

AŠŠUR AND HIS FRIENDS: A STATISTICAL ANALYSIS OF NEO-ASSYRIAN TEXTS

*Tero Alstola, Shana Zaia, Aleksi Sahala, Heidi Jauhiainen, Saana Svärd,
and Krister Lindén (University of Helsinki)*

Abstract

Large digital datasets of cuneiform sources lend themselves to computational analysis that can complement and improve upon traditional philological work. The present article applies social network analysis to an electronic corpus of 1,532 texts to study the god Aššur and his position in divine networks in the Neo-Assyrian period. Our results show that the performance of social network analysis can be improved by using a small window size and calculating tie strengths with pointwise mutual information. This allows us to study the co-occurrences of gods in semantic contexts. From a network perspective, Aššur is not a very central god in our corpus despite his importance in Assyrian royal theology, but he rather joins the existing networks of gods without altering them.

1. Introduction

In this article, we use computational methods to study the role of the god Aššur in the Mesopotamian pantheon. We map his position in the network of deities and the contexts in which he appears with other gods in the first millennium BCE. This article does not aim to offer a comprehensive study of Aššur's role in the pantheon; rather, it investigates how certain computational methods can contribute to the study of cuneiform texts and Mesopotamian culture. Aššur was chosen as the focal point of our study for two reasons. First, Aššur's prominence in the Neo-Assyrian texts reflects the ideology and identity-building processes in the Assyrian Empire. Other gods still remained important, however, and we aim to find out if Aššur truly obtained a central, prominent position in the pantheon from a statistical perspective. Second, Aššur has been thoroughly studied in previous scholarship, which allows us to anchor the results of our computational analysis in careful philological work.¹

We use social network analysis (SNA) to analyze a corpus of 1,532 texts from the Neo-Assyrian period (from the late tenth to the late seventh century BCE), obtained from the Open Richly Annotated Cuneiform Corpus (Oracc). We use two software tools, UCINET and Gephi, to analyze and visualize the net-

1. The authors gratefully acknowledge the financial support from the Academy of Finland in the writing of this article. Furthermore, we would like to thank Niek Veldhuis for helpful feedback and cooperation regarding Oracc and Johannes Bach and Tommi Jauhiainen for their comments on the draft versions of this article. Abbreviations follow the Chicago Assyrian Dictionary and CDLI Wiki (http://cdli.ox.ac.uk/wiki/doku.php?id=abbreviations_for_assyriology).

works. SNA is used on two levels. First, we produce networks of gods by looking at their co-occurrences in a text. The resulting network shows which gods were relevant enough for the ancient author to mention them in one and the same text. Second, we analyze co-occurrences in a ten-word window, meaning that two gods are connected if they are separated by no more than eight other words in the text. This allows us to map co-occurrences in more lexical semantic contexts and produce networks that show which gods are actually mentioned together in our corpus.

To further analyze the semantic contexts in which Aššur occurs, we use pointwise mutual information (PMI) as a measure when calculating tie strengths between the actors in our network. Our results show that PMI is well suited for revealing semantic relationships between gods, and it enhances the traditional SNA methods in certain respects.

Finally, in order to better understand the nature of Aššur's co-occurrences with other gods in the network, we study the contexts in which he appears together with some of the major Mesopotamian deities. We do this by using the corpus search tool Korp to display text windows in which Aššur co-occurs with certain gods. These contexts are then analyzed using traditional philological methods.

Our results show that, although the Assyrians actively sought to establish Aššur as the king of the gods in the Mesopotamian pantheon, he remained an outsider in the networks of gods. The gods with whom Aššur frequently interacts are the most important deities of the Mesopotamian pantheon in the first millennium. At the same time, Aššur does not, from a statistical point of view, have a central position in the networks, and he does not hold the networks together in any significant way. A noteworthy exception is a group of Arabian gods who were deported to Assyria and who are connected to the divine networks only via Aššur. This shows that gods from peripheral regions could remain marginal and unconnected, while gods from the core areas of the Assyrian Empire formed a tightly knit network.

The corpus of texts used in this article is available in Korp, an online service provided by the Language Bank of Finland.² It allows the user to make simple or more complicated queries in the text corpus, to create concordances, and to export the data in various formats. At the time of writing of this article, Korp contains 15,421 texts from Oracc, including all of the texts used in this article. All of the data used for and generated during our research is freely available online.³

The article is structured as follows. In section 2, we sum up previous research on Aššur and in section 3 we explain how our dataset was created and what texts it contains. Section 4 discusses the methods and tools used in this study, namely, social network analysis, graph visualization with Gephi, pointwise mutual information (PMI), and the text corpus search tool Korp. Section 5 is the main part of the article, illustrating the results we have gleaned from our material. Finally, section 6 summarizes our main findings and offers suggestions for further research.

2. Aššur's Biography

Aššur is an unusual deity in the ancient Mesopotamian milieu for several reasons that are related to his origins. The god, who is attested from the Old Assyrian period onwards, was initially a *numen loci*, intrinsically tied to a particular rocky outcrop in Ashur, modern Qal'at Sherqat in Iraq (Lambert 1983). This outcrop developed into a city named after Aššur and housed the god's temple, Ešarra, which remained the center of Aššur veneration until the end of the Assyrian Empire. Essentially, Aššur was one of the few "purely" Assyrian gods, so intimately tied to the religious and (for a long time) political center of Assyria

2. The "Oracc in Korp" corpus, provided by FIN-CLARIN: <http://urn.fi/urn:nbn:fi:lb-2018071122>; for metadata on Oracc in Korp; see Jauhiainen, Sahala, and Alstola 2017.

3. <https://doi.org/10.5281/zenodo.2620130>.

that worship of the god never spread outside of Assyria.⁴ As Assyria grew to eventually become a massive empire that spanned most of the Near East and into Egypt, Aššur's nature changed and expanded alongside it. The god's close association with Assyrian royal power meant that he was responsible for all aspects of kingship, from nomination and investiture, to commanding and protecting kings in their political and military affairs, to ensuring the king's long and successful reign (Maul 1999; 2017: 348–50; Parker 2011: 364–67). These characteristics were fleshed out in epithets and descriptions of the god's actions in state texts, giving the reader the consistent impression of Aššur as the ultimate king of the gods as well as the representation of Assyrian power.

But it is not only Aššur's unequivocal Assyrian nature that sets him apart from other Mesopotamian gods, who tended to be shared or identified with similar deities amongst several regional or local panthea, or between Sumerian and Akkadian panthea—Aššur's development from *numen loci* to full *deus persona* meant that quite significant changes had to occur in order to assimilate him with how other gods were conceptualized in the ancient Near East. That is, Aššur was one of few gods who did not originally have a divine family, genealogy, or divine circle of companions, nor did he seem to have an anthropomorphic divine image until quite late in the Assyrian Empire's development—Sennacherib (704–681 BCE) was the first king to mention creating an image of the god—and his original depictions as a god (rather than as a divinized natural feature) may have been therioanthropomorphic (part human and part animal) and perhaps related to bulls (Frahm 1997: 283; von der Osten-Sacken 2010: 746–47, 809–10; Lassen 2017). Unlike older, better-established gods, Aššur does not appear in lexical lists of divine names, which were created for scholarly purposes and did not necessarily reflect active cult but were a source for identifications between gods across language and geography, meaning that Aššur had to be inserted somehow into preexisting hierarchical relationships. As a result, these features were artificially created for Aššur over time: he is never given parents, but becomes himself the father of Ninurta, Šērū'a, and Zababa at various points.⁵

Aššur also takes on a consort, though it is not clear which goddess was thought to be his (primary) spouse at any given point, as Mullissu seems to have been the original goddess to be called his consort, with the later transition of Šērū'a from a daughter to a wife also a possibility. Furthermore, Ištar is a candidate, as she herself was associated closely with Mullissu, and there is a suggestion that Ištar of Nineveh and Ištar of Arbela were consorts of Aššur in their respective cities (Frahm 2014: 215–16; Meinhold 2014: 144–47; see also Allen 2015: esp. 177–88). Some of these growing pains are a result of cases in which Aššur was simply identified with other high-status deities and took on their attributes and divine relationships; for instance, Aššur is identified with Enlil, the high-ranking patron god of Babylonia's ancient religious center at Nippur, and thus takes on his status (“Enlil of the gods”) as well as his consort, Ninlil, later called Mullissu.⁶ Aššur is also linked with or replaces such gods as the primordial god Anšar or the god of Babylon, Marduk.⁷ Otherwise, Aššur is simply the head god of the Assyrian pantheon, and his links to other

4. Ashur was the religious and political center of Assyria until Ashurnasirpal II (883–859 BCE) moved the political capital to Calah (modern Nimrud, Iraq), though Aššur's temple did not move and Ashur remained the traditional religious center until the end of the empire. There is evidence for other cult locations, such as a local shrine in Kārum Kaneš in Anatolia during the Old Assyrian period, which would have serviced the Assyrian trading outpost located there (Larsen 2015: 265), and there have been suggestions that a temple or cult space to the god existed in Kār-Tukulti-Ninurta (under Tukulti-Ninurta I) and Nineveh (under Sennacherib), plus a temple in Uruk during the Neo-Babylonian period after the fall of the empire. See Menzel 1981, 1: 120–21; George 1993: 117 no. 687; Beaulieu 1997; Frame 1997: 55–64, 1999: 12–13; Holloway 2002: 65–66, 74; Meinhold 2014: 144.

