

In Search of Lost Time - Estimating Dates of Broken Cuneiform Documents

Stav Klein

`stav.klein@msmail.ariel.ac.il`

Ariel University

Avital Romach

Yale University

Shai Gordin

Ariel University

Article

Keywords:

Posted Date: April 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9033722/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

In Search of Lost Time - Estimating Dates of Broken Cuneiform Documents

March 13, 2026

Abstract

This paper tackles a simple but crucial question: can we estimate the dates of ancient texts when their original dates are missing? Historians have long relied on prosopography, the method of dating a text by checking whether the people it names appear in other, securely dated documents. The challenge is that this process quickly becomes unmanageable and inconsistent when applied across thousands of records. Our contribution is to bring computational tools and, in particular, network analysis, to bear on this traditional approach, allowing dates to be estimated systematically, quickly, at scale, and with full transparency and reproducibility.

1 Introduction

The cuneiform writing system, the first writing system in the world attested from the latter fourth millennium BCE to the first centuries CE, marks one of humanities' earliest transitions from pre-history to history. Developed in Mesopotamia, modern day Iraq, and spread throughout the Near East, clay tablets bearing cuneiform signs are the first written evidence we have for the events of the time, from economic, administrative, and legal documents to letters, scientific compendia, works of literature and much more. Some genres of cuneiform documents were dated with great precision, giving the day, month, and regnal year of the ruling king. For cuneiform documents dated to the first millennium BCE, that information, when complete, can be converted into an exact Julian year (see also the date converter tool on the eBL platform (Parker and Dubberstein, 1971)). This is vital information for the reconstruction of historical events and their proper sequence.

Yet the medium itself of cuneiform writing poses a challenge. Clay tablets are by far more durable than other ancient writing surfaces, such as leather or papyri, which are not preserved except in

desert environments such as Egypt. However, over millennia the clay tablets can become brittle, cracked, worn, or lose parts of their surface; environmental conditions after excavation may accelerate deterioration if tablets are not conserved properly (Gütschow, 2012). Beyond physical loss, the record itself is uneven: tablets are not discovered in equal numbers across all periods, places, or individuals' lifetimes. Together, these gaps mean that many texts cannot be securely placed in time, limiting how they can be connected to one another or to the wider historical picture.

This contribution focuses on the Neo-Babylonian and Achaemenid periods (c. 620–330 BCE), for which we have abundant archival documents dated to exact days of the month and year, according to the Babylonian lunar based calendar (Stern, 2012, pp. 71–124). From this period many archives were found, of some of the main temples of the period as well as archives of prominent families in Babylonian society, rich landowners and successful entrepreneurs. Generations of the same families can be reconstructed from these documents, their connections to the temple and palace institutions, and the ups and downs of their fortunes (Jursa, 2005; Jursa, 2010). A digital prosopography of individuals attested in these sources, that covers a fair share of the published documents, was made available under the auspices of the *Prosobab* project led by Waerzeggers and Groß, which is the baseline dataset that made this research possible (Waerzeggers, Groß, and al., 2019).

The specificity of Babylonian dating is especially important for studying micro-historical phenomena, individual careers, family histories, or small-scale events, where even a few years' difference can alter the historical interpretation (Kleber, 2013; Waerzeggers, 2014a; Podany, 2022). Thus, documents with broken or missing dates create gaps that can drastically effect event reconstruction. To bridge these gaps, the main tool of the historian is prosopography: the systematic study of individuals and their lives across multiple documents (Stone, 1971; Radner, 1998; Zadok, 2018). If an undated tablet mentions people who are also found in securely dated texts, then those overlaps can be used to approximate when the undated tablet was written. Additional information such as titles, offices, marital status, or birth of children, helps to better pinpoint when a specific document was more likely to be written.

This method has been indispensable for reconstructing fragmentary archives and tracing individuals over time (Kümmel, 1979; Bongenaar, 1997; Zadok, 2002; Zadok, 2003). Yet, it also has clear limitations. Most importantly, the method cannot be applied manually across thousands of documents without becoming slow and unreplicable. Furthermore, in practice, prosopography involves weighing contextual factors (kinship ties, offices or titles, etc.), and it is difficult to equally and consistently weigh so many varied factors when performed manually. One reconstruction can consequently weigh another one, and can lead to circular reasoning. All of these factors led to the development of digital approaches to prosopography (Bradley and Short, 2005; Pasin and Bradley, 2015; Akoka et al., 2020; Waerzeggers, 2021; Laurie E Pearce et al., 2026).

We provide a solution to this problem that will enable consistent, large-scale date reconstruction by modeling the relationships between individuals and documents as a network. In this framework, documents are linked to one another through the people they mention, and individuals are anchored in time by the dated tablets in which they appear. An undated text can then be positioned within this web according to its connections with dated material. Framing prosopography as a network problem has two advantages: (a) it makes the logic of dating explicit and transparent, and (b) it allows the method to be scaled up consistently across thousands of tablets. Instead of relying on case-by-case judgments, the network structure provides a systematic way to estimate dates for large bodies of texts, while keeping the underlying assumptions open to inspection and adjustment.

Using this approach, we show that the majority of undated tablets can be assigned a plausible date range, typically within about a decade. First, we assess the reliability of the estimates by applying the method to tablets whose dates are preserved, treating them as if the date were missing and then comparing our results to the actual values. This validation confirms that the network-based approach can recover 99% of original dates with a useful degree of about 70% accuracy with a median deviation of 3 years, demonstrating its potential as a large-scale complement to traditional historical analysis. Following this validation, we use the method to provide estimated dates to 949 undated Neo-Babylonian documents, which are 89% of the total number of undated tablets. The efficacy of this method, however, is more prominent when looking at the archive level, where 4 tablet groups now have estimations for over 50% of their documents: the Išū-abūšu A from Borsippa, Šumu-iddin/Aplāya from Babylon, Sîn-uballiṭ from Ur, and Esagilāya B from Babylon.

An additional strength of this method lies in its iterative nature. Once an undated tablet has been assigned a plausible range, that estimate can be used to anchor the individuals it mentions more securely in time. In turn, those individuals provide new connections to other undated texts, extending the reach of the network. This process of refinement continues until no further links can be added, allowing the network to expand chronologically in a consistent and transparent way. The result is not only more dated texts, but also a richer, interconnected framework for understanding how these documents and individuals relate to one another.

More broadly, this study illustrates how computational methods can extend traditional scholarship in the humanities. Prosopography has long been central to reconstructing fragmentary historical records, but its application has been limited by scale and consistency. By reframing the problem as one of network analysis, we make it possible to apply the same reasoning systematically across thousands of texts, while keeping every step transparent and reproducible. The result is not only a practical tool for dating damaged cuneiform tablets, but also a demonstration of how digital approaches can turn scattered fragments into evidence that can be integrated into wider historical narratives. At the same time, adopting computational methods does more than provide answers to traditional questions: it opens new questions about what becomes visible at scale, how uncertainty

can be represented, and how ambiguous situations, such as a disputed regnal year or a co-regency, might be re-examined through the combined testimony of many documents.

