
Inferring Social Rank in an Old Assyrian Trade Network

David Bamman
School of Computer Science
Carnegie Mellon University
dbamman@cs.cmu.edu

Adam Anderson
Department of Near Eastern
Languages and Civilizations
Harvard University
aganders@fas.harvard.edu

Noah A. Smith
School of Computer Science
Carnegie Mellon University
nasmith@cs.cmu.edu

1 Introduction

In the early 20th century, the attention of Assyriologists and archaeologists was directed to a number of cuneiform tablets coming from a remote archaeological tell in Kültepe, Turkey. After the first series of excavations, archaeologists discovered a large collection of texts and the remains of a Bronze Age trade colony, referred to in the texts as *kārum Kaneš*. Once these initial ca. 5,000 texts were deciphered, the field of Old Assyrian studies was born. In 1948 official Turkish excavations began at Kültepe and added over 17,000 tablets to the Old Assyrian text corpus, which now totals ca. 23,000 cuneiform tablets [5]. These texts document the intricacies of thriving Bronze Age trade networks, comprised of Old Assyrian merchants from the ancient city of Assur approximately 4,000 years ago (ca. 1950-1750 BCE) [1]. The texts further show how the merchants acted as the middle-men in a large series of inter-connected networks which, among other things, linked the natural resources of tin (in Iran and Afghanistan) and copper (in Turkey) in order to produce bronze in Anatolia.

However, one thing the texts do not make clear is the scope and structure of the colonial trade network, in terms of the people involved and their organization. Although the high degree of literacy among the inhabitants of the colony at Kaneš helped create an extremely rich source of texts illustrating the daily life of the people involved, the practice of *papponomy* (naming a son after his grandfather) has obscured the identities of the merchants for modern scholarship. Thus, due to the density and ambiguity of the names mentioned in these texts, it has been too difficult to gain an understanding of the scope of the colonial society on the basis of the textual record at Kültepe.

Our work therefore focuses on jointly inferring the unique individuals as well as their social rank within the Old Assyrian trade network, using a novel probabilistic latent-variable model that exploits partial rank information contained in the texts.

2 Data

Of the 23,000 tablets unearthed at Kültepe (Kaneš), 5,691 published texts (known as the “old texts” [7]) have been digitized and transcribed into machine-actionable text as part of the Old Assyrian Text Project [6]. While these tablets include mostly economic and legal transactions, 2,094 of them are letters between merchants. Along with the body of the Akkadian text, each of these letters includes a highly stylized introductory formula (an “epistolary formula”) which lists the senders and recipients using strict dominance rules concerning the order of the names. For example:

| | | | | | | |
|------|-----------|----------------|-----|-------------|--------|-------------------|
| umma | Aššur-idī | Aššur-nādā | ana | Amur-Ištar | Alāhum | Aššur-taklāku |
| from | Aššur-idī | and Aššur-nādā | to | Amur-Ištar, | Alāhum | and Aššur-taklāku |

These formulae have a consistent internal structure from which we can draw relative social ranks among the individuals involved. Each formula can be divided into two parts (a *receiving* rank and a *sending* rank), and an individual placed linearly after another *within* one of these ranks cannot be socially higher than any mentioned before (whether the first is higher or equal is ambiguous). Additionally, one individual (who may be either among the senders or recipients) is mentioned first in the letter, a marked position signifying the highest social status of those mentioned in either rank.

$$\overbrace{\text{Aššur-idī} + \text{Aššur-nādā}}^{\text{first mentioned}} \rightarrow \underbrace{\text{Amur-Ištar} + \text{Alāhum} + \text{Aššur-taklāku}}_{\text{receiving rank}} \quad (1)$$

sending rank
receiving rank

These partial orderings provide a rich source of evidence for the global social structure; from this example, we can extract seven pairwise partial social orderings: Amur-Ištar \geq Alāhum and Aššur-taklāku; Alāhum \geq Aššur-taklāku; and Aššur-idī \geq all four of the others.