5. E.g., SAA 12: 87 2' (Zababa); Lambert 1983: 82.

6. Von der Osten-Sacken 2010: 752–55, 761–62; Maul 2017; Mullissu's name is written ⁴NIN.LÍL.

7. Frahm 1997: 282–88; Frame 1999: 8–9; Pongratz-Leisten 2015: 416–26; Maul 2017: 352. The association with Marduk includes Assyrian kings who were seeking to unite Assyria with Babylonia by ideologically pairing Aššur and Marduk together as well as Sennacherib replacing Marduk with Aššur in rituals, literature, and other cult practices, resulting in Aššur taking on some of Marduk's characteristics and roles.

gods are framed within that characteristic, though some gods are geographically and theologically better linked to him because they have cult spaces in Ashur or in the Ešarra itself.

Overall, one gets the impression that Aššur's personality is fairly static in Assyrian state texts since he is consistently the high god of Assyria, for whom the kings accomplish their various deeds. Most of Aššur's roles in these texts feature him acting in independent capacities, especially in regards to Assyrian power and kingship, or sometimes paired (or in a triad) with other gods such as Marduk or Ištar, depending on the context and the king to whom the inscriptions belong. In extended deity sequences, Aššur always appears first in the list, whereas other gods are prone to more variation. Indeed, there are no cases in which a deity sequence in an Assyrian royal inscription does not begin with Aššur.⁸ Aššur's characteristics in these texts, as expressed by epithets, are consistently related to his sovereign status but remain relatively vague, emphasizing his own divine kingship and lordship over the universe, destiny, and judgment but largely lacking the more colorful epithets given to deities such as Ištar or Adad or other "great gods." Because of Aššur's ubiquitous role in Assyrian state texts and his primacy in the pantheon, all of the deities in the Assyro-Babylonian pantheon are somehow related to the Assyrian high god, however, whether directly or not, though these connections are difficult to untangle using traditional philological methods, in part because of the amount of data available. Nonetheless, a study of these connections and contexts can better illuminate how Aššur was perceived and how ancient sources relate him to other gods in the pantheon across time and different text genres. Since manual statistical analysis of these connections and contexts would be quite time-consuming, we can instead turn to computational methods, which can considerably speed up the process and make it more accurate.

3. Dataset and Preparation for Analysis

As already noted, our data come from the Open Richly Annotated Cuneiform Corpus (Oracc).⁹ Oracc is one of the largest electronic corpora of Sumerian and Akkadian texts, and it is regularly updated. Our data are a snapshot of the corpus from September 2017, extracted from the files provided by Oracc.¹⁰ The dataset we have consists of 17,619 texts and almost two million words. The text material used for this article contains all texts tagged in Oracc as being from, and only from, the Neo-Assyrian period. We did not include texts tagged as, for instance, "Neo-Assyrian, Neo-Babylonian period." We rely, thus, solely on the metadata provided for each text in Oracc. The texts tagged as being "lexical" in Oracc were omitted since they refer to a great number of rare divine names, many of which are only attested in the lexical lists. We also left out the Oracc subprojects designated with the words "Scores" and "Sources," which only duplicate texts that are available elsewhere in the parent projects.

In Oracc, a text is identified by a P-number if it is in the CDLI catalogue and by a Q-number if it is a composite text (Robson 2017). For example, instead of being identified by its museum number BM 61433, this literary text is known as P404913 in Oracc, and, instead of its publication number *RINAP* 4: 1, this royal inscription is known as Q003230 in Oracc. If the text has not yet been assigned a P- or a Q-number, it has a temporary X-number (Tinney 2017). Because Oracc operates with P-, Q-, and X-numbers, these designations are used in the datasets of this article instead of the more conventional museum or publication numbers.

8. On Aššur's roles and epithets in the Neo-Assyrian royal inscriptions, see Zaia 2017.

9. <http://oracc.museum.upenn.edu>.

10. <http://oracc.museum.upenn.edu/doc/opendata/index.html>.

We noticed that some texts are included in more than one Oracc project. For example, the text Q003364 is edited both in the Royal Inscriptions of Babylonia Online¹¹ and Royal Inscriptions of the Neo-Assyrian Period¹² projects. Because P- and Q-numbers are unique, they always refer to one and the same text across all of the Oracc projects. However, we realized that X-numbers are not unique, and the same X-number can refer to different texts in different projects. We decided to include all texts with X-numbers in our analyses, because there was no simple automated way of removing the possible duplicates.

In order to include each text with a P- or Q-number only once in our dataset, we kept the version that had the most words with dictionary-form annotation. If several texts had the same number of these lemmas—the dictionary or citation forms of the words—we chose the one for which we had the most translated lines.¹³ In the cases in which there were still several candidates, we chose the text that was alphabetically first—for example, *ribo/babylon6/Q003364* over *rinap/rinap4/Q003364*—with no regard to which of the publications is more recent. Because the online publications of such texts were very similar to each other and because we were mainly interested in the names of the gods, minor differences in text editions were insignificant. Duplicate texts with different identification numbers were not removed. We ended up with a list of 130 texts that were left out of our analysis.¹⁴

We used Oracc metadata to identify gods in our data: all of the words tagged as “divine name” were considered as gods for the purpose of this article. Since Akkadian is an inflecting language, we decided to use only the lemmas of the words. Annotation of a text is always an interpretation by a scholar and, as Oracc is composed of a number of subprojects, the way a word is annotated in different texts is not always consistent. Since we are investigating the interaction of Aššur with other deities, it was imperative that all the lemmas pointing to the same god are written in the same way. We hence standardized the divine names.

Lemmas referring to divine names were standardized manually with some help from automated data processing. The point of departure was a list of 1,901 dictionary forms of divine names in our full Oracc dataset of 17,619 texts. Of these names, 1,312 had no variant spellings and they were left as they were. The remaining 589 divine names were standardized to 256 lemmas.¹⁵ To help the process of combining the lemmas, we processed lists that grouped lemmas according to strings of letters that they shared.¹⁶ The full lists of lemmas were manually analyzed with the aid of the lists of grouped lemmas, using transliterations in the Oracc data to check if two or more lemmas actually referred to one and the same Akkadian or Sumerian word in the transliteration. When several lemmas in Oracc referred to a single Akkadian or Sumerian word, we chose the most common lemma in Oracc to represent the word and replaced all other lemmas with the most common lemma. However, different ancient names of a deity were kept distinct. For instance, the lemma Ištar was used to replace Issar and IŠTAR but not Inanna. Accordingly, the names Bēl and Marduk were not combined, although Bēl was a common byname of Marduk.

In order to limit the number of variables in our dataset and focus only on certain research questions in this article, we decided not to include geographical information in our analysis. Thus, we did not distinguish between homonymous gods who were associated with different geographical names, such as Ištar of Arbela and Ištar of Nineveh. As gods are, however, often associated with certain cities, it is a potentially fruitful future research avenue to statistically analyze divine names in relation to geographical names.

11. <http://oracc.museum.upenn.edu/ribo/babylon6/Q003364>.

12. <http://oracc.museum.upenn.edu/rinap/rinap4/Q003364>.

13. The translations of the Oracc texts we have derive from the files we received from Oracc in October 2016. We would like to thank Niek Veldhuis and Steve Tinney for these files.

14. A list of the texts omitted from the analysis as duplicates: the file */Texts/TextsOmitted* at <https://doi.org/10.5281/zenodo.2620130>.

15. A list of standardized divine names: the file *Standardized_divine_names*.

16. The script used to find similar lemmas: the file */Scripts/Merge_lemmas*.

As mentioned above, the text material for this article contains all texts tagged in Oracc as being from the Neo-Assyrian period. As a result, our data are uneven in terms of the chronological and genre distribution of the texts. Our texts span from the reign of Aššur-dān II (934–912 BCE) until the fall of the Assyrian Empire in the late seventh century BCE. The majority of the texts originate from the late eighth and seventh centuries, from the reigns of Sargon II, Sennacherib, Esarhaddon, and Ashurbanipal. Our dataset consists of 5,056 texts.¹⁷ Over 4,100 of those texts come from Nineveh and 3,563 of them have been tagged as being written purely in the Neo-Assyrian dialect. Of the other texts, 407 are stated to have been written in Akkadian, 832 in Neo-Babylonian, and 219 are bilingual texts in Sumerian and Akkadian.

The most common genres in our dataset are administrative letters (1,654 texts), legal transactions (758), royal inscriptions (679), and astrological reports (563). The distribution of our texts in different genres inevitably affects our results: for example, the moon god Šin and the sun god Šamaš are attested much more often than most other deities in astrological reports. The various genres and the number of texts from each genre are presented in table 1. The information on text genres is extracted from Oracc metadata.