1.1 Literature Review

To our knowledge, this is the first attempt to reconstruct dates with network-based prosopography on ancient historical documents. Dating is a fundamental task in historical research, when dealing with either texts or archaeological artifacts. While archaeological chronological modeling often relies on Bayesian statistics (Hamilton and Krus, 2018; Otárola-Castillo and Torquato, 2018; Dye et al., 2023), dating documents when a date is not written or is missing due to damage over time requires other methods.

Chronological attribution of ancient texts. Recent studies model chronological attribution directly from textual and palaeographic evidence at scale (Sommerschild et al., 2023, pp. 713–714). For Greek and Latin inscriptions, neural networks such as *Ithaca* and its Latin successor *Aeneas* learn to restore damaged passages while jointly performing geographical and chronological attribution. The latter also retrieves evidentiary parallels that ground predictions for historians. Reported dating errors range from < 30 years in Greek to ~ 13 years for Latin, and human–AI collaboration significantly improves expert performance (Assael, Sommerschild, Shillingford, et al., 2022; Assael, Sommerschild, Cooley, et al., 2025). In a complementary direction, Popović et al. (2025) trained *Enoch*, a style-based regressor built on letter-shape features anchored to radiocarbon measurements, which then propagates this physical calibration to previously undated Dead Sea Scrolls, yielding date probability distributions for manuscripts spanning the third century BCE to the first century CE and prompting several proposed redatings of key Jewish and Christian texts. The model achieves mean absolute errors of ~ 28 – 31 years, representing approximately 7–8% of the roughly 400-year temporal span of the Qumran corpus, a level of precision that is all the more striking given that the underlying training set comprised only twenty-four radiocarbon-dated samples. These works demonstrate that chronological signals can be captured from text and script alone when paired with absolute chronologies, and that even severely data-constrained models can produce historically meaningful resolution across a multi-century archive.

Generalization and domain robustness. On cuneiform tablets, recent work foregrounds the challenge of cross-collection generalization: visual/textual classifiers can reach $\sim 84\%$ accuracy when training and testing within the same museum collections. However, performance drops with regards to entire tablet collections completely excluded from training and validation and used only for testing, due to confounds and distribution shift.¹ Adversarial regularization and domain adaptation help but

¹For example, train on British Museum and Louvre tablets, then test on the Yale Babylonian Collection; revealing the impact of domain/collection shift.

do not fully close the gap, underscoring the need for bias-aware training and evaluation protocols in computational dating (Chen et al., 2024).

Kapon (2024) uses a larger dataset of 2D cuneiform tablet images from the same source (CDLI), to show that *tablet morphology* carries chronological signal at scale: models trained only on tablet silhouettes reach a macro- F_1 of $\sim 61\%$, and Variational Autoencoder (VAE) based explorations provide interpretable trajectories of shape change. Together, these results constitute the first large-scale empirical studies of visual and textual cuneiform dating at scale. They also motivate multi-view dating pipelines that fuse script/text and shape, evaluated under bias-aware, cross-collection protocols.

Chronology as an auxiliary objective. In Ancient Chinese, jointly predicting an era label alongside translation improves both tasks. This shows that even coarse dynastic-period classification from internal linguistic features improves translation quality, and is itself computationally recoverable with moderate accuracy, from the diachronic semantic properties of individual sentences (Chang et al., 2021).

Network analysis and chronological inference. Network analysis in historical studies has long been a pivotal methodology (Ahnert et al., 2021), and a recent thematic issue of the *Journal of Historical Network Research* placed *time* at the center of network reasoning (Petz et al., 2025). The editors establish the theoretical framework: historical networks are temporally diffuse and incomplete, and time cannot always be reduced to clean chronometric labels, requiring models that accommodate vagueness and uncertainty across multiple temporal registers. Of particular relevance is Thiery and Mees (2025), who use correspondence analysis and Allen’s interval algebra to establish the chronological ordering of Roman archaeological sites from material co-occurrence patterns in stamped pottery assemblages, formalizing the result as a directed graph structure. Their approach demonstrates that temporal sequence can be recovered from co-occurrence data alone, a principle that underlies the prosopographic logic of the present study.

Cuneiform social network analysis (SNA). The application of social network analysis (SNA) in cuneiform studies specifically has been growing rapidly in the past decade, moving from theoretical propositions to archive-specific case studies and, more recently, to corpus-wide computational approaches (Waerzeggers, 2014b; Gonçalves, 2021).² The closest use of networks for chronological reconstruction is Wagner et al. (2013) on the Murašû archive from Persian-period Nippur. They demonstrated that clustering documents purely by shared individuals—*without* using date information—can reproduce and clarify generational phases, implying that chronological signals are latent in

²See also Schmitz and L. Pearce (2013), Wagner et al. (2013), Laurie E. Pearce and Schmitz (2014), A. Anderson (2017), A. G. Anderson (2020), Kulikov et al. (2021), Waerzeggers (2021), Goddeeris (2022), Jones (2023), and Bennett et al. (2024)

an archive’s social fabric and that undated documents can be positioned probabilistically by network structure alone.

Positioning the present study: from network to chronological inference. Prosobab offers an extensive, machine-readable dataset of first-millennium individuals, attestations, and documents (Waerzeggers, Groß, et al., 2019). Building on this resource, Jauhiainen and Alstola (2025) have made an important contribution by extracting the dataset into SNA-ready formats, producing co-occurrence and directed social edge lists that open the corpus to network-level analysis. The present study departs from this dataset in a significant respect. Rather than adopting the Prosobab-derived edge lists directly, we reconstruct the prosopographic network independently from the underlying attestation data, applying a systematic pre-processing pipeline that identifies and resolves chronological inconsistencies, corrects erroneous date assignments, and eliminates duplicated attestations that would otherwise introduce noise into relational inference. These corrections are not merely procedural: unresolved chronological errors in an edge list propagate directly into any temporal analysis built upon it, rendering downstream dating inferences internally inconsistent. By constructing a source-critical network from the ground up, we ensure that the relational fabric encoding our chronological signal is itself historically coherent, a prerequisite for the probabilistic dating framework we develop here.

More fundamentally, our study breaks new ground in the computational dating of ancient documents. While prior work has demonstrated that textual, paleographic, and relational signals can carry chronological information within individual, well-delimited archives, no previous study has operationalized this insight into a scalable, empirically validated framework applicable to large, heterogeneous corpora drawn from multiple source archives, scribal traditions, and time periods simultaneously. We present the first such framework for Babylonian cuneiform: a network-based prosopographic approach that derives probabilistic date estimates for previously undated documents at corpus scale. In doing so, the present study transforms a methodological premise, that chronological signals are latent in the social fabric of archives, into a reproducible, bias-aware analytical procedure with demonstrable historical results.

2 Data

2.1 The Prosobab Dataset

The analysis in this paper is based on the Prosobab dataset (Waerzeggers, Groß, et al., 2019), an open-access relational database of cuneiform texts from first-millennium BCE Babylonia. The database is organized around three kinds of entities: tablets, attestations, and individuals.

