; 13; 16).

Standard approaches typically entail scrupulous manual effort by biblical scholars and linguists on corpora of limited size (e.g., 6; 22; 10). These are prone to subjective reasoning and arguments, which are sometimes hard to disentangle and establish whether they are clear-cut. Here, statistical and computational techniques offer a critical leap forward in detecting distinctive literary features, providing a more empirical way of measuring the significance of such underpinnings.

We examine the hypothetical distinction between texts of priestly (P) and non-priestly (nonP) origin in Genesis and Exodus, for which exists a surprisingly large agreement amongst biblical scholars (e.g., 12; 15; 7). Examining this distinction with an independent, unsupervised computational methodology would establish a measure of confidence therein and encourage its application to additional instances of biblical texts, especially those of greater controversy, where our approach could help tilt the scale in favor of one hypothesis over another.

2. METHODOLOGY

We intertwine descriptive and inferential statistics. The first are used in text classification and interpretability analyses, whereas the latter, to quantify uncertainty through hypothesis testing. While descriptive statistics were successfully applied to specific texts (e.g., 11; 17), we are unaware of similar studies where uncertainty quantification was considered. Furthermore, identification of literary features *responsible* for the classification, as opposed to cluster-wise significant feature detection (e.g., 9; 2; 23), is novel for stylometry.

2.1. Corpus

We use STEP Bible¹ (digitized Leningrad codex), with its morphological and semantic tags for all words, prefixes, and suffixes. We consider two representations of the text: word-wise and a grammatical representation by phrase-dependent parts-of-speech (pdps).

We obtained a scholarly labeling assigning each verse in Genesis and Exodus as P/nonP.