Since gods who appear only once or twice in the dataset do not offer any useful information for our purposes but only add noise to the data, we decided to analyze only the gods who appear at least three times in our texts. Only 1,532 of the texts mention at least two deities who are attested at least thrice in the corpus. These 1,532 texts make up the corpus we use with computational methods.

Table 1. The number of texts of different genres in the corpus used for analysis. The information on genres, including the categories themselves, is extracted from Oracc metadata.

Genre	Number of all Neo-Assyrian texts	Number of texts with two or more gods
Administrative	10	0
Administrative letter	1,654	347
Administrative record	284	14
Astrological report	563	228
Astronomical	9	3
Hemerological	12	6
Historical	4	2
Hymn-prayer	20	16
Incantation-ritual	31	22
Legal	16	1
Legal transaction	758	40
Letter	1	1
Literary	87	47
Mathematical	1	0
Medical	9	0
Omen	65	22
Prayer/incantation	20	10
Priestly letter	203	135
Prophecy	9	6
Ritual	138	41
Royal inscription	679	247
Scholarly	5	2

17. Our dataset of Neo-Assyrian texts: the folder Texts.

Genre	Number of all Neo-Assyrian texts	Number of texts with two or more gods
Scholarly letter	386	305
School	23	6
Technical procedure	14	3
Treaties	14	12
Uncertain/unspecified	41	16
Total number of texts	5,056	1,532

4. Methods

4.1. Social Network Analysis

Social network analysis (SNA) is a powerful set of methods for analyzing the patterns of interaction between people in contemporary and past societies.¹⁸ It builds on mathematical graph theory and analyzes interaction as a network of connections between related actors. Social relationships can be visualized in beautiful graphs, but the strength of SNA lies in its mathematical procedures. They can be used to analyze complex social structures and reveal patterns and characteristics that may go unnoticed by a human observer alone. This is made possible by transforming the connections between actors in a given social network into matrices that a computer can analyze, giving results in a form that can be more easily used by a researcher.

SNA has primarily been used in social sciences, but it has recently attracted growing interest in historical studies because digital sources allow easy implementation of computational methods. Some of the possibilities of SNA have already been explored in Assyriology,¹⁹ but the full potential of this method has not yet been realized in the study of the large text corpora from the Neo-Assyrian, Neo-Babylonian, and other well-documented periods.

Social networks consist of *actors* and *ties* between them. When networks are analyzed as graphs, actors are called *nodes* and ties *edges*. As full networks can be laborious to analyze, they are often split into smaller units for a closer study. For example, an *ego network* consists of a central actor (*ego*) and actors directly connected to the ego (*alters*, Prell 2012: 7–12, 118–33).

Density measures how tight a network is, and it is calculated by dividing the actual number of ties in a network by the total number of ties if all the actors were connected to each other. Accordingly, value 0 means there are no ties between the actors and value 1 means that all the potential ties are present in the network (Prell 2012: 166–71).

Centrality refers to an actor's position in the network and her potential to reach other actors.²⁰ We used three different centrality scores. First, *degree centrality* measures the number of other actors to whom an actor is directly connected. Degree centrality can answer questions like “how many friends does a per-

18. On the history, theoretical background, and applications of social network analysis, see Scott 2000; Scott and Carrington 2011; Prell 2012.

19. Waerzeggers 2014a applies SNA to Assyriological data for the first time and explores the possibilities of the method. Waerzeggers 2014b uses SNA and the concept of structural holes to explain the exceptional personal success of a Babylonian man. Wagner et al. 2013 uses SNA to classify documents from the Murašû archive according to their chronological phase. Still 2016 applies SNA to study the marriage system among the priestly families in Borsippa.

20. On network centrality, see Prell 2012: 95–117 with further literature. Network centrality also connects to the concept of power and the idea that not all power networks are hierarchical, but that relational power between individuals should be taken into account in research (see further, Svärd 2015: 15–38, 147–59).

son have?” Second, *eigenvector centrality* is a more sophisticated version of degree centrality, as it takes into account the degree centrality of the alters. A person with well-connected friends would have a high eigenvector centrality. Third, *betweenness centrality* measures how often a certain actor is located on the shortest path between a pair of actors in the network. It shows who keeps different parts of the network together and controls the flow of information in the network.

Usually, not all actors in a network are connected to each other, and thus some actors are dependent on others if they wish to reach certain parts of the network. In other words, there are *structural holes* in the network, which certain actors may be able to bridge. Such actors often benefit from this role, as they have better access to resources in the network and because they can regulate communication between different parts of the network.²¹

We use two social network analysis software packages in this article. UCINET is primarily used to calculate the values for ego networks (Borgatti, Everett, and Freeman 2002), whereas Gephi is used for calculating scores for network centrality, density, and modularity, and for visualizing networks (Bastian, Heymann, and Jacomy 2009).

4.2. Visualizing Networks

To visualize our results, we used Gephi for building graphs and networks. After experimenting with the four standard graph-drawing algorithms available in Gephi as well as with some add-on algorithms, we chose to use the one called ForceAtlas2 (Jacomy et al. 2014) because its results seemed to be most interpretable considering our research questions. ForceAtlas2 is a so-called force-directed graph drawing algorithm, where the nodes (gods in our case) repulse each other; that is, they try to get as far away from one another as possible, whereas the edges (the ties between the gods) make nodes attract each other pulling connected nodes closer (Jacomy et al. 2014). The distance between nodes is the smallest number of edges between them. The smaller the distance is, the stronger the attraction.

Ties in a network are not normally evenly distributed, but a full network often contains several communities. *Community* refers to a subgroup within the network, the actors of which have many ties between one other but fewer ties to the rest of the network. Community structure can be detected using various computational methods that aim to find the best possible way to divide a network into communities. With force-directed graphs, one can visualize communities as grouped nodes within the network (Noack 2009). Gephi provides an option to use a mode called LinLog with the ForceAtlas2 algorithm (Jacomy et al. 2014). The LinLog calculates the attraction between nodes by taking the logarithm from the distance between nodes, which makes the clusters in a graph tighter (Jacomy et al. 2014).

A measure called *modularity* is used to evaluate the quality of the division of a network into communities (Newman and Girvan 2004). It is calculated by taking the fraction of all edges that are within communities and subtracting from it the fraction of all edges that would be within the same communities if the distribution of all the edges were random. The modularity value 0 means that the division of a network into communities does not differ from a random division, and the maximum value 1 indicates that a method has detected a very strong community structure. In general, it is considered that values above 0.3 indicate that an algorithm has detected a strong community structure in the network (Newman and Girvan 2004).

We used Gephi to split the network into communities and to analyze the modularity of communities. An algorithm called the Louvain Method (Blondel et al. 2008) is used for splitting networks in Gephi. This

21. On the concept “structural hole,” see Burt 1992. See Prell 2012: 122–25 for practical examples.

community extraction method first assigns each node to its own community. Then it iteratively evaluates the neighbors of each node and considers whether placing the node in the community of a neighbor would be beneficial for the modularity of the network. The best results are gained by using a random order when choosing the next node to consider. This is repeated until no improvements are gained.

4.3. Pointwise Mutual Information

Pointwise mutual information (PMI) is a measure of association that was adapted to language technology by Church and Hanks (1990) as a statistical method for finding collocations and associations between words. As a measure, PMI compares the expected chance of two words co-occurring independently with the actual probability of their co-occurrence calculated from the corpus. If the actual probability of the co-occurrence is significantly greater than the expected chance, the words may share some kind of mutual information. The nature of this mutual information is often dependent on the window size that defines the maximum allowed distance between the words. Small window sizes are appropriate for recognizing fixed expressions or compound words (e.g., stock market, bus station), whereas larger windows may provide insight into words related by semantic concepts (markets ~ finance; subway ~ transportation, Church and Hanks 1990: 23).

Procedurally explained, the traditional method proposed by Church and Hanks finds all possible pairs of words (w_1, w_2) that fit inside the given forward-looking window. The PMI score is a logarithmic ratio of the joint probability $p(w_1, w_2)$ to the expected probability that the words are observed together independently $p(w_1)p(w_2)$, as in Equation 1.

$$\text{PMI}(w_1, w_2) = \log_2 \frac{p(w_1, w_2)}{p(w_1)p(w_2)} \quad (1)$$

As PMI has a tendency to give exceedingly high association scores for low-frequency words (Bouma 2009; Croft, Metzler, and Strohman 2010), we chose to use the PPMI² measure²² as defined in Equation 2 (Role and Nadif 2011). This measure is based on the PMI^k family proposed by Daille (1994), which greatly reduces the unwanted low-frequency bias by raising the numerator to the power of k . We chose to use a commonly preferred k -factor of 2.

$$\text{PPMI}^2(w_1, w_2) = 2^{\log_2} \frac{p(w_1, w_2)^2}{p(w_1)p(w_2)} = \frac{p(w_1, w_2)^2}{p(w_1)p(w_2)} \quad (2)$$

The minimum and maximum scores for PPMI² are 0 and 1, which indicate that the words are either never or always attested together within the given window. When referring to PMI in this article, we mean PPMI².