- Tablet entries describe the documents themselves, with metadata including their type, provenance, and—if preserved—the original Babylonian date formula as well as its modern equivalent.
- Attestation entries record each occurrence of a person within a tablet, together with their function in the transaction.
- Individual entries represent disambiguated persons who may be attested in multiple texts, an interpretive step carried out by specialists on the Prosobab team. Each individual is assigned a unique identifier (PID), which links together all of their attestations and the tablets on which they occur. <https://prosobab.leidenuniv.nl/guide.php>.

This dual structure—raw attestations on the one hand, and interpretively linked individuals on the other—makes it possible to move between the immediate evidence of a single document and the reconstructed lifespans of the people named in it. It also provides the foundation for modeling connections between documents through shared individuals.

The corpus contains 21,121 unique individuals linked to 4,589 tablets through more than 38,606 recorded attestations. 3,522 tablets preserve a complete date, while the rest are only partially dated or undated. Figure 1 shows the chronological distribution of both pre-dated and our newly estimated tablets from the Prosobab dataset. The pre-dated tablets appear here after the corrections and validations described in Section 2.2. The distribution reflects the cleaned and internally consistent dataset on which all subsequent analyses are based. The earliest dated text is a purchase document from Uruk in 699 BCE, and the latest is a receipt from Ur in 317 BCE.




Figure 1: Distribution of existing dated (blue) and our newly estimated (orange) tablets in the Prosobab dataset. Source: Created by the author.

On average, each tablet mentions about 8.24 individuals, while each individual appears in approxi-

mately 1.79 tablets. The majority of people (about 82%) occur only once, whereas a smaller group of highly connected individuals appear in many texts, sometimes over extended periods.

This long-tail distribution - many short-lived, sparsely attested names and a few highly recurrent ones - creates the characteristic network structure that makes computational modeling both possible and necessary. The frequent reappearance of certain individuals provides chronological anchors that connect otherwise isolated documents, while the abundance of single attestations reflects the fragmentary and archival nature of the evidence.

These properties make the Prosobab dataset unusually well suited for testing large-scale prosopographic and network-based approaches. Its combination of dense local connections and long chronological coverage enables both micro- and macro-historical analyses, while the open-access format ensures full reproducibility of our results.

2.2 Pre-processing and Data Verification

The Prosobab dataset was systematically reviewed and cleaned to ensure that the working corpus is internally consistent and suitable for computational analysis. This process combined automated validation procedures with manual verification, within a fully documented and reproducible workflow. All scripts and correction logs are publicly available in the ProsoppyBase repository. This base repository contains the original data together with all implemented corrections (including scripts and logs) and functions as the authoritative, most up-to-date version from which this project and future projects derive their data. The data and results generated for the present estimation task are available in the ProsoppyEstimation repository.

2.2.1 Handling missing and duplicate entries

The dataset was first checked for missing values and structural inconsistencies. Rows with missing or invalid entries in key fields (such as individual identifiers or Julian dates) were removed. For cases where the Julian year contained two possible values separated by a slash (e.g., “535/534”), the first number of the two was systematically selected, providing a consistent representation while preserving chronological coherence.

Potential duplicate attestations were then investigated. Each attestation in Prosobab should represent a unique occurrence of a person in a specific tablet, but some apparent duplicates were found to differ only in auxiliary metadata such as the types and objects field, which is not relevant for chronological reconstruction. All tablets sharing the same identifier were verified to have identical years, confirming that no meaningful duplicates remained in the data used for date estimation.

2.2.2 Verification of chronological consistency

To detect inconsistencies in the dating information, all Babylonian date formulas were cross-checked against a trusted regnal chronology derived from the electronic Babylonian Library (eBL; <https://www.ebl.lmu.de/tools/date-converter>). Two complementary verification functions were applied. The first compared each recorded Julian year with the known reign of the ruler named in the tablet. This procedure identified a number of clear discrepancies - typically typographical errors such as “51” instead of “531” - which were logged in the file `ProsopyBase/data/corrections/initial_prosobab_inconsistencies_julian_date.csv`. A second, computational verification recalculated Julian dates directly from the Babylonian day-month-year formula, using the same regnal data as reference. This additional check captured subtler mismatches that still fell within a king’s reign but conflicted with the expected calendar conversion. Also, this check was used to fill in missing Julian years where the Babylonian date existed. These cases were reported in `ProsopyBase/data/corrections/julian_dates_calculator_mismatches.csv` and `ProsopyBase/data/corrections/babylonian_date_fills.csv`. Together, the two approaches ensured comprehensive coverage of both out-of-range and in-range discrepancies and to handling missing data.

The missing Julian years (i.e. where a babylonian date exists but the corresponding Julian year does not) were corrected automatically, but the inconsistent years (i.e. where both babylonian date and julian year exist but are not aligned) were documented for inspection and traceability. Each entry was compared against the corresponding record in the online Prosobab database and, where necessary, against the published transliterations of the original text. Only errors that could be confirmed unambiguously in the source material were corrected, following a conservative approach to data modification. The complete list of manual corrections is provided in `ProsopyBase/data/corrections/final-fixes.csv`.

3 Methods

3.1 Theoretical Upper Bounds for Estimations

The estimation process treats the relationships between individuals and documents as a network through which chronological information can propagate. Each tablet links together all the people named within it, forming a small cluster of co-occurring individuals. Across thousands of documents, these clusters generate a dense individual-individual co-occurrence network, in which people are connected whenever they appear on the same tablet. Chronological information - whether directly attested or inferred - spreads through this network along the same links that structure the social relationships recorded in the texts.

Because the estimation process relies on the network of individuals and documents, the connectivity

of that network itself sets hard limits on how far chronological information can travel. If a tablet is completely isolated from every dated source, no method—ours or anyone else’s - can recover its date through prosopography alone. To quantify these structural limits, we compute U_{iter} , the best-case ceiling representing the largest share of undated tablets that are theoretically reachable from the dated corpus through shared individuals. This upper bound provides a benchmark for evaluating how effectively the iterative estimation method approaches the dataset’s inherent potential.

To visualize this idea, one can imagine laying out every tablet and every person as points, and drawing a line whenever a person is mentioned on a tablet. The result is a network that naturally divides into small “islands” (connected components) of connected tablets and individuals - groups in which every element is linked, directly or indirectly, to all the others, but disconnected from the rest of the corpus. If a component includes at least one dated tablet, chronological information can, in principle, circulate within it: from the dated document to the people it names, and from those people to other tablets in which they appear. Components composed entirely of undated tablets, however, lack any such entry point and remain chronologically isolated. The measure U_{iter} quantifies this structural property by identifying the share of undated tablets that belong to components containing at least one dated document. We refer to this upper bound as U_{iter} - short for iterative upper bound - because it represents the limit that would be reached if the iterative estimation procedure were allowed to propagate chronological information perfectly through the network. This proportion thus defines the theoretical ceiling on what any people-based, iterative approach can date, given the way the sources are interconnected.