If all such partial orderings were to be trusted, if each observed name in such a formula were unambiguous, and if social power were a stationary quality that remained constant over time, inferring a consistent global rank over all individuals would be easy (though more than one such rank may be possible). Unfortunately, however, none of these assumptions are true. The rank we observe in one letter is a subjective judgment by the author, and we can easily imagine that complex social dynamics are involved in the choice of who to rank highest (which can vary by author); names are indeed ambiguous with one name potentially referring to multiple people, and the same person can be known by several names; and the letters span a period of ca. 200 years, over which time a young individual with low social rank can age and accrue power.¹

3 Technical Approach

Our goal is to find the social ranking over individuals (possibly not in a one-to-one relation with names) that best explains the observed data. To illustrate the intuition behind our approach, consider a simple example. Suppose we are trying to establish the temporal rank of a set of individuals with the names ADAMS, JEFFERSON and MADISON, and we have the following evidence (where $>$ indicates “was president before”).

- ADAMS $>$ JEFFERSON
- JEFFERSON $>$ MADISON

Assuming transitivity, a sound global rank among these three is: ADAMS $>$ JEFFERSON $>$ MADISON; while we never directly witness a statement of the sort ADAMS $>$ MADISON, we can infer it through intermediary relations. Now suppose we observe an additional piece of evidence:

- MADISON $>$ ADAMS

If we assume that the three names only refer to three distinct individuals, transitivity breaks down: putting all three statements together results in a circular rank, leading to the contradiction that ADAMS $>$ JEFFERSON while at the same time JEFFERSON $>$ ADAMS. However, we can establish a sound global order if we allow the three names to refer to four individuals (e.g., two people both have the name ADAMS), resulting in the rank: ADAMS₁ $>$ JEFFERSON $>$ MADISON $>$ ADAMS₂.² In fact, given the inconsistency of the evidence under the assumption of only three people, the existence of four underlying people is in fact more likely. Here, our method offers an informed hypothesis—that ADAMS refers to two distinct individuals rather than one—that can be verified (or refuted) in consultation with the data. In this simple case, the hypothesis is supported by the fact that ADAMS can refer to both JOHN ADAMS and JOHN QUINCY ADAMS.

In the case of our Old Assyrian dataset, the evidence takes the form of 4,191 pairwise observed ranks of 717 individual names in 1,657 letters, along with the Akkadian text of those letters. Our task is to find the

¹Most of the letters in our corpus have not been dated to a finer level of granularity than century; a potential extension to the model described here would exploit this information when available.

²Alternatively, JEFFERSON₁ $>$ ADAMS $>$ MADISON $>$ JEFFERSON₂ is also valid, as is MADISON₁ $>$ JEFFERSON $>$ ADAMS $>$ MADISON₂. In data where the orderings are not strict (i.e., ADAMS \geq JEFFERSON), global ranks involving equalities are also possible.

most likely overall social rank of a fixed set of actors that best explains the pairwise ranks in letters that we see. In casting the problem in this way, we are building a *probabilistic generative model* of the data. Latent Dirichlet allocation [2] is a familiar example of a generative model that seeks to explain data in text documents by inferring latent topic assignments to individual words. In our case, the latent variables are a.) the identities of the people named in each letter; and b.) the social rank of those people, represented as continuous values.

Figure 1 illustrates the graphical model in detail using standard notation alongside our inference algorithm. In brief, we use a randomized algorithm known as Monte Carlo Expectation Maximization [8]. The algorithm alternately a.) samples a value of the latent entity z for each instance of a name x in a given letter, conditioned on current values of all other latent variables and parameters, cycling through the name instances, and b.) uses those accumulated samples to optimize the values of the social rank β that generated the pairwise ranks y in evidence. In this way, we alternate between picking probable latent individuals referred to in letters (given some fixed social rank), and determining the best social rank given that estimate of who those names refer to.

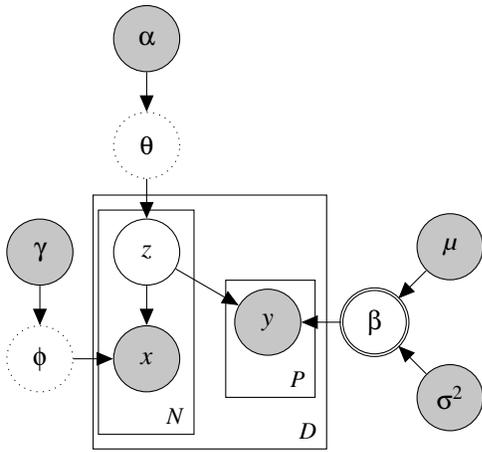


Figure 1: Graphical model and MCEM algorithm. Shaded circles represent observed or fixed variables, empty circles are latent variables, dotted circles represent variables that are integrated out via collapsed Gibbs sampling, and the double-circled β is optimized via MCEM. $\phi, \theta \sim \text{Dir}$; $z \sim \text{Categorical}$; $x \sim \text{Multinomial}$; $y \sim \text{Logistic}$; and $\beta \sim \text{Normal}$. In our case, we fix the hyperparameters $\alpha = 100, \gamma = 0.01, \mu = 0, \sigma^2 = 1$, and the number of possible latent entities to 1000.