4.4. Korp

Korp is an online service provided by the Language Bank of Finland.²³ Korp allows users to make queries in text corpora, whether in one corpus or in several corpora at the time. The results of the query are presented as *concordances*, that is, listing line by line every instance of the word matching the query with

22. Role and Nadif (2011) refer to this formula ambiguously as PPMI (Positive PMI). We chose to call it PPMI² instead, because PPMI often refers to another formula, $\max(\text{PMI}, 0)$ (Jurafsky and Martin 2017: §15.2).

23. <https://korp.csc.fi/#?lang=en>; <https://www.kielipankki.fi/support/korp/>.

all its neighbors. Furthermore, the results are, by default, presented in the so-called *KWIC* (keyword in context) view, meaning that each result is displayed with the matching words highlighted in the middle of the display. Korp is a useful tool for studying the contexts in which words appear.

Korp software was originally developed by the Language Bank of Sweden (Språkbanken) for their own online service (Borin, Forsberg, and Roxendal 2012), but it is the Finnish Korp service that contains the “Oracc in Korp” corpus.²⁴ At the time of writing, the Finnish Korp contains the texts from Oracc as of September 2017, including all of the texts used in this article. We are, however, working with Oracc to implement regular updates of Oracc in Korp. The entry for each word in Oracc in Korp contains the information about transcription, dictionary form, translation, part-of-speech tag, etc., provided the information is available in Oracc.

There are three levels of search in Korp. With simple search, one can look for a word, a sequence of words, or a string of characters. Within the Oracc in Korp corpus, the simple search looks for matches in transliterated texts. With extended search, it is possible to query for several consecutive words at once as well as metadata or annotations, for example, a translation or a part-of-speech tag of a word. In the advanced search, one can use the CQP (corpus query processor)²⁵ query language and make more complicated queries in the material. In Oracc in Korp, it is possible to make searches using Unicode transliterations, transcriptions, and dictionary forms. It is also possible to use ASCII characters in queries. All the results of a query also contain a link to the relevant text in Oracc.

4.5. Combining the Methods

In this article, we use social network analysis (SNA) to analyze the co-occurrences of divine names in texts from the Neo-Assyrian period. SNA is used for analyzing three different datasets. First, we produce networks of gods by looking at their co-occurrences in a text. It does not matter how many times two gods are mentioned in a single text, but it is significant in how many texts they co-occur. The strength of the tie between two gods equals the number of texts in which the gods are attested together. Because the information given by words other than divine names is irrelevant, we created a simple list linking a text to a god attested in it. Texts that only contained one god were discarded. This node list was imported into the social network analysis software UCINET and processed into a two-mode network linking gods to the texts in which they are attested. The two-mode network was further transformed into a one-mode undirected network linking gods who appear in the same texts. Afterwards, the network was exported to Gephi for further analysis and visualization. UCINET was used to calculate the values for ego networks, whereas Gephi was used for calculating scores for network centrality, density, and modularity, and for visualizing networks. The resulting network shows which gods were relevant enough for the author of the text mention them in one and the same text. We use this dataset to analyze centrality and ego network structure in sections 5.1 and 5.2.

Second, we use SNA to analyze the co-occurrences of gods in a ten-word window, meaning that two gods are connected if they are separated by no more than eight other words in a text. Tie strength equals the sum of windows in which two gods co-occur. We implemented a program to produce a list of co-occurring gods from the texts in our corpus.²⁶ The program created an edge list expressing connected nodes and the strength of the tie between them. This list was imported into UCINET and Gephi and transformed into an undirected network to analyze the data according to the procedures explained in the previous

24. <http://urn.fi/urn:nbn:fi:lb-2018071122>

25. Evert and the CWB Development Team 2016.

26. Source code: the file /Scripts/Pmizer-assur at <https://doi.org/10.5281/zenodo.2620130>.

paragraph. Using a small window allows us to map co-occurrences in more semantic contexts. This dataset is used for analyzing centrality and ego network structure in sections 5.1 and 5.2.

Third, after analyzing the results of SNA with a ten-word window, we realized that using the raw number of co-occurrences as tie strength did not capture the semantic dependency of two gods very well. This was indicated by low modularity scores and by network visualizations that did not provide very useful information for our research on network community structure. Consequently, we decided to use pointwise mutual information (PMI) to calculate the tie strengths in ten-word windows. In our previous research (Svård et al. 2018) and in the course of writing this article, we have observed that using relatively small windows with PMI is well suited for analyzing semantic contexts in our corpus. We used the same corpus that had been created for analysis of simple co-occurrences in ten-word windows and we only calculated PMI scores for gods who were attested at least three times in the corpus. The PMI scores were calculated with a window size of 10 for each pair of co-occurring divine names.²⁷ We only took into account the pairs that were attested at least twice to leave out unique pairs. Due to the forward-looking window used for calculating the PMI scores, the scores are not symmetrical in nature. In other words, the PMI score for Aššur and Marduk takes into account the cases when the word Marduk is attested in a text after Aššur with no more than eight words between them. The scores for both directions (i.e., Aššur and Marduk, and Marduk and Aššur) were summed up and this value was used as the tie strength. This allowed us to create an edge list indicating the connected nodes and the tie strength between them. The edge list was imported into UCINET and Gephi and transformed into an undirected network to analyze the data according to the procedures explained above. We use this dataset to analyze community structure in ego networks in section 5.4.

For all our visualizations we used the ForceAtlas2 algorithm with LinLog mode. The size of a node indicates its degree, that is, the number of nodes to which it is directly connected. The colors of the nodes show communities assigned by the modularity algorithm in Gephi. The method of visualization of our networks that we used is only one of the possible ways of doing the visualization and is in no way absolute. The visualization was mainly used for directing our research, that is, choosing which gods to analyze, and for raising further research questions. All of the graphs and the exact parameters used for creating them are available in the online repository.²⁸ We used the corpus search tool Korp to explore and analyze the contexts in which certain gods appear together.

5. The Relationship of Aššur to Other Gods

In the following, we discuss the results of our statistical analysis and what light this sheds on Aššur's connections with other gods as well as the nature of these connections. We use social network analysis to study how Aššur relates to other gods on two different levels: texts and semantics.²⁹ On the level of texts, two gods are related if they appear in the same text. On the level of semantic contexts, two gods are related if they co-occur within a window of ten words. Moreover, we use the number of co-occurrences and PMI scores to indicate tie strength between two co-occurring gods. We also analyze the contexts in which Aššur is attested together with some of the major Mesopotamian gods using the corpus search tool Korp.

We begin with studying Aššur's position in different networks. We analyze centrality in full networks and then focus on the structure of Aššur's ego networks. Then, we analyze the community structure of the divine networks, aiming to find patterns that illuminate Aššur's status in relation to other gods. Finally,

27. The source code of our PMI script: the file /Scripts/Pmizer-assur.

28. The folder Visualizations at <https://doi.org/10.5281/zenodo.2620130>.

29. Network datasets: the folder Network-data at <https://doi.org/10.5281/zenodo.2620130>. Visualizations: the folder Visualizations.

we provide two concrete examples of how SNA can be used to identify interesting patterns in the network structure and how network data and corpus search tools can be used to address the questions arising from these observations.

5.1. Centrality in Full Networks

We calculated three different centrality scores for each god in the networks created with *full texts* and a *small window* of ten words: degree centrality, eigenvector centrality, and betweenness centrality. Because these three measures do not take tie strength into account, we did not need to calculate co-occurrences or PMI values for this purpose. Different centrality measures offer different perspectives on Aššur's position in the divine networks. Degree centrality measures the number of gods with whom he is directly connected, showing that they appear in the same text or semantic context. Eigenvector centrality takes into account the connectedness of Aššur's immediate connections and thus shows whether he is well-connected to the central gods in the network. Betweenness centrality analyzes centrality from a different perspective: a high score would indicate that Aššur bridges structural holes in the network and connects different groups of gods to each other.

Surprisingly, Aššur is only the second most-often attested god in our corpus after Šin.³⁰ This stems from the fact that Šin is frequently mentioned in astrological reports, a common text genre in the corpus. More surprising, however, is that Aššur's network centrality scores are relatively low in comparison with the number of his attestations.³¹ In SNA with full texts, he is placed at position 13–16 among the gods, and deities such as Ninurta and Nergal always have higher scores than Aššur. SNA with a small window provides us with similar results for degree and eigenvector centrality (the eleventh highest score in both), but here Aššur's betweenness centrality score is the fourth highest in the corpus. However, this is explained by a number of Esarhaddon's inscriptions referring to the statues of the gods of the Arabs, which were captured during Sennacherib's campaigns.³² As three of these six gods appear within a small window with Aššur, he is the only link between them and the rest of the network. The other three gods of the Arabs attested in these passages do not fit into a ten-word window with Aššur, and they are connected to the network only via the gods linked to Aššur. This greatly affects Aššur's betweenness centrality score. When the six gods of the Arabs are removed from the network, Aššur only has the twenty-second highest betweenness centrality score when using a small window.