Formally, we define the measure U_{iter} as follows. Let the dataset be represented as a bipartite graph:

$$(G = T, P, E)$$

where T is the set of tablets, P the set of individuals and $E \subseteq T \times P$ the set of attestations linking people to the tablets on which they appear. This graph can be decomposed into a set of connected components:

$$C = \{C_1, C_2, \dots, C_n\}$$

each representing a distinct cluster of tablets and individuals connected through shared appearances. Let T^{dated} and $T^{undated}$ denote the subsets of dated and undated tablets, respectively. For each component C_i define an indicator function:

$$I(C_i) = \begin{cases} 1 & \text{if } C_i \cap T^{dated} \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

That is, $I(C_i) = 1$ if the component contains at least one dated tablet and 0 otherwise. The upper




Figure 2: U_{iter} **for undated tablets**. In this example, five tablets are undated (dashed), but only three of them belong to a connected component that contains at least one dated tablet (here T_1), Only these three undated tablets can, in principle, receive chronological information through the co-occurrence network; the others are structurally isolated. Consequently, the theoretical upper bound for any prosopographic dating method on this network is $3/5=60\%$ of the undated corpus. Component B, which contains only undated tablets and no link to any dated document, is unreachable by any method based solely on shared individuals. **For the dated tablets** only component C contains at least two dated tablets so that when leaving-one-out another the other tablet is used as anchor. In total the U_{iter} for the dated tablets in this example is $2/3=67\%$. See section 4.1 for precise formula and calculation. Source: Created by the author.

bound U_{iter} is then given by:

$$U_{iter} = \frac{\sum_{i=1}^n |C_i \cap T^{\text{undated}}| I(C_i)}{|T^{\text{undated}}|}.$$

This expression calculates the proportion of undated tablets that belong to components containing at least one dated tablet. In other words, U_{iter} measures the share of undated documents that are theoretically reachable from the dated corpus through shared individuals, representing the best-case coverage any iterative, people-based estimation could achieve given the network's structure.

Applying this measure to the cleaned Prosobab dataset yielded a value of $U_{iter} = 0.9503$, meaning that approximately 95% of the undated tablets are, in principle, reachable from the dated corpus through shared individuals. In absolute terms, this corresponds to 1,014 out of 1,067 undated documents that lie in network components containing at least one dated tablet. The remaining 5% of undated tablets belong to components entirely disconnected from the dated evidence and can therefore never be dated through prosopographic links alone. This result defines the structural ceiling for our estimation

approach: even under ideal conditions of information propagation, no more than about 95% of the undated corpus can be assigned a date without additional external evidence. It is important to emphasize, however, that this represents a highly permissive upper limit. In this calculation, we imposed no restrictions on which individuals could serve as anchors (see Section 4.3), constraints that are essential in the actual estimation procedure to maintain historical plausibility and chronological precision.

3.2 The Babylonian Yellow Pages

The *Babylonian Yellow Pages* functions as a methodological reference tool for organizing and updating the chronological information associated with each individual. It is not part of the network itself, but an external, dynamically evolving lookup table that records temporal attributes for person-nodes: their earliest and latest attested years of activity, and later, their estimated years as these become available through iterations. As the estimation progresses, the *Babylonian Yellow Pages* is continuously updated, reflecting the growing body of chronological information inferred from the network. For each undated tablet, the individuals it mentions are matched against the current version of the *Babylonian Yellow Pages* to retrieve their temporal attributes. Only those whose activity spans fall below a specified maximum are retained as valid anchors, ensuring that the estimation relies on individuals whose lifespans can meaningfully constrain a date. The algorithm then combines the earliest and latest years of all valid anchors to calculate an estimated date range for the tablet, based on a weighted average of the chronological intervals associated with those individuals.

The estimation proceeds iteratively. After each round, **newly estimated** tablets are used to update the *Babylonian Yellow Pages*. These entries are clearly labeled as estimated and function as a second-tier source of evidence, used only after all directly attested information has been incorporated. This hierarchical design prevents circular reasoning, since inferred dates are never used to revise or justify the estimates from which they originated. Once a tablet has received an estimated date, it is not reprocessed in subsequent iterations. Instead, the newly estimated tablets expand the set of individuals represented in the Yellow Pages: people who had previously lacked temporal information now acquire estimated activity spans. These additional person-level entries, in turn, allow further undated tablets to be evaluated in the next iteration. The process continues until a complete pass through the corpus yields no new estimates or new individuals to add to the Yellow Pages, at which point it is considered to have converged. The final output consists of an updated *Babylonian Yellow Pages*, containing both attested and estimated temporal attributes for individuals, and a corresponding table of cuneiform tablets with newly assigned date ranges.

3.3 Why use weighted average?

These values are obtained through a weighted average, where each individual's contribution is proportional to the number of documents in which they appear. Formally, let N denote the number of individuals in the set, and for each individual i , let n_i be the number of documents in which they are attested, and $[L_i, U_i]$ the lower and upper bounds of their recorded activity span. The estimated lower and upper bounds are then computed as

$$\hat{L} = \frac{\sum_{i=1}^N n_i L_i}{\sum_{i=1}^N n_i}, \quad \hat{U} = \frac{\sum_{i=1}^N n_i U_i}{\sum_{i=1}^N n_i}.$$

The resulting estimated interval is therefore given by $\hat{I} = [\hat{L}, \hat{U}]$.

This weighting reflects a pragmatic assumption: people who are attested in many documents provide more stable chronological information. In the Babylonian onomastic system - where personal identification typically includes the individual's name, patronymic, and often family or professional affiliation - true homonymy is less common. As a result, the weighting does not primarily reflect uncertainty of identification, but rather confidence in the reliability of an individual's attested record.

However, this choice involves an inherent trade-off. Individuals who appear in many documents tend also to have longer recorded activity spans, simply because repeated appearances over time extend their temporal coverage. Such individuals thus gain greater influence in the weighted average while simultaneously broadening the estimated range. The algorithm therefore balances two opposing tendencies - greater reliability from frequent attestations versus reduced precision from extended lifespans. The use of a weighted average offers a simple, transparent, and fully reproducible way to formalize this balance, while remaining close to the intuitive reasoning of traditional prosopographic analysis. Nevertheless, alternative formulations - such as weighting by the temporal overlap of activity spans, by the density of intermediate attestations, or through probabilistic inference - could be explored in future work to optimize how chronological information is aggregated within the network.

3.4 The activity span constraint

A crucial parameter in this framework is the maximum allowed activity span for individuals included as anchors. If the threshold is set too low, many tablets cannot be dated because they lack qualifying individuals; if set too high, long-lived anchors inflate the estimated ranges and diminish chronological precision. To find a workable balance, the algorithm incorporates a tiered fallback procedure. When no valid anchors remain under the initial threshold X , the constraint is automatically relaxed to $2X$, and, if still insufficient, to $3X$, ensuring that tablets connected only through longer-lived individuals can still be considered without halting the estimation process entirely. To evaluate the overall effect of

these thresholds, we ran the estimation algorithm across several span limits (1 to 30 years). For each run, we recorded both the proportion of undated tablets that could be dated and the average width of their estimated ranges. The results revealed a clear trade-off: increasing the threshold expanded coverage but rapidly degraded precision. A limit of 17 years produced the most balanced outcome, yielding a large number of estimable tablets while maintaining average date ranges of roughly a decade—narrow enough to be meaningful for historical interpretation. The quantitative details of this trade-off are discussed in greater depth in Section 4.3.