¹ <https://github.com/STEPBible/STEPBible-Data>

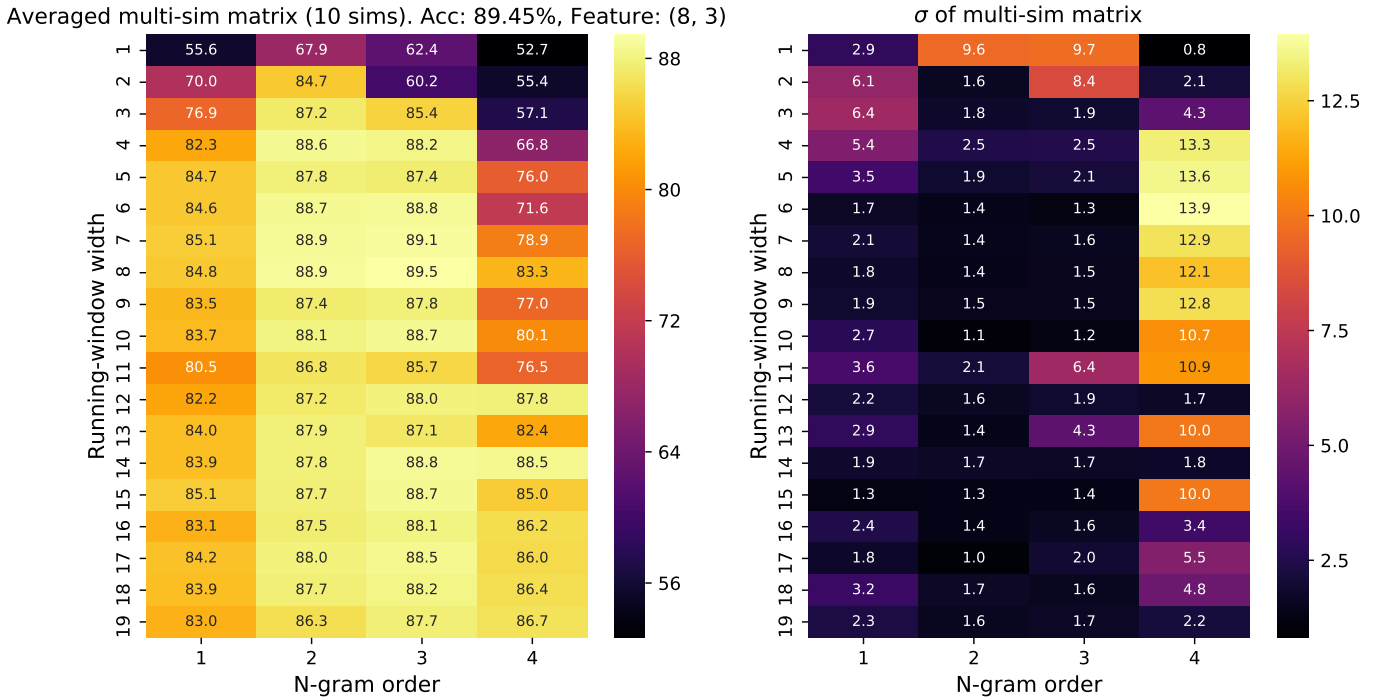


Figure 1. Cross-validated grid-search optimization for Exodus (pdps). **Left Matrix:** averaged accuracy over 10 simulations with respect to combinations of window sizes (y -axis) and n -gram sizes (x -axis). The best-fit combination is 8 and 3, respectively, yielding $\approx 89.5\%$ accuracy. **Right Matrix:** pixel-wise standard deviation of the left matrix.

2.2. Parameterization and Embedding

Our underlying assumption is that significant literary differences between texts manifest in simple linguistic parameters. Therefore, we consider three parameters, distinct combinations of which result in different classifications. These are: (1) word-/pdp-wise representations, (2) n -gram size, the length of sequences of consecutive words/pdps, and (3) running-window size, the number of verses surrounding the original, providing additional context.

We use tf-idf to encode each verse, assigning a relevance score to each feature in the context (1). The critical consideration behind choosing this traditional embedding is that it allows interpretability of the results, unlike neural-net-based language models, which are convoluted (e.g., 3; 5).

2.3. Optimization

We use k -means to classify the embedded verses and use an unbalanced accuracy measure to quantify the goodness of classification. We perform cross-validated grid-search on a range of running-window and n -gram sizes for words/pdps, identifying the combination that yields the highest accuracy (Fig. 1).

2.4. Testing and Validating

Through hypothesis testing, we aim to establish statistical significance of the distinction. We perform two tests, under the null hypothesis that our labels were randomly assigned: (1) arbitrary permutations, and (2) *cyclic* permutations, where we generate the null by shifting the labeling cyclically, seeking to conserve implicit correlations between consecutive verses.

2.5. Feature Importance

Minimizing k -means loss is equivalent to maximizing *inter-cluster* variances. Leveraging this, we extract a vector of feature-wise importance that maximizes the inter-cluster variance found by 2-means. This vector allows us to trace the features most responsible for the classification (Fig. 2).




Figure 2. Feature importance bar plot for Exodus embedded as trigrams of pdps with window width 8. Here, 100% of the distinction is made by features that are abundant in the nonP cluster.

3. CONCLUSIONS

We examined the hypothesized P/nonP distinction in Genesis and Exodus and introduced a novel computational and statistical methodology for text stylometry that is essentially independent of, but in synergy with, established philological practices. We sought an optimal single feature—a combination of running-window and n -gram sizes, and extracted features that contribute most to classification and their respective proportions. We achieve a 73% and 90% (balanced) accuracy for Genesis (word unigrams with window of 6) and Exodus (pdp trigrams with window 8), respectively. The difference in accuracy between the two seems to arise from the more sporadic distribution of P in Genesis, as opposed to a more formulaic one in Exodus.

Comparing our results to (4), who performed a synonym-words-based analysis of the P/nonP distinction, we improve accuracy for Genesis by $\approx 12\%$, and attain similar accuracy for Exodus.

REFERENCES

- [1] Akiko Aizawa. An information-theoretic perspective of tf-idf measures. *Information Processing & Management*, 39(1):45–65, 2003.
- [2] Deng Cai, Chiyuan Zhang, and Xiaofei He. Unsupervised feature selection for multi-cluster data. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 333–342, 2010.