4 Results

The input to our algorithm is a set of pairwise ranks between names mentioned in a letter (of the form $A\check{s}\check{s}ur-id\bar{i} \geq A\check{s}\check{s}ur-n\bar{a}d\bar{a}$, as above), along with an upper bound on the number of the latent entities we expect (K); the output is twofold: a.) a global rank of those K latent entities, along with the names in letters associated with them most often; and b.) a distribution over all possible latent entities for each name mentioned in each letter.

We apply our model to *hypothesis generation*: given a set of evidence, the algorithm offers hypotheses it finds likely, which a domain expert can then validate according to established methods in the field. One such lead generated by our method about a well-studied individual concerns the name of Innāya.

Innāya

In 1991, Cécile Michel produced a two-volume work on two merchants in the colony of Kaneš named Innāya [3, 4]. On the basis of two attested patronyms it was apparent that there were at least two individuals who were known by this name. By charting their family trees and the structure of their respective businesses, she reconstructed the separate archives and identities of these two merchants. The first and best attested individual is Innāya son of Elāli with 142 texts, who appears to have a more complete textual record. The second individual, Innāya son of Amurāya, is only attested in 74 texts; while these two merchants overlapped

chronologically, the latter appears to have been a minor figure in the colony [3, p. 48]. While there were a number of texts which Michel was unable to determine (ca. 57), her study illustrated the complexity involved in the Old Assyrian archives due to an active use of homonyms at that time. While the work on Innāya is in no way complete, Michel provided a basis on which future scholarship might build, and will serve as a proving ground for the purpose of this study.

Figure 2 shows the overall distribution for latent ranks associated with the name Innāya in all of the letters, as learned by our model. Given the evidence that we have seen, our algorithm has learned that this name is generally associated with a very high-ranking individual (in this, and all other plots, the highest rank is 1, with 1000 being the lowest).

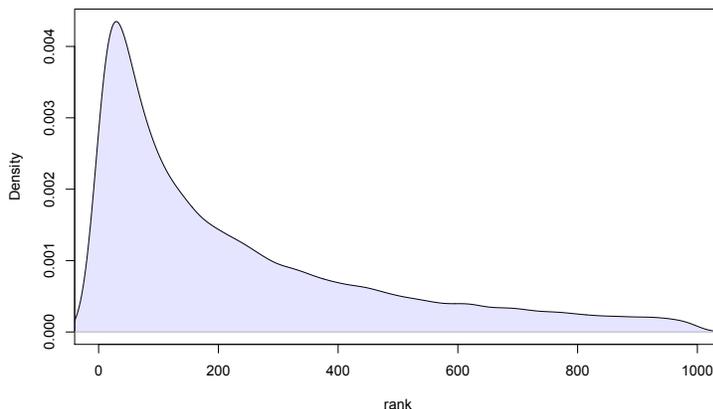


Figure 2: Overall distribution of latent ranks for the name Innāya. The highest rank is 1; 1000 is the lowest.

If we look at the individual letters themselves, however, a more nuanced picture emerges. Figure 3 plots the distribution of latent entities for a set of 6 of the 190 letters in which Innāya is mentioned. While the majority of letters in this set recapitulate the overall distribution shown above, several letters are noticeable outliers. For example, in letter TC1,33 and BIN6,109, our model has high belief that the real person associated with the name Innāya is not in fact the high-ranking individual at all, but rather a much lower ranked one. Consulting these letters, we see that Innāya is dominated by one or more individuals with a relatively low rank, which is not consistent with a high-ranking individual. If we look at the intersection of the publicly available letters in our collection and those inspected by Michel (a total of 142 texts), we find that our method agrees with Michel’s assessment of the identity of Innāya in each specific letter 80.9% of the time (discounting the level of agreement due to chance, this leads to a Cohen’s κ of 0.435). These results clearly support the conclusions Michel has drawn for Innāya – two individuals, at least, each with differing social networks and hierarchical ranks – and provide evidence for the validity of our approach.