4 Results

4.1 Validation upper bounds

To assess the efficacy of the method, we first validate whether we can successfully date tablets in which the dates are preserved - in other words, using ground truth data. We first determine how many dated tablets are structurally testable. Because the validation procedure relies on prosopographic connections, a dated tablet can only be validated if it shares individuals with at least one other dated tablet. To formalize this requirement, we compute a validation-specific version of the upper bound U_{iter} . Representing the dated portion of the corpus as a bipartite graph of tablets and individuals, we identify its connected components $C_1, \dots, C_n$ and define a component as verifiable if it contains at least **two** dated tablets—that is, if withholding the date of one tablet still leaves chronological information available within its component. Let T^{dated} denote the set of all fully dated tablets. For each component C_i we define the indicator

$$I_{val}(C_i) = \begin{cases} 1, & \text{if } |C_i \cap T^{dated}| \geq 2, \\ 0, & \text{otherwise.} \end{cases}$$

The validation upper bound is then given by

$$U_{iter}^{dated} = \frac{\sum_{i=1}^n |C_i \cap T^{dated}| I_{val}(C_i)}{|T^{dated}|}. \quad (1)$$

This quantity measures the proportion of fully dated tablets that are, in principle, recoverable under leave-one-out conditions. Out of 3,522 fully dated tablets, 3,517 (99.86%) satisfy this criterion, establishing the structural maximum for any prosopography-based validation procedure.

4.2 Validation results

With this context, we performed a full leave-one-out validation experiment. In each iteration, a single dated tablet was temporarily treated as undated, and the algorithm attempted to estimate its date using all remaining dated evidence. To prevent leakage, all temporal attributes associated with the target tablet were removed from the *Babylonian Yellow Pages* prior to estimation, ensuring that none of its own attested information contributed to its predicted range. This procedure directly tests the method’s ability to reconstruct known chronological information based solely on prosopographic links.

Under this configuration, the algorithm produced verifications for 3,473 dated tablets, leaving only 49 (1.4%) without sufficient anchors. Among the verified cases, 2,404 were successfully recovered, while 1,069 diverged from their true dates. The deviations for the correct verification are very narrow with a 2-year median. This corresponds to a success rate of 69.22% (2404/3473) across the entire verifiable corpus (3,473 tablets).

Among the 1,069 incorrectly verified tablets, the median deviation is 7.5 years, which very likely stem from probable issues with the disambiguation of individuals in the original dataset. This raises the need for further inquiry of these tablets on a case-by-case basis.

Taken together, these results show that the estimation procedure performs close to the structural ceiling defined by $U_{iter}^{dated}=0.9875$ (3474/3517) for the dated subset. The method successfully recovers a substantial majority of known dates with chronologically meaningful accuracy and demonstrates robustness across a wide variety of prosopographic configurations. This provides strong empirical support for its application to undated tablets under historically motivated constraints.

4.3 Results of Tablet Estimation

Applying the iterative estimation algorithm with a maximum activity-span threshold of 17 years produced estimated dates for 949 of the 1,067 undated tablets in the corpus, corresponding to 89% coverage from the undated corpus. The process converged after two iterations, with successive rounds estimating 943 and 6 additional tablets before no further updates were possible. The average width of the estimated date ranges was 12.43 years, indicating that the majority of reconstructed ranges fall within a chronologically meaningful span for historical interpretation. Relative to the structural limit defined by $U_{iter} = 0.95$ (1014/1067), the method recovered 93.6% (949/1014) of all undated tablets that are theoretically reachable through prosopographic links, demonstrating that the algorithm effectively saturates the portion of the corpus that can be dated under realistic constraints.

The threshold experiment further illustrates the balance between coverage and precision. For example, a stricter limit of 10 years yielded 812 estimates (80% coverage) with an average range of 6.16




Figure 3: The threshold experiment showing the tradeoff between precision in terms of average number of years per estimation and the number of tablets that can be estimated at all. This threshold is done solely on the undated tablets to establish which activity span should serve as cut-off. Source: Created by the author.

years, while a more permissive threshold of 20 years increased coverage to 98.22% (996 estimates) but broadened the average range to 15.37 years. We believe that 17-years-max constraint provides a good compromise between coverage and precision, which is useful for historical, political, economic, and administrative analyses. Nonetheless, the framework we suggest allows the flexibility of choosing different parameters in future iterations.

5 Discussion and Conclusions

The results demonstrate that large-scale date estimation based on prosopographic connections is not only feasible but methodologically tractable when approached through a network framework. By representing the relationships among tablets and individuals as a structured graph, it becomes possible to quantify the reach of chronological information and to reconstruct temporal contexts for thousands of documents that were previously undatable. The method approaches the theoretical dating potential of the corpus — as measured by the U_{iter} upper bounds — while maintaining a level of precision that remains historically meaningful. The following discussion considers the implications of these results along two complementary dimensions: their contribution to the historical understanding of first-millennium Babylonian society, and their methodological significance for computational approaches to ancient documentary corpora.

Historical contribution

The most immediate historical result is the systematic recovery of temporal context for a substantial body of previously undatable cuneiform documentation. Dates in the ancient Near East are among the most consequential pieces of metadata for any text: they situate transactions, relationships, and institutional practices within the specific political and economic conjunctures of Babylonian history. By making this temporal dimension computationally recoverable at scale, the present study opens large portions of the corpus to historical analysis that could not previously proceed beyond a general period attribution.

The error analysis simultaneously yields a secondary historical dividend. For example, the 102-year deviation observed for Nbk 42 (Prosobab Tablet ID 7504), dated to Julian year 600, but has a network-based estimate of approximately 498, reveals that the individual PID 3034 has been assigned to two distinct persons: Amurru-šamā' and Ana-Bēl-upāqu, whose attestation patterns fall in different temporal ranges. This is not a failure of the estimation procedure but a diagnostic output: the very magnitude of the deviation flags a prosopographic ambiguity in the underlying dataset that would otherwise have remained invisible. Computational date estimation thus functions here not merely as a dating tool but as a source-critical instrument, surfacing latent inconsistencies in expert-curated prosopographic records and generating new hypotheses about individual identity for subsequent philological investigation.

The network framework also opens productive avenues for the social history of the Babylonian documentation. Preliminary analysis reveals, for instance, that Nabû-zēru-iqīša (PID 19678, attested 12.III.11 Nbn – 12.X.06 Dar I) participates as a witness with a frequency and a thematic specificity, clustering notably around transactions involving silver (*kaspu*). This distinguishes him from the general population of documentary witnesses. Whether this reflects a form of professional specialization, a position within an institutional network, or a particular relationship to creditors and debtors in the silver economy remains to be determined. These questions will be pursued in a forthcoming study devoted to the social properties of the witness network — including patterns of geographic mobility, frequency of appearance, and archival affiliation — as a window onto specialized roles within Babylonian legal and economic practice.