Centrality scores indicate that Aššur remained an outsider in the traditional network of long-standing Mesopotamian gods, even as late as the Neo-Assyrian period. Although the texts in our corpus continuously refer to him, Aššur is only related directly to a limited number of gods and he does not connect different parts of this traditional network to one other. Since this network existed before Aššur rose to a prominent position, he essentially enters the existing network without having much impact on its structures.

At the same time, Aššur connects the gods of the Arabs to the network. This shows that these gods remained outside of the Mesopotamian pantheon and came into contact with it only after the Assyrian military campaigns against the Arabs. Aššur's connective role highlights his position both in the network and in the Assyrian royal ideology: he was not a very central god from a network perspective, but he was the

30. A list of gods and their frequencies in our corpus of 1,532 texts: the file *Frequencies_gods*.

31. Network centrality scores: the file *Centralities*.

32. *RINAP* 4: 1 iv 1–16; 6 iii 1'–10'; 97 7–14a. Esarhaddon claims that he wrote the "might of Aššur" on these statues and returned them to Hazael, the king of the Arabs. See Retsö 2004: 153–61; Zaia 2015: 34–36.

most important god to the Assyrian Empire and the king. Accordingly, Aššur is the single Mesopotamian deity attested together with the foreign gods of the Arabs, spoils of war, who were deported to Assyria.

5.2. Ego Network Structure

Full networks shed light on Aššur's position in general, but they are too complicated for a detailed analysis of his immediate connections. In order to better understand Aššur's role in the divine networks, we focused on his ego network, that is, the network consisting only of the gods to whom he is directly connected. Aššur's ego network created with full texts comprises ninety-seven other deities, and, with a small window, thirty-two other deities.³³ In order to study the nature of these ego networks, we analyzed their density and Aššur's betweenness centrality in them. These values help us to trace possible structural holes in the ego network: low network density and the ego's high betweenness centrality may indicate that there are structural holes in the network that the ego (here, Aššur) fills.³⁴ These measures are calculated for binary data, not taking tie strength into account.³⁵

With full texts, the density of Aššur's ego network, with ninety-seven other deities, is 43.0 and Aššur's normalized betweenness centrality in the ego network is 4.7. As these values have little meaning on their own, we compared them with those of other ego networks of comparable size. Ten other gods or groups of divine beings in our sample—the Anunnaki, Asalluhi, Bēlet-ilī, the Igigi, Ištar, Nergal, Ninurta, Nusku, the Sebetu, and Uraš—have an ego network of 80–122 alters. The average density of their ego networks is 44.4 and the average of their normalized betweenness centrality in their ego networks is 5.4. These numbers are close to those of Aššur, and so they do not indicate that there are structural holes in Aššur's ego network or that his ego network is different from other ego networks of comparable size in our corpus.

With a small window, the density of Aššur's ego network, with thirty-two other deities, is 47.4 and Aššur's normalized betweenness centrality in his ego network is 27.1. Eight other gods in our sample—Mullissu, Nabû, Nergal, Ninurta, Nusku, Šamaš, Šin, and Uraš—have an ego network of 26–39 alters. The average density of their ego networks is 47.3 and the average of their normalized betweenness centrality in their ego networks is 23.9. In the light of betweenness centrality scores, Aššur's position in his ego network seems to be somewhat stronger than that of the egos in the control group. However, we notice again that the figures are greatly affected by the gods of the Arabs in Aššur's ego network. When these gods are removed, the density of Aššur's ego network is 57.1 and Aššur's normalized betweenness centrality in his ego network is only 11.7. Accordingly, Aššur is instrumental in connecting a group of foreign deities to the rest of the gods, but otherwise he does not bridge structural holes in the network.

Our analysis of Aššur's ego network structure essentially confirms the conclusions made in the previous section on network centrality. At a general level, Aššur's ego network is densely connected and it is not structurally dependent on him. Because of his prominent role in the Assyrian royal theology, Aššur primarily interacts with the most important Mesopotamian deities of the first millennium BCE, and numerous minor deities remain outside his ego network. There are very few gods who are primarily connected to the network via Aššur. As he originally had no divine genealogy, he did not introduce a host of family members or attendants to the Mesopotamian divine networks. The gods of the Arabs, however, are connected to the networks only via Aššur. We may conclude that Aššur entered the existing network of Mesopotamian deities without taking a central role in it.

33. Visualizations: the files *Text_Assur-ego* and *W10_co-occurrences_Assur-ego* in the folder *Visualizations*.

34. On analyzing structural holes in ego networks, see Prell 2012: 124–25 with further literature.

35. Basic measures of ego networks in our datasets: the folder *Egonets*.

5.3. Modularity and Visualization

In addition to Aššur's position in the networks in general, we wanted to study the network structure in more detail. First, we analyzed the community structure of our networks. We were interested in seeing if the networks split into communities, revealing something about Aššur's relationship to the other gods of the Mesopotamian pantheon. Second, we used the ForceAtlas2 algorithm to visualize network structures in Gephi. The resulting visualizations were used to observe Aššur's position in the networks vis-à-vis other gods on a level that is not feasible for a human observer using only textual or statistical data.

In order to visualize the networks effectively and to analyze their community structure, we needed data that presents tie strength between two actors. It turned out that using the simple number of co-occurrences as tie strength did not yield useful results as the modularity scores were relatively low and thus the network could not be reliably split into communities. The scores for the full network with full texts were below 0.18 and with a small window below 0.24. For Aššur's ego network, the scores were below 0.14 with full texts and below 0.18 with a small window. As the scores were rather low, the resulting divisions of the networks were likely somewhat arbitrary.

In search for a better way to measure tie strength, we calculated PMI values for each pair of gods appearing in the small window of ten words. Using PMI values to indicate tie strength yielded useful results. The modularity score was very strong (0.77–0.78) in the full network and strong (0.57–0.59) in Aššur's ego network. As stated in section 4, modularity scores above 0.3 indicate that the algorithm has detected a significant community structure. Accordingly, using PMI values as tie strength highlights the community structure of our networks much better than the number of co-occurrences does. Moreover, using PMI values as tie strength also produced more informative and understandable graphs with ForceAtlas2. In what follows, we discuss the modularity and structure of networks created with PMI values as tie strength.

5.4. Community Structure in Ego Networks

Aššur's ego network created with a small window and PMI values has a high modularity score (0.57–0.59), and it splits into four communities of unequal sizes (fig. 1). The composition of the communities is not stable, and the modularity algorithm, when run multiple times, divides the network in three different ways.³⁶ However, as discussed below, the different divisions are not random or arbitrary. On the contrary, they neatly reflect the structure of the network.

Two communities stay the same in all three divisions of the network. Aššur himself belongs to the first community, consisting of twelve gods including Aššur (Community A). This group primarily consists of prominent Mesopotamian deities: Šin, Šamaš, Ištar, Adad, Nergal, Marduk, Bēl, Nabû, Zarpanītu, and Tašmētu. Additionally, Tiamat belongs to this community, because she is only connected to Aššur, Marduk, and Adad in the ego network.

The second stable community in the network consists of the aforementioned three gods of the Arabs, Ruldāwu, Abirillu, and Atar-qurumâ (Community B). They form a tight group that is distant from the rest of the network, being only connected to each other and Aššur.

The remaining eighteen gods in Aššur's ego network do not constitute a single community since the modularity algorithm divides them into two communities in three different ways. However, the eighteen gods belong to four stable subunits. A *subunit* refers here to a group of actors that is never split into two communities, but the whole subunit can move from one community to another. It appears that the bor-

36. Visualizations: the files PMI_Assur-ego_mod1/2/3 in the folder Visualizations.

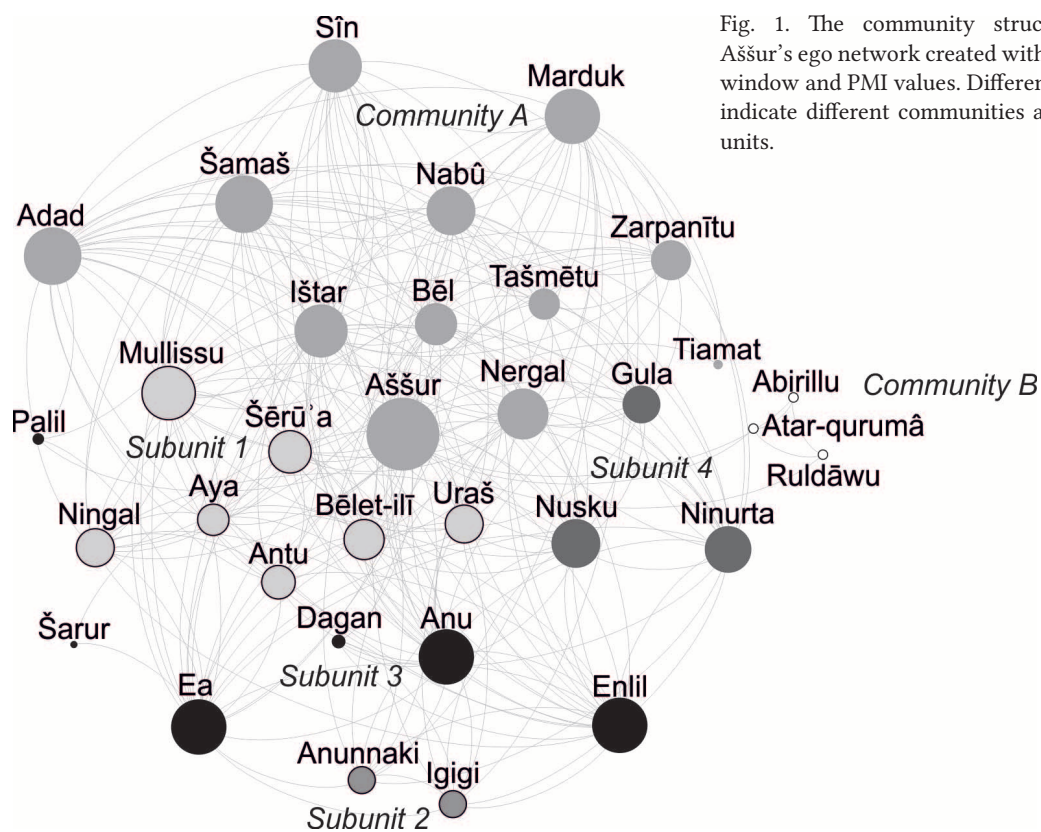


Fig. 1. The community structure of Aššur's ego network created with a small window and PMI values. Different colors indicate different communities and subunits.