- [3]Supriyo Chakraborty, Richard Tomsett, Ramya Raghavendra, Daniel Harborne, Moustafa Alzantot, Federico Cerutti, Mani Srivastava, Alun Preece, Simon Julier, Raghuvier M Rao, et al. Interpretability of deep learning models: A survey of results. In *2017 IEEE Smartworld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (Smartworld/SCALCOM/UIC/ATC/CBDcom/IOP/SCI)*, pages 1–6. IEEE, 2017.
- [4]Idan Dershowitz, Nachum Dershowitz, Tomer Hasid, and Amnon Ta-Shma. Orthography and biblical criticism. In *Proceedings of Digital Humanities (DH 2014, Lausanne, Switzerland)*, pages 451–453, 2014.
- [5]Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [6]Martin Ehrensverd. Once again: The problem of dating biblical Hebrew. *Scandinavian Journal of the Old Testament*, 11(1):29–40, 1997.
- [7]Avraham Faust. The world of P: The material realm of priestly writings. *Vetus Testamentum*, 69(2):173–218, 2019.
- [8]Hermann Gunkel. *Creation and Chaos in the Primeval Era and the Eschaton: Religio-Historical Study of Genesis 1 and Revelation 12*. Eerdmans, Grand Rapids, MI, 2006. Trans. by K. William Whitney, Jr.; original edition 1895.
- [9]Eduardo R. Hruschka and Thiago F. Covoos. Feature selection for cluster analysis: an approach based on the simplified silhouette criterion. In *International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC’06)*, volume 1, pages 32–38. IEEE, 2005.
- [10]Jarod Jacobs. The balance of probability: statistics and the diachronic study of ancient Hebrew. *Journal for Semitics*, 25(2):927–960, 2016.
- [11]Mike Kestemont, Justin Stover, Moshe Koppel, Folgert Karsdorp, and Walter Daelemans. Authenticating the writings of Julius Caesar. *Expert Systems with Applications*, 63:86–96, 2016.
- [12]Israel Knohl. *The Sanctuary of Silence: The Priestly Torah and the Holiness School*. Eisenbrauns, 2007.
- [13]Israel Knohl. *The Divine Symphony: The Bible’s Many Voices*. Jewish Publication Society, 2010.
- [14]Ernest Nicholson. *The Pentateuch in the Twentieth Century*. Oxford University Press, November 2002. doi:10.1093/acprof:oso/9780199257836.001.0001.
- [15]Thomas Römer. From the call of Moses to the parting of the sea: Reflections on the priestly version of the Exodus narrative. In *The Book of Exodus*, pages 121–150. Brill, 2014.
- [16]Thomas Römer. *The Invention of God*. Harvard University Press, November 2015. doi:10.4159/9780674915732.
- [17]Mayuri Verma. Lexical Analysis of Religious Texts using Text Mining and Machine Learning Tools. *International Journal of Computer Applications*, 168(8):39–45, June 2017. doi:10.5120/ijca2017914486.
- [18]Gerhard von Rad. *Genesis – A commentary*. S.C.M. Press, London, 3rd rev. ed. edition, 1972. Trans. by John H. Marks.
- [19]Gerhard von Rad. *Old Testament Theology: The theology of Israel’s historical traditions*, volume 1. Westminster John Knox Press, 2001.
- [20]Julius Wellhausen, J. Sutherland Black, Allan Menzies, and William Robertson Smith. *Prolegomena to the History of Israel*. Cambridge University Press, 2009. doi:10.1017/cbo9781139383363.
- [21]Yair Zakovitch. A study of precise and partial derivations in biblical etymology. *Journal for the Study of the Old Testament*, 5(15):31–50, 1980.
- [22]Yair Zakovitch. Inner-biblical interpretation. *A Companion to Biblical Interpretation in Early Judaism*, pages 27–63, 2012.
- [23]Guo-Niu Zhu, Jie Hu, Jin Qi, Jin Ma, and Ying-Hong Peng. An integrated feature selection and cluster analysis techniques for case-based reasoning. *Engineering Applications of Artificial Intelligence*, 39:14–22, 2015.