Digital humanities contribution

Underlying both the dating results and the error-detection outcomes is a methodological foundation that deserves explicit recognition: the systematic, source-critical pre-processing of the prosopographic data. Existing datasets such as Prosobab represent years of expert manual curation and are indispensable resources for their field; yet precisely because they aggregate knowledge accumulated over decades and across multiple scholarly conventions, they inevitably contain structural inconsistencies

that are difficult to detect through manual inspection alone: chronological assignments, regnal year conversions, and the resolution of ambiguous individual identities. In the present study, automated checks for chronological coherence, cross-validation of date fields, and consistent parsing of attestation records were indispensable prerequisites for meaningful network construction. Pre-processing is therefore not a peripheral technical step but an epistemological one: it determines whether the relational fabric encoding the chronological signal is historically coherent in the first place.

The network framework itself contributes a further methodological instrument in the form of the U_{iter} upper bound, which quantifies the theoretical limit of what can be inferred from prosopographic co-occurrence alone. By making this ceiling explicit, the study introduces a principled measure of epistemic transparency. It specifies where the network contains sufficient relational density to sustain temporal inference and where human expertise, such as archival knowledge, philological judgment, or contextual interpretation, remains irreplaceable. Computational methods and traditional scholarship are thus positioned not as alternatives but as complementary operations with distinct and quantifiable domains of competence.

Future development will extend this framework toward a knowledge-graph representation of the corpus, in which individuals, institutions, places, roles, and transaction types are treated as typed entities connected by semantically differentiated relationships. Within such a structure, date estimation would become one layer of a broader inference system, enabling queries that integrate temporal, spatial, social, and institutional dimensions simultaneously. These can include reasoning about the succession of offices within a family lineage, the geographic diffusion of scribal practices, or the institutional mediation of property transmission across generations. Such a direction connects the project to wider developments in digital humanities toward interoperable, semantically enriched representations of historical knowledge.

Broader impact

Taken together, the historical and methodological dimensions of this study illustrate a model for how computational reasoning and traditional historical scholarship can be mutually reinforcing rather than methodologically competitive. The reconstruction of dates at scale for a multi-archive, multi-century corpus of cuneiform documentation demonstrates that chronological inference need not depend on explicit date formulae or palaeographic intuition alone: it can be derived from the relational structure of the social networks encoded in the documents themselves. At the same time, the study shows that the credibility of such inference rests entirely on the rigor of the evidentiary base. Computational precision is only as reliable as the source-critical work that precedes it. In formalizing prosopographic reasoning as a network problem, the present study makes explicit the structures through which historical knowledge is inferred, the assumptions on which it rests, and the boundaries within which it operates. Beyond the specific case of Babylonian documentation, this constitutes a transferable

framework: a demonstration that large-scale, bias-aware chronological estimation from relational evidence is achievable across ancient documentary corpora, and that its outputs are not only statistically tractable but historically productive.

References

- Ahnert, Ruth, Sebastian E. Ahnert, Catherine Nicole Coleman, and Scott B. Weingart (2021). *The Network Turn: Changing Perspectives in the Humanities*. Elements in Publishing and Book Culture. Cambridge: Cambridge University Press. ISBN: 978-1-108-79190-8. DOI: 10.1017/9781108866804. URL: <https://www.cambridge.org/core/elements/network-turn/CC38F2EA9F51A6D1AFCB7E005218BBE5>.
- Akoka, Jacky, Isabelle Comyn-Wattiau, Stéphane Lamassé, and Cédric Du Mouza (2020). “Contribution of Conceptual Modeling to Enhancing Historians’ Intuition - Application to Prosopography”. en. In: *Conceptual Modeling*. Ed. by Gillian Dobbie, Ulrich Frank, Gerti Kappel, Stephen W. Liddle, and Heinrich C. Mayr. Vol. 12400. Series Title: Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 164–173. ISBN: 978-3-030-62521-4 978-3-030-62522-1. DOI: 10.1007/978-3-030-62522-1_12. URL: http://link.springer.com/10.1007/978-3-030-62522-1_12 (visited on 10/24/2023).
- Anderson, Adam (Aug. 2017). “The Old Assyrian Social Network: an analysis of the text from Kültepe-Kanesh (1950-1750 B.C.E.)” PhD thesis. DOI: 10.13140/RG.2.2.29787.92967.
- Anderson, Adam G. (2020). “Death of an Old Assyrian Salesman”. en. In: *Chatreššar 11 / 2020*. Accepted: 2020-11-30T10:09:15Z, pp. 5–34. URL: <https://dspace.cuni.cz/handle/20.500.11956/123610> (visited on 03/02/2026).
- Assael, Yannis, Thea Sommerschild, Alison Cooley, Brendan Shillingford, John Pavlopoulos, Priyanka Suresh, Bailey Herms, Justin Grayston, Benjamin Maynard, Nicholas Dietrich, Robbe Wulgaert, Jonathan Prag, Alex Mullen, and Shakir Mohamed (Sept. 2025). “Contextualizing ancient texts with generative neural networks”. en. In: *Nature* 645.8079, pp. 141–147. ISSN: 1476-4687. DOI: 10.1038/s41586-025-09292-5. URL: <https://www.nature.com/articles/s41586-025-09292-5> (visited on 03/01/2026).
- Assael, Yannis, Thea Sommerschild, Brendan Shillingford, Mahyar Bordbar, John Pavlopoulos, Marita Chatzipanagiotou, Ion Androutsopoulos, Jonathan Prag, and Nando de Freitas (Mar. 2022). “Restoring and attributing ancient texts using deep neural networks”. en. In: *Nature* 603.7900, pp. 280–283. ISSN: 1476-4687. DOI: 10.1038/s41586-022-04448-z. URL: <https://www.nature.com/articles/s41586-022-04448-z> (visited on 03/01/2026).
- Bennett, Ellie, Lena Tambs, and Krister Lindén (Sept. 2024). “Letters have weight: weighted k-shells in a Neo-Assyrian co-attestation network”. en. In: *Journal of Historical Network Research*