ders between the subunits are rather fluid in general, because all of them can join other subunits to form communities. However, Subunits 1 and 2 are mutually exclusive and never fit into the same community.

One stable subunit in all three divisions contains Uraš and several female deities (Subunit 1): Anu's spouse Antu, Šamaš's wife Aya, Bēlet-ilī (who is often paired with Ea), Sîn's spouse Ningal as well as Mullissu and Šērū'a. The last two are female deities related to Aššur, but there is no scholarly consensus on which one of them was regarded as his wife.³⁷ Uraš's appearance in this group results from Esarhaddon's Succession Treaty, including its Tell Tayinat version, in which Uraš is attested a few times together with Bēlet-ilī, Mullissu, and Šērū'a.

Another stable subunit consists of the Igigi and Anunnaki (Subunit 2). They are somewhat different from the other gods in the network because the Igigi and Anunnaki do not refer to specific gods but to the gods of the heaven and the netherworld, respectively. The strong tie between the Igigi and Anunnaki results from the fact that they are frequently mentioned together in stock phrases in our corpus.

Anu, Ea, Enlil, Dagan, Paḷil, and Šarur belong to the third subunit (Subunit 3). Three of them, Anu, Enlil, and Ea, were ancient great gods in the Mesopotamian pantheon. In the first millennium they often appear, like the Igigi and Anunnaki, in stock phrases, representing the universe: Anu was the sky god, Enlil the king of the earth, and Ea the god of the subterranean waters. The three other gods were less important and relatively rarely attested in our corpus. Their position in this subunit results from their links to one

37. On this problem, see Meinhold 2014.

of the three great gods. Dagan is often attested with Anu, and Šarur appears in the *Epic of Anzû* in several contexts with Ea. Although Palil is connected to Aššur, Šamaš, and Adad in another community, his closer connections to Ea draw him into this subunit.

The fourth stable subunit consists of Ninurta, Nusku, and Gula (Subunit 4). Their close relationship reflects ancient theological beliefs: Gula was considered to be Ninurta's wife, and Nusku was perceived as the attendant of Enlil, the father of Ninurta. These relationships are also reflected in our texts in which Ninurta and Gula are often attested as a pair and Enlil, Ninurta, and Nusku as a triad, especially in blessing formulas in letters.

The analysis of the community structure of Aššur's ego network yields a number of observations, which can be summarized as follows. First, the community to which Aššur belongs is stable and does not change when the modularity algorithm is run multiple times. This community is distinctively different from the rest of the network. This divide is reasonable from an Assyriological perspective: many of the gods in Aššur's community are related to kingship and are frequently attested in royal inscriptions, whereas the gods in the rest of the network are a more diverse group, many of them frequently attested in fossilized expressions and stock phrases.

Second, it is somewhat surprising that, while Nergal belongs to Aššur's community, Ninurta is part of another subunit. These two gods co-occur frequently in royal inscriptions and were both associated with war. Moreover, according to some scholars, Ninurta had a very central role in the Assyrian royal theology, which would suggest a close connection to Aššur. Ninurta's absence in Aššur's community suggests that the picture is more complicated, at least from a network perspective, and that Ninurta and Nergal have a different set of close ties in the network that results in them being in separate subunits.

Third, the position of several female deities in the network is unexpected. Zarpanītu and Tašmētu are indeed close to their husbands Marduk and Nabû in Community A, but Antu, Aya, Bēlet-ilī, Ningal, Mullissu, and Šērū'a belong to Subunit 1 while their husbands are in other subunits or communities. It is noteworthy that these female deities are grouped together so closely, because they are not normally listed as a group of their own in Assyrian texts. However, sequences like "Aššur, Mullissu, Šērū'a, Sīn, Ningal, Šamaš, Aya, Anu, Antu" are relatively common in our texts. This results in strong PMI scores between the goddesses, who are rather rarely attested in the texts in general.

The previous examples show how the division of networks into communities and their visualization can be used to obtain a simplified view of a complex SNA dataset. Because we were able to meaningfully explain the community structure with philological methods, we deem it to be of good quality. As demonstrated above, a significant utility of SNA is its ability to highlight peculiarities in the network structure. This can help scholars to perceive interesting patterns in the data and analyze them in more detail. In the following section, we focus on two features of Aššur's ego network and study them using digital and traditional philological methods.

5.5. Understanding the Network Divisions

Aššur's ego network created with PMI values reflects the semantic contexts in which gods co-occur. Accordingly, the closeness of two gods and the statistical probability of their co-occurrence are taken into account. This explains, for instance, why the gods of the Arabs constitute a very tight group in the graph. In this section, we focus on two features of Aššur's ego network that were highlighted in the previous section: the case of Nergal and Ninurta and the question of Aššur's wives. We aim to show how one can use the results from SNA and the corpus search tool Korp to understand the network divisions and utilize them in Assyriological research.

In order to conveniently study the contexts in which different gods occur, we used Korp to list all attestations of certain gods with Aššur within the window size of ten words. We took into account the following cases: “Aššur + 0–8 words + god X” and “god X + 0–8 words + Aššur.” Korp is well-suited for tasks like this, because the search engine allows for more complicated queries than the search tool in Oracc.

According to previous Assyriological research,³⁸ Nergal and Ninurta were both associated with war and military campaigns in the Neo-Assyrian period as they inflicted death on the enemies of the Assyrian king. However, their roles are perceived to be somewhat different, as Nergal is viewed as the god of plague and violent death, whereas Ninurta is characterized as a warrior god. Moreover, Annus (2002: 95–101) and Pongratz-Leisten (2015: 219–70) have suggested that Ninurta had a special role in the Assyrian royal theology; namely, that the Assyrian king, although he was perceived as a human, was equated with Ninurta. The two gods are almost equally often attested in our corpus, Nergal 192 and Ninurta 207 times. As mentioned above, the modularity algorithm used in this article assigns them to different groups, Nergal to Community A with Aššur and ten other gods and Ninurta to Subunit 4 with Gula and Nusku.

Nergal appears fifty-two times together with Aššur in the small window.³⁹ He is often attested in lists of several deities to whom the Assyrian king prayed or whose images he fashioned. A number of legal transactions refer to Nergal and Aššur in a peculiar context: the one who breaks a contract or contests a legal transaction “shall tie two white horses at the feet of Aššur and bring four *ḫarbakkannu* horses to the feet of Nergal.”⁴⁰ These texts appear to be very influential in linking Nergal closely to Aššur in the network.

Surprisingly, Ninurta is only attested twenty-nine times with Aššur. They appear together in contexts in which Aššur and Ninurta support the Assyrian king in his royal duties, but a great deal of these texts are dated to the reigns of Adad-nirari II and Ashurnasirpal II, that is, in the earlier phase of the Neo-Assyrian Empire. The occurrences in the inscriptions of Ashurnasirpal II relate to the fact that the king transferred the capital of the Empire to Calah, an important cult center of Ninurta. As the bulk of our texts date to the Sargonid kings and as their inscriptions rarely mention Aššur and Ninurta together, the number of their co-occurrences is small in total. At the same time, Ninurta is strongly linked with his wife Gula and with Nusku, the attendant of Ninurta’s father Enlil. Consequently, Ninurta is drawn into a separate subunit away from Aššur and the connection between the two is quite weak.

This seems to be in contrast with Annus and Pongratz-Leisten’s claims for “Ninurta theology,” that is, that Ninurta had a central position in the Assyrian royal pantheon and royal ideology. Pongratz-Leisten (2015: 225, 261) writes, for instance, that Ninurta “emerged as *the* model of kingship” and “Ninurta and the king performed similar functions in their defense of the political and cosmic order.” The identification is based largely on Ninurta’s characterization as a heroic warrior and mythological works such as *Lugal-e* and *Anzû*, which Annus and Pongratz-Leisten view as having a direct impact on shaping royal ideology, as depicted in state texts such as the royal inscriptions (Pongratz-Leisten 2015: 258). While Ninurta’s importance in the royal ideology is visible in the individual texts, the quantitative perspective of the available Assyrian corpora given by network analysis does not give the same image of the god. Ninurta theology is not visible or made explicit in the network of divine relationships, which features other gods as more strongly connected to the Assyrian high god and, thus, part of the innermost circle of major deities.