5 Conclusion

The case of Innāya illustrates one of the greatest obstacles facing the field of Old Assyrian studies to date. Before any definitive statements can be made about the organization and makeup of the Old Assyrian trade network on the basis of the texts found in Kültepe, we must first determine the scope and structure of the Old Assyrian colonial society. Unfortunately, due to an active use of homonyms in the textual record, the scope of this colonial society has been obscured by a level of ambiguity too complex for any single specialist, or for that matter any group of specialists, to untangle.

As part of a solution, we present a method for aggregating small, local pieces of information—pairwise social differences between names in a cuneiform tablet—into a single underlying social order that offers the best explanation for the data we have. The latent variable model that we design allows us to be clear about our assumptions (the relationship between variables is encoded in the structure of the model; whether or not one variable is allowed to exert a direct effect on another is transparent). These kinds of generative models also allow us to add other forms of evidence; one possible extension would tie the choice of latent entity for each name in a letter with all of the other text observed (so that, for example, if one Innāya is often

associated with letters mentioning *tin* while another trafficks in *textiles*, we have further evidence that the two are different individuals).

In applying this method to our Old Assyrian dataset, we are engaging in exploratory data analysis, offering informed hypotheses that are driven by data, and that our model believes are the most promising avenues for directed research by Assyriologists. Grounding these hypotheses in the data allows us to return to the source of our induction – the letters themselves – and validate whether or not we have sufficient evidence to support our claims. In our particular case, the agreement between our model’s beliefs and those in the published literature are encouraging. We leave to future work to explore the differences that remain.