- 10.1, pp. 150–197. ISSN: 2535-8863. DOI: 10.25517/jhnr.v10i1.95. URL: <https://jhnr.net/articles/10.25517/jhnr.v10i1.95> (visited on 03/02/2026).
- Bongenaar, A. C. V. M. (1997). *The Neo-Babylonian Ebbabar Temple at Sippar: its Administration and its Prosopography*. Type: Book. Leiden.
- Bradley, John and Harold Short (Jan. 2005). “Texts into Databases: The Evolving Field of New-style Prosopography”. en. In: *Literary and Linguistic Computing* 20.Suppl, pp. 3–24. ISSN: 0268-1145. DOI: 10.1093/llc/fqi022. URL: <https://academic.oup.com/dsh/article/20/Suppl/3/1027872> (visited on 06/13/2019).
- Chang, Ernie, Yow-Ting Shiue, Hui-Syuan Yeh, and Vera Demberg (Aug. 2021). “Time-Aware Ancient Chinese Text Translation and Inference”. In: *Proceedings of the 2nd International Workshop on Computational Approaches to Historical Language Change 2021*. Ed. by Nina Tahmasebi, Adam Jatowt, Yang Xu, Simon Hengchen, Syrielle Montariol, and Haim Dubossarsky. Online: Association for Computational Linguistics, pp. 1–6. DOI: 10.18653/v1/2021.lchange-1.1. URL: <https://aclanthology.org/2021.lchange-1.1/> (visited on 03/02/2026).
- Chen, Danlu, Jiahe Tian, Yufei Weng, Taylor Berg-Kirkpatrick, and Jacobo Myerston (Aug. 2024). “Classification of Paleographic Artifacts at Scale: Mitigating Confounds and Distribution Shift in Cuneiform Tablet Dating”. In: *Proceedings of the 1st Workshop on Machine Learning for Ancient Languages (ML4AL 2024)*. Ed. by John Pavlopoulos, Thea Sommerschild, Yannis Assael, Shai Gordin, Kyunghyun Cho, Marco Passarotti, Rachele Sprugnoli, Yudong Liu, Bin Li, and Adam Anderson. Hybrid in Bangkok, Thailand and online: Association for Computational Linguistics, pp. 30–41. DOI: 10.18653/v1/2024.ml4al-1.4. URL: <https://aclanthology.org/2024.ml4al-1.4/> (visited on 03/01/2026).
- Dye, Thomas S., Caitlin E. Buck, Robert J. DiNapoli, and Anne Philippe (May 2023). “Bayesian chronology construction and substance time”. In: *Journal of Archaeological Science* 153, p. 105765. ISSN: 0305-4403. DOI: 10.1016/j.jas.2023.105765. URL: <https://www.sciencedirect.com/science/article/pii/S0305440323000432> (visited on 02/26/2026).
- Goddeeris, Anne (2022). “Women and Their Weight: Incorporating Weighted Edges in a Network Analysis of the Central Redistributive Household of Nippur (Eighteenth Century BCE)”. eng. In: *The mummy under the bed: essays on gender and methodology in the ancient Near East*. Ed. by Katrien De Graef, Agnès Garcia-Ventura, Anne Goddeeris, and Beth Alpert Nakhai. wEdge 1. Münster: Zaphon, pp. 369–406. ISBN: 978-3-96327-088-8.
- Gonçalves, Carlo (2021). “Social network analysis and kinship in the old babylonian diyala. Fathers and sons in the archive of nūr-šamaš”. In: *H2D/ Digital Humanities Journal* 3.1. DOI: <https://doi.org/10.21814/H2D.3470>.
- Gütschow, C. (2012). *Methoden zur Restaurierung von ungebrannten und gebrannten Keilschrifttafeln - Gestern und Heute*. Berliner Beiträge zum Vorderen Orient. Type: Book. Gladbeck: PeWe.

- Hamilton, W. Derek and Anthony M. Krus (Apr. 2018). “The Myths and Realities of Bayesian Chronological Modeling Revealed”. en. In: *American Antiquity* 83.2, pp. 187–203. ISSN: 0002-7316, 2325-5064. DOI: 10.1017/aaq.2017.57. URL: <https://www.cambridge.org/core/journals/american-antiquity/article/myths-and-realities-of-bayesian-chronological-modeling-revealed/F02B247E777529F4500181216C307284> (visited on 02/26/2026).
- Jauhiainen, Heidi and Tero Alstola (Jan. 2025). “Social Networks of Prosobab”. en. In: *Journal of Open Humanities Data* 11, p. 56. DOI: 10.5334/johd.384. URL: <https://openhumanitiesdata.metajnl.com/articles/10.5334/johd.384> (visited on 11/06/2025).
- Jones, Christopher W. (Dec. 2023). “The Decline and Fall of the Assyrian Court Scholar: A Social Network-Based Examination”. en. In: *Avar: An Interdisciplinary Journal of Life and Society in the Ancient Near East* 2.2, pp. 297–354. ISSN: 2752-3535. DOI: 10.33182/aijls.v2i2.2839. URL: <https://avarjournal.com/avar/article/view/2839> (visited on 03/02/2026).
- Jursa, Michael (2005). *Neo-Babylonian Legal and Administrative Documents. Typology, Contents and Archives*. Guides to the Mesopotamian Textual Record 1. Type: Book. Münster: Ugarit-Verlag.
- (2010). *Aspects of the economic history of Babylonia in the first millennium BC: economic geography, economic mentalities, agriculture, the use of money and the problem of economic growth*. Alter Orient und Altes Testament 377. Type: Book. Münster: Ugarit-Verlag.
- Kapon, Danielle (July 2024). *Shaping History: Advanced Machine Learning Techniques for the Analysis and Dating of Cuneiform Tablets over Three Millennia - Datasets*. eng. DOI: 10.5281/zenodo.12787746. URL: <https://zenodo.org/records/12787746> (visited on 12/31/2025).
- Kleber, Kristin (Jan. 2013). “Famine in Babylonia: A microhistorical approach to an agricultural crisis in 528–526 BC”. en. In: *Zeitschrift für Assyriologie und Vorderasiatische Archäologie* 102.2, pp. 219–244. ISSN: 1613-1150. DOI: 10.1515/za-2012-0013. URL: <https://www.degruyterbrill.com/document/doi/10.1515/za-2012-0013/html?lang=en&srsltid=AfmB0opC6iQ2Q6yCNYhytkJz-1MQh60HcirVjrBSGzoGNnMcBmzFG> (visited on 02/04/2026).
- Kulikov, Anya, Adam Anderson, and Niek Veldhuis (July 2021). “Sumerian Networks: Classifying Text Groups in the Drehem Archives”. en. In: *IDEAH*. DOI: 10.21428/f1f23564.6da394df. URL: <https://ideah.pubpub.org/pub/q228591x> (visited on 11/01/2021).
- Kümmel, Hans Martin (1979). *Familie, Beruf und Amt im spätbabylonischen Uruk*. Abhandlungen der Deutschen Orient-Gesellschaft. Type: Book. Berlin: Gebr. Mann Verlag.
- Otárola-Castillo, Erik and Melissa G. Torquato (2018). “Bayesian Statistics in Archaeology”. In: *Annual Review of Anthropology* 47, pp. 435–453. ISSN: 0084-6570. URL: <https://www.jstor.org/stable/48550916> (visited on 02/26/2026).
- Parker, Richard Anthony and Waldo Herman Dubberstein (1971). *Babylonian Chronology 626 B.C.-A.D. 75*. en. Brown University Studies 19. Google-Books-ID: WExmAAAAMAAJ. Providence, R.I.: Brown University Press. ISBN: 978-0-598-21298-6.