Unlike the majority of Mesopotamian deities, Aššur did not have a well-defined divine genealogy. Consequently, Assyrian sources are ambivalent about Aššur’s wife as well, and he is associated closely with

38. On Nergal, see von Weier 1971; Lambert 1973; Wiggermann 1999; on Ninurta, see Annus 2002.

39. The number is the sum of co-occurrences in the following cases: “Aššur + 0–8 words + Nergal” and “Nergal + 0–8 words + Aššur.”

40. See, for example, SAA 6: 87 rev. 3–4.

two goddesses, Šērū'a and Mullissu, in the Neo-Assyrian period. Their exact relationship with Aššur is disputed, and it remains unclear if one of them was his primary or sole wife (Mullissu) and the other was his secondary consort or daughter (Šērū'a; Frahm 2014: 215–16; Meinhold 2014: 145–47; see also Allen 2015: esp. 177–88). During the late Neo-Assyrian period, Mullissu is increasingly identified as Aššur's consort as Aššur takes on to some extent the roles and attributes of Enlil; Mullissu (Ninlil) consequently shifts from being attested with Enlil to being mentioned alongside Aššur, especially in contexts that pair the gods in parental or consort roles. To complicate matters, Mullissu also becomes more strongly identified with Ištar of Nineveh and Meinhold (2014: 144) has argued that Ištar was Aššur's primary consort in Nineveh and Arbela, both of which had local forms of Ištar as their patron goddess. Šērū'a and Mullissu belong to Subunit 1 in Aššur's ego network together with Antu, Aya, and Ningal, the wives of Anu, Šamaš, and Šin, respectively. Their position is thus different from that of Zarpanītu and Tašmētu, who are located close to their husbands Marduk and Nabū in Community A.

In our corpus, Mullissu appears much more often together with Aššur within the small window (66 times) than Šērū'a does (10 times). Mullissu shares many contexts with Aššur, and they appear together in short greeting formulas in letters, in cultic contexts, and in royal inscriptions giving commands to the Assyrian king. Although attested much more rarely, Šērū'a shares several different contexts with Aššur: they appear in treaties, they receive offerings, and Sennacherib fashioned their statues. The raw number of co-occurrences and the stronger PMI value in our data indicates that Mullissu was closer to Aššur than Šērū'a. Moreover, Mullissu and Šērū'a are several times attested together, and they occasionally appear together with Aššur in contexts such as “May Aššur, Mullissu, Šērū'a ... curse you.” This results in a strong tie between Mullissu and Šērū'a. This is interesting but not conclusive because it could be because they were both consorts, thus in a triad of the form male god-primary consort-secondary consort, or because they represent a divine familial triad in the form father-mother-child. Nonetheless, the closer connection between Aššur and Mullissu may support the conclusion that Mullissu was the god's original consort. Moreover, while the strong relationship between Mullissu and Šērū'a leaves open the question of whether the latter transitioned from Aššur's daughter to his second-ranking wife at a late stage, this indicates Šērū'a's importance in the divine family despite her relative paucity of attestations in the corpus (Meinhold 2014: 149).

The case studies presented above illustrate how SNA can highlight unexpected patterns in the relationships between actors, in this case Mesopotamian gods. We chose to study two peculiar cases in closer detail and used the corpus search tool Korp to make the relevant texts easily available for traditional philological work. By combining the network data, PMI values, and information collected from Korp, we were able to contribute to the discussion of two problems related to the Assyrian religion. First, we noticed that Ninurta's link to Aššur is surprisingly weak, something that we may not have observed using traditional philological methods, as he is still prominent in individual mythological texts. Instead, we found that other gods were more strongly associated with Aššur, to the point that Ninurta belongs to a different subunit than Aššur. This indicates that “Ninurta theology” is largely invisible from an aggregate perspective and may suggest that, at least by the Neo-Assyrian period, its nature is such that it is explicated only in certain genres or individual texts. Even more surprising is that Nergal, with whom Ninurta is routinely mentioned in (for instance) deity sequences in the royal inscriptions, belongs to the subunit with Aššur. Using a network approach can thus lead us to investigate why this distinction, which would be difficult to detect otherwise, would be the case. Second, we studied the problem of Aššur's genealogy, specifically the question about his wife's identity. The two main candidates, Mullissu and Šērū'a, are closely connected to each other and to Aššur, but Mullissu is much more frequently attested together with her alleged husband. Our findings seem to support the theory that Šērū'a was regarded as secondary to Mullissu in terms of the goddesses' relationships with Aššur.

6. Conclusions

The aim of this article was to use computational methods to analyze the role of the god Aššur in the Mesopotamian pantheon in the first millennium BCE. The methods employed in this study benefit from large datasets and so we selected as our corpus the 1,532 digitized texts from the Neo-Assyrian period, obtained from the Open Richly Annotated Cuneiform Corpus (Oracc). We used social network analysis (SNA) to map Aššur's position in the networks of gods on two levels by analyzing his co-occurrences with other gods in full texts and in windows of ten words. The first option shows which gods were most often mentioned in the texts and with whom, and the second option sheds light on the semantic contexts in which certain gods appear. As the simple number of co-occurrences did not capture semantic relationships well enough, we used pointwise mutual information (PMI) to calculate tie strengths between the gods. Finally, we used the online tool Korp to search for and analyze the contexts in which Aššur appears together with other gods.

Our results show that the late integration of Aššur into the Mesopotamian pantheon is quite visible even in the texts written in the Neo-Assyrian period. Although Aššur is very frequently attested in the texts, he remains an outsider: many other deities have a much more central position in the networks of gods, even those that appear in the texts relatively rarely. Despite the Assyrian efforts to foster relationships and associations between Aššur and the most important gods such as Šin, Marduk, Šamaš, and Ištar, Aššur never became a truly central deity in the pantheon's network. In fact, it is likely because Aššur is so closely associated with Assyria's power and kingship and consequently appears in a wide range of contexts, often independently, that he remains relatively isolated in the network of gods, which does not rely upon him to connect major groups with one another. In essence, he sits at the top of a divine hierarchy whose gods already have long-established historical relationships and rich characters.

We hope that the open availability of the datasets and tools used in this article will facilitate future research on the rich textual material in Oracc. Currently, the more than five thousand texts from the Neo-Assyrian period lend themselves to diverse research projects applying computational methods. Our methodology yielded good results with our dataset, and it is worth investigating if the same is true for other comparable corpora as well, especially since new texts are made available in Oracc regularly and these texts are updated into Korp. Our results show that the performance of SNA can be improved by using a small window size and calculating tie strengths with PMI. Korp's flexible search tools allow the user to look for the attestations of a single word or to create more complicated queries and export the results. We demonstrated the utility of SNA and Korp with two case studies on the gods Ninurta and Nergal and the goddesses Mullissu and Šērū'a. The results of these case studies suggest that Ninurta's role in the Assyrian royal theology may have been weaker than sometimes assumed, and that our data seem to support the theory that Šērū'a was regarded as secondary to Mullissu in terms of the goddesses' relationships with Aššur. The abundant geographical information in our texts was not taken into account in this study, but this might be a fruitful avenue for future research as gods normally had close ties to certain cities.

In general, these online databases and digital methods have the potential to complement and improve upon traditional philological work. As we have demonstrated, digital methods can produce interesting results when studying a large number of sources on an aggregate level, which would be too cumbersome to go through manually. These methods can be applied to other research questions such as those that relate to the semantic fields of certain terms and relationships between other beings or places. Moreover, as more texts are added to Oracc, it will be possible to conduct diachronic and comparative studies across different periods, text genres, and geographical locations. Overall, our analysis shows that statistical and computational methods are useful tools in studying Mesopotamian culture and history and that this is an approach with substantial future potential.