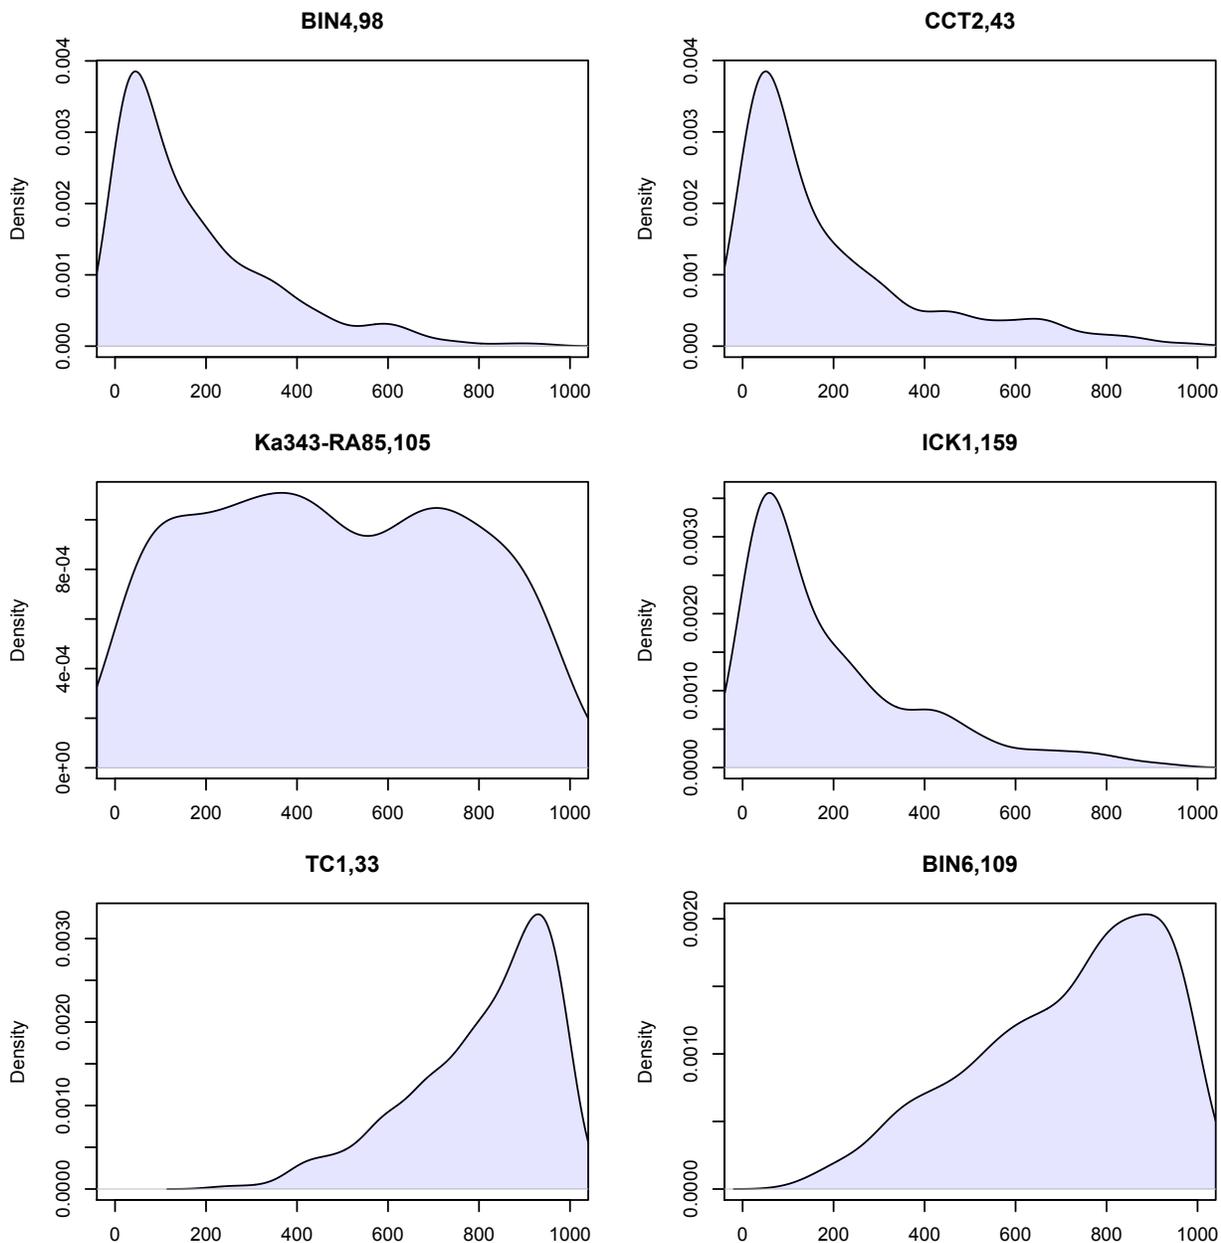


Figure 3: Distribution over latent ranks associated with the name Innāya in various letters. The highest rank is 1; 1000 is the lowest.

6 Acknowledgments

We would like to thank Mogens Larsen, Thomas Hertel, and other contributors to the Old Assyrian Text Project for providing the digitized texts we analyze. The research reported in this article was supported in part by Google (through the Worldly Knowledge Project at CMU) and by an ARCS scholarship to D.B.

References

- [1] Gojko Barjamovic, Thomas Hertel, and Mogens T. Larsen. *Ups and Downs at Kanesh: Chronology, History and Society in the Old Assyrian Period*. PIHANS (Publications de l'institut historique-archéologique néerlandais de Stamboul), vol. 120. Nederlands Instituut voor het Nabije Oosten, Leiden, 2012.
- [2] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, March 2003.
- [3] Cécile Michel. *Innāya dans les Tablettes paléo-assyriennes I: Analyse*. Editions Recherche sur les Civilisations, Paris, France, 1991.
- [4] Cécile Michel. *Innāya dans les Tablettes paléo-assyriennes II: Edition des textes*. Editions Recherche sur les Civilisations, Paris, France, 1991.
- [5] Cécile Michel. The Alāhum and Aššur-taklāku Archives found in 1993 at Kültepe Kaniš. *Altorientalische Forschungen*, 35:53–67, May 2008.
- [6] Old Assyrian Text Project. <http://oatp.ku.dk>.
- [7] K. R. Veenhof and J. Eidem. *Mesopotamia: The Old Assyrian Period*. Orbis Biblicus Et Orientalis 160/5. Academic Press (Fribourg), Vandenhoeck and Ruprecht (Göttingen), 2008.
- [8] Greg C. G. Wei and Martin A. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *Journal of the American Statistical Association*, 1990.