- Pasin, Michele and John Bradley (Apr. 2015). “Factoid-based prosopography and computer ontologies: towards an integrated approach”. en. In: *Literary and Linguistic Computing* 30.1, pp. 86–97. ISSN: 0268-1145. DOI: 10.1093/llc/fqt037. URL: <https://academic.oup.com/dsh/article/30/1/86/352888> (visited on 06/13/2019).
- Pearce, Laurie E, Saana Svärd, Caroline Waerzeggers, Tero Alstola, Adam Anderson, Heather D. Baker, Elena Devecchi, Carlos Gonçalves, Shai Gordin, Melanie Groß, Heidi Jauhainen, John P. Nielsen, Émilie Pagé-Perron, Seth Richardson, Erica Scarpa, and Jonathan Tenney (2026). “White paper on prosopographical studies in the Ancient Near East”. In: *Cuneiform Digital Library Journal*.
- Pearce, Laurie E. and Patrick Schmitz (2014). “Berkeley Prosopography Services”. In: *ISAW Papers* 7.1. URL: <https://dlib.nyu.edu/awdl/isaw/isaw-papers/7/pearce-schmitz/>.
- Petz, Cindarella, Aline Deicke, and Julian Jarosch (Dec. 2025). “Foreword: Time and Temporality as Categories of Connectedness”. en. In: *Journal of Historical Network Research* 12.1, pp. 1–11. ISSN: 2535-8863. DOI: 10.25517/jhnr.v12i1.122. URL: <https://jhnr.net/articles/10.25517/jhnr.v12i1.122> (visited on 02/04/2026).
- Podany, Amanda H. (Dec. 2022). *Weavers, Scribes, and Kings: A New History of the Ancient Near East*. Oxford: Oxford University Press. ISBN: 978-0-19-005904-0.
- Popović, Mladen, Maruf A. Dhali, Lambert Schomaker, Johannes van der Plicht, Kaare Lund Rasmussen, Jacopo La Nasa, Ilaria Degano, Maria Perla Colombini, and Eibert Tigchelaar (June 2025). “Dating ancient manuscripts using radiocarbon and AI-based writing style analysis”. en. In: *PLOS ONE* 20.6, e0323185. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0323185. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0323185> (visited on 03/01/2026).
- Radner, K. (1998). *The Prosopography of the Neo-Assyrian Empire, Volume 1, Part 1: A*. Type: Edited Book. Helsinki.
- Schmitz, Patrick and Laurie Pearce (Sept. 2013). “Berkeley prosopography services: ancient families, modern tools”. In: *Proceedings of the 1st International Workshop on Collaborative Annotations in Shared Environment: metadata, vocabularies and techniques in the Digital Humanities*. DH-CASE ’13. New York, NY, USA: Association for Computing Machinery, pp. 1–8. ISBN: 978-1-4503-2199-0. DOI: 10.1145/2517978.2517980. URL: <https://dl.acm.org/doi/10.1145/2517978.2517980> (visited on 03/02/2026).
- Sommerschield, Thea, Yannis Assael, John Pavlopoulos, Vanessa Stefanak, Andrew Senior, Chris Dyer, John Bodel, Jonathan Prag, Ion Androutsopoulos, and Nando De Freitas (Sept. 2023). “Machine Learning for Ancient Languages: A Survey”. en. In: *Computational Linguistics* 49.3, pp. 703–747. ISSN: 0891-2017, 1530-9312. DOI: 10.1162/coli_a_00481. URL: <https://direct.mit.edu/coli/article/49/3/703/116160/Machine-Learning-for-Ancient-Languages-A-Survey> (visited on 05/24/2024).

- Stern, Sacha (2012). *Calendars in antiquity: empires, states, and societies*. eng. 1st ed. Oxford New York: Oxford University Press. ISBN: 978-0-19-174617-8 978-0-19-162622-7.
- Stone, L. (1971). “Prosopography”. In: *Daedalus* 100. Type: Journal Article, pp. 46–79.
- Thiery, Florian and Allard W. Mees (Dec. 2025). “Dating Dated Sites: Using Correspondence Analysis to handle Chronologies as Graphs”. en. In: *Journal of Historical Network Research* 12.1, pp. 166–196. ISSN: 2535-8863. DOI: 10.25517/jhnr.v12i1.105. URL: <https://jhnr.net/articles/10.25517/jhnr.v12i1.105> (visited on 02/04/2026).
- Waerzeggers, Caroline (2014a). *Marduk-rēmanni. Local Networks and Imperial Politics in Achaemenid Babylonia*. Orientalia Lovaniensia Analecta 233. Type: Book. Leuven: Peeters.
- (2014b). “Social network analysis of cuneiform archives: A new approach”. In: *Documentary Sources in Ancient Near Eastern and Greco-Roman Economic History: Methodology and Practice*. Ed. by H. D. Baker and M. Jursa. Type: Book Section. Oxford: Oxbow Books, pp. 207–233. ISBN: 1-78297-758-9.
- (Dec. 2021). “Digital prosopography of Babylonia: new horizons”. In: *Bridging the gap: disciplines, times, and spaces in dialogue - Volume 1: sessions 1, 2, and 5 from the conference Broadening Horizons 6 held at the Freie Universität Berlin, 24-28 June 2019*. Ed. by C. W. Hess and F. Manuelli. Oxford: Archaeopress, pp. 81–96. URL: <https://hdl.handle.net/1887/3281616>.
- Waerzeggers, Caroline, Melanie Groß, et al. (2019). *Prosobab: Prosopography of babylonia (c. 620-330 BCE)*.
- Waerzeggers, Caroline, Melanie Groß, and et al. (2019). *Prosobab: Prosopography of Babylonia (c. 620-330 BCE)*. URL: <https://prosobab.leidenuniv.nl> (visited on 02/20/2021).
- Wagner, Allon, Yuval Levavi, Siram Kedar, Kathleen Abraham, Yoram Cohen, and Ran Zadok (2013). “Quantitative Social Network Analysis (SNA) and the Study of Cuneiform Archives: A Test-case based on the Murašû Archive”. English. In: *Akkadica* 134.2, pp. 117–134. ISSN: 1378-5087. URL: <https://cris.iucc.ac.il/en/publications/quantitative-social-network-analysis-sna-and-the-study-of-cuneifo> (visited on 10/29/2024).
- Zadok, Ran (2002). “Contributions to Babylonian Geography, Prosopography and Documentation”. In: *Ex Mesopotamia et Syria lux. Festschrift für Manfred Dietrich zu seinem 65. Geburtstag*. Ed. by O. Loretz, K. A. Metzler, and H. Schaudig. AOAT. Type: Book Section. Münster.
- (2003). “The Representation of Foreigners in Neo-and Late-Babylonian Legal Documents (Eighth through Second Centuries B.C.E.)” In: *Judah and the Judeans in the Neo-Babylonian Period*. Ed. by Oded Lipschits and Jospheh Blenkinsopp. Type: Book Section. Winona Lake, Indiana: Eisenbrauns, pp. 471–589.
- (2018). *A prosopography of the Israelites in Old Testament traditions: a contextualized handbook*. English. Archaeological Center Publications. ISBN: 978-965-7162-23-1. (Visited on 01/14/2024).