References

- Allen, S. L.
2015 *The Splintered Divine: A Study of Ištar, Baal, and Yahweh Divine Names and Divine Multiplicity in the Ancient Near East*. SANER 5. Berlin: de Gruyter.
- Annus, A.
2002 *The God Ninurta in the Mythology and Royal Ideology of Ancient Mesopotamia*. SAAS 14. Helsinki: The Neo-Assyrian Text Corpus Project.
- Bastian, M.; Heymann, S.; and Jacomy, M.
2009 Gephi: An Open Source Software for Exploring and Manipulating Networks. *Proc. Third International AAAI Conference on Weblogs and Social Media*. Online: <https://www.aaai.org/ocs/index.php/ICWSM/09/paper/view/154>
- Beaulieu, P.-A.
1997 The Cult of AN.ŠĀR/Aššur in Babylonia after the Fall of the Assyrian Empire. *SAAB* 11: 55–73.
- Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E.
2008 Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics: Theory and Experiment*: P10008.
- Borgatti, S. P.; Everett, M. G.; and Freeman, L. C.
2002 *Ucinet 6 for Windows: Software for Social Network Analysis*. Harvard, MA: Analytic Technologies.
- Borin, L.; Forsberg, M.; and Roxendal, J.
2012 Korp—the Corpus Infrastructure of Språkbanken. Pp. 474–78 in *Proc. Eighth International Conference on Language Resources and Evaluation*. Online: http://www.lrec-conf.org/proceedings/lrec2012/pdf/248_Paper.pdf.
- Bouma, G.
2009 Normalized (Pointwise) Mutual Information in Collocation Extraction. Pp. 31–40 in *Von der Form zur Bedeutung: Texte automatisch verarbeiten: Proceedings of the Biennial GSCL Conference 2009*, ed. C. Chiarcos, R. E. de Castilho and M. Stede. Tübingen: Gunter Narr.
- Burt, R. S.
1992 *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.
- Church, K. W., and Hanks, P.
1990 Word Association Norms, Mutual Information, and Lexicography. *Computational Linguistics* 16: 22–29.
- Croft, W. B.; Metzler, D.; and Strohman, T.
2010 *Search Engines: Information Retrieval in Practice*. Boston: Addison-Wesley.
- Daille, B.
1994 Approche mixte pour l'extraction automatique de terminologie: Statistiques lexicales et filtres linguistiques. PhD diss., Université Paris.
- Evert, S. and the CWB Development Team
2016 *CQP Query Language Tutorial*. Online: http://cwb.sourceforge.net/files/CQP_Tutorial.pdf.
- Frahm, E.
1997 *Einleitung in die Sanherib-Inschriften*. AfO Beiheft 26. Vienna: Institut für Orientalistik der Universität Wien.
2014 Family Matters: Psychohistorical Reflections on Sennacherib and His Times. Pp. 163–222 in *Sennacherib at the Gates of Jerusalem: Story, History, and Historiography*, ed. I. Kalimi and S. Richardson. CHANE 71. Leiden: Brill.
- Frame, G.
1997 The God Assur in Babylonia. Pp. 55–64 in *Assyria 1995: Proceedings of the 10th Anniversary Symposium of the Neo-Assyrian Text Corpus Project, Helsinki, September 7–11, 1995*, ed. S. Parpola and R. M. Whiting. Helsinki: The Neo-Assyrian Text Corpus Project.
1999 My Neighbor's God: Aššur in Babylonia and Marduk in Assyria. *Canadian Society of Mesopotamian Studies Bulletin* 34: 5–22.
- George, A.
1993 *House Most High: Temples of Ancient Mesopotamia*. MC 5. Winona Lake, IN: Eisenbrauns.
- Holloway, S. W.
2002 *Aššur Is King! Aššur Is King! Religion in the Exercise of Power in the Neo-Assyrian Empire*. CHANE 10. Leiden: Brill.
- Jacomy, M.; Venturini, T.; Heymann S.; and Bastian, M.
2014 ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLoS ONE* 9: e98679.

- Jauhiainen, H.; Sahala, A.; and Alstola, T.
 2017 Open Richly Annotated Cuneiform Corpus, Korp Version, September 2017 [text corpus]. The Language Bank of Finland, FIN-CLARIN. Online: <http://urn.fi/urn:nbn:fi:lb-2018071121>.
- Jurafsky, D., and Martin, J. H.
 2017 *Speech and Language Processing*, 3rd ed. draft. Online: <https://web.stanford.edu/~jurafsky/slp3/>.
- Lambert, W. G.
 1973 *Studies in Nergal*. *BiOr* 30: 355–63.
 1983 The God Aššur. *Iraq* 45: 82–86.
- Larsen, M. T.
 2015 *Ancient Kanesh: A Merchant Colony in Bronze Age Anatolia*. New York: Cambridge University Press.
- Lassen, A. W.
 2017 The “Bull-Altar” in Old Assyrian Glyptic: A Representation of the God Assur? *Proceedings of the 2nd Kültepe International Meeting, Kültepe, 26–30 July 2015*, ed. F. Kulakoğlu and G. Barjamovic. Subartu 39. Kültepe International Meetings 2. Turnhout: Brepols.
- Maul, S. M.
 1999 Der assyrische König—Hüter der Weltordnung. Pp. 201–14 in *Priests and Officials in the Ancient Near East*, ed. K. Watanabe. Heidelberg: Winter.
 2017 Assyrian Religion. Pp. 336–58 in *A Companion to Assyria*, ed. E. Frahm. Malden, MA: Wiley-Blackwell.
- Meinhold, W.
 2014 Die Familie des Gottes Aššur. Pp. 141–49 in *La famille dans le Proche-Orient ancien: Réalités, symbolismes et images: Proceedings of the 55th Rencontre Assyriologique Internationale, Paris 6–9 July 2009*, ed. L. Marti. Winona Lake, IN: Eisenbrauns.
- Menzel, B.
 1981 *Assyrische Tempel*. 2 vols. Studia Pohl, Series Maior 10. Rome: Biblical Institute Press.
- Newman, M. E. J., and Girvan, M.
 2004 Finding and Evaluating Community Structure in Networks. *Physical Review E* 69: 026113.
- Noack, A.
 2009 Modularity Clustering Is Force-Directed Layout. *Physical Review E* 79: 026102.
- Parker, B. J.
 2011 The Construction and Performance of Kingship in the Neo-Assyrian Empire. *Journal of Anthropological Research* 67: 357–86.
- Pongratz-Leisten, B.
 2015 *Religion and Ideology in Assyria*. SANER 6. Berlin: de Gruyter.
- Prell, C.
 2012 *Social Network Analysis: History, Theory and Methodology*. Los Angeles: Sage.
- Retsö, J.
 2004 *The Arabs in Antiquity: Their History from the Assyrians to the Umayyads*. London: RoutledgeCurzon.
- Robson, E.
 2017 Project Catalogues. *Oracc: The Open Richly Annotated Cuneiform Corpus*. Online: <http://oracc.museum.upenn.edu/doc/help/managingprojects/catalogues/>
- Role, F., and Nadif, M.
 2011 Handling the Impact of Low Frequency Events on Co-occurrence Based Measures of Word Similarity: A Case Study of Pointwise Mutual Information. Pp. 226–31 in *Proc. International Conference on Knowledge Discovery and Information Retrieval (KDIR-2011)*. Online: <https://doi.org/10.5220/0003655102260231>.
- Scott, J.
 2000 *Social Network Analysis: A Handbook*, 2nd ed. London: Sage.
- Scott, J., and Carrington, P. J., eds.
 2011 *The SAGE Handbook of Social Network Analysis*. London: Sage.
- Still, B.
 2016 The Social World of the Babylonian Priest. PhD diss., Leiden University.
- Svärd, S.
 2015 *Women and Power in Neo-Assyrian Palaces*. SAAS 23. Helsinki: The Neo-Assyrian Text Corpus Project.
- Svärd, S.; Jauhiainen, H.; Sahala, A.; and Lindén, K.
 2018 Semantic Domains in Akkadian Texts. Pp. 224–56 in *CyberResearch on the Ancient Near East and Neighboring Regions: Case Studies on Archaeological Data, Objects, Texts, and Digital Archiving*, ed. V. Bigot Juloux, A. R.

- Gansell, and A. Di Ludovico. *Digital Biblical Studies 2*. Leiden: Brill.
- Tinney, S.
 2017 ATF Linkage. *Oracc: The Open Richly Annotated Cuneiform Corpus*. Online: <http://oracc.museum.upenn.edu/doc/help/editinginatf/linkage/>
- von der Osten-Sacken, E.
 2010 Aššur, großer Berg, König von Himmel und Erde. *UF 42*: 732–822.
- von Weiher, E.
 1971 *Der babylonische Gott Nergal*. AOAT 11. Kevelaer: Butzon & Bercker; Neukirchen-Vluyn: Neukirchener Verlag.
- Waerzeggers, C.
 2014a Social Network Analysis of Cuneiform Archives—a New Approach. Pp. 207–33 in *Documentary Sources in Ancient Near Eastern and Greco-Roman Economic History: Methodology and Practice*, ed. H. D. Baker and M. Jursa. Oxford: Oxbow.
 2014b *Marduk-rēmanni: Local Networks and Imperial Politics in Achaemenid Babylonia*. OLA 233. Leuven: Peeters.
- Wagner, A.; Levavi, Y.; Kedar, S.; Abraham, K.; Cohen, Y.; and Zadok, R.
 2013 Quantitative Social Network Analysis (SNA) and the Study of Cuneiform Archives: A Test-Case Based on the Murašû Archive. *Akkadica* 134: 117–34.
- Wiggermann, F. A. M.
 1999 Nergal. *RIA* 9: 215–26.
- Zaia, S.
 2015 State-Sponsored Sacrilege: “Godnapping” and Omission in Neo-Assyrian Inscriptions. *JANEH* 2: 19–54.
 2017 Official Religion in the Neo-Assyrian Royal Inscriptions. PhD diss., Yale University.