

Semantic and Cultural Networks

Sarah Shugars and Sandra González-Bailón

1. INTRODUCTION

When Queen Victoria chose to wear white to her wedding in 1840, she seeded a lasting connection in Western minds. While Western brides had previously worn a variety of colours, the Queen's colour choice sparked a fashion trend which quickly evolved into a cultural standard. The concept of 'wedding dress' now immediately evokes the concept of 'white'. There is no inherent, natural reason why these concepts must be connected to one another; this connection was instead forged through a process of imitation and social diffusion that ultimately crystallised into a cultural norm.

There are many such connections we could draw between ideas, beliefs and other types of cultural constructs. These connections might be observed through tangible artefacts – that is, words, images, or videos – or they may primarily exist as latent conceptual ties – that is, as assumptions, associations, or justifications. Despite this diversity in scope, there is a great deal of conceptual and methodological similarity in the analysis of these socially constructed networks. As with other types of networks, a key decision is which constructs are to be included as nodes, and what type of 'connection' the edges represent. In the

case of semantic and cultural networks, both can be understood as socially constructed networks in which tangible or intangible cultural artefacts are connected to one another according to some meaningful measure. By 'cultural artefacts' we mean information units – words, ideas, beliefs, images – whose meanings are co-created and used by individuals to communicate or otherwise navigate a shared culture.

As we discuss in this chapter, these socially constructed structures include conceptual and knowledge networks – for example, networks of ideas; networks of words and meanings; and networks of images, audio, or video (a more nascent research area). What all these networks have in common is that the nodes are culturally created artefacts and the edges capture socially conceived connections between those artefacts. If 'white' and 'wedding dress' are connected, it is because someone once put them together, and this co-occurrence became more frequent over time, to the point that the connection became reified in the collective imagination. This example also highlights the importance of the cultural frame of reference: in China, for instance, wedding dresses are predominantly red, a colour which symbolises good luck and happiness.

2. WHY STUDY SEMANTIC AND CULTURAL NETWORKS?

We usually think of social networks as models describing relationships between social actors, such as individuals, organisations, or countries. However, social network analysis (SNA) is more broadly understood as the structural examination of social systems (Freeman, 2004): the goal isn't merely to understand how actors are connected to one another, but to shed insight into why actors form connections and the implications of these social embeddings. The study of socially constructed networks – such as semantic and cultural networks – is deeply germane to this goal: studying cultural artefacts and the perceived connections between them has the potential to build understanding into how people interpret, represent and communicate their realities. These models can therefore allow researchers to examine why actors engage in certain behaviours and interactions and to further interrogate the collective implications of those actions.

The study of semantic and cultural networks can, for example, cast light into the biases of a given society. When semantic networks tell us that the word 'programmer' is more closely connected to the word 'male' than to the word 'female' (Bolukbasi et al., 2016) or when the concept of 'gender' is defined as a binary, thus hiding other identities, we are uncovering important information about a society, its structure and its collective understanding of the world. In other words, semantic and conceptual networks help us make tangible the biases, values and associations of the society that generated them. This is particularly important since advances in machine learning and data analysis increase the potential to exacerbate the biases embedded in these socially constructed networks (Mehrabi et al., 2021). Identifying a gendered connotation to the word 'programmer' tells us something relevant about a culture. But when automated tasks of information retrieval or machine translation leverage the data contained in semantic networks, it not only reinforces existing social biases; it also creates the illusion that these semantic connections are not merely socially constructed but a reflection of some natural truth. Uncovering these connections for what they are (i.e., socially constructed meaning) is a crucial step in deploying these automated tasks responsibly.

Socially constructed networks can also serve as powerful tools for mapping processes of human problem solving and meaning-making. For example, the education one might receive in an MBA or other professional programme is not typically about memorising a list of facts, but is more richly

geared towards learning a profession's approach to reasoning and decision-making (Shaffer et al., 2009; Shavelson, 1974). This is why case studies can be such a useful pedagogical tool – rather than learning discrete, isolated facts, students develop generalised knowledge by learning to weigh relevant trade-offs between connected ideas and approaches (Flyvbjerg, 2006). Notably, this is an inherently networked understanding of the learning process (Lynn & Bassett, 2020) – it's not just about the ideas (nodes), but also about the connections (or edges) bringing those ideas together. As students learn their profession's understanding of those connections, they are better able to address the novel problems they encounter.

These examples also allow us to introduce some of the challenges of working with semantic and cultural networks. Human processes such as learning, remembering, reasoning, arguing, explaining and communicating are often *implicitly* networked, but may not be explicitly so. In many cases, *both* the nodes and the edges are socially constructed and, therefore, are latent network properties that researchers need to infer. In the analysis of semantic networks, for instance, nodes might be easily defined as 'words' which can be directly observed; but this does not solve important conceptual issues: for instance, should homographs (i.e., words which share a spelling but not a meaning) be considered as the same node? What about words which are spelled differently but share the same meaning? Might some nodes be multiword phrases rather than single unigrams? What rules do we use for determining which words are connected? Must they be within the same span? The same sentence? The same paragraph? All of these are questions researchers must consider within the context of their data and their research question. In other words, when examining culturally constructed networks, researchers must give careful thought to how they conceptualise and operationalise their basic units of analysis (nodes and edges), and they should be mindful of the many degrees of freedom that go into their observations and measurements.

While the challenges of measuring and analysing socially constructed networks may at first seem overwhelming, it is important to remember that many of these issues also exist within other areas of SNA as well. Arguably, when the nodes in a social network are actors, the definition of the population is more concrete, as long as network boundaries are clearly delineated. Yet, the connections between those actors are often open to interpretation and may at times be intangible as well. An observed tie in a social network may represent a concrete interaction such as an email exchange, a transfer of funds, or being physically proximate

with a radio frequency identification (or RFID) badge; but these ties may equally be intangible: feelings of trust, having memorable conversations, or perceiving someone as a leader. Furthermore, ties may mean different things even if their measurement is equivalent: ties to a spouse, a dear friend, or a sibling may be equally strong (in terms of frequency of contact) but different in terms of emotional, supportive, or informational value; likewise, after years of absence some social ties may be permanently severed while others can be effortlessly reconnected. Just as there are important conceptual questions when studying semantic and cultural networks, there are similarly deep conceptual questions in actor-driven SNA that require careful consideration.

The challenges faced by researchers aiming to study semantic and cultural networks are therefore neither new nor unique. The difference is not that these networks of cultural artefacts are inherently more ineffable than other forms of social networks; rather there has simply been less time spent collectively grappling with and navigating these conceptual concerns. While scholars have raised questions about semantic and cultural networks for nearly as long as they have been examining actor-driven social networks, it is only relatively recently that computational tools have made the study of these networks tractable and scalable, raising the stakes of failing to examine the implicit biases and assumptions embedded in cultural artefacts and their relations.

In what follows, we first provide a brief overview of the philosophical and theoretical background to existing work in this area. We then discuss recent research organised along two distinct strands: networks of concepts, beliefs and knowledge; and networks of cultural artefacts such as text and images. We then provide some practical advice to help scholars navigate the methodological decisions and available computational tools.

3. THEORETICAL BACKGROUND

Some of the earliest work in SNA focused not on measuring naturally occurring social network structures, but on studying how networks of communication patterns affected group problem solving and performance (Bavelas, 1950; Guetzkow & Simon, 1955). In a series of lab-based studies, the arrangement of a social network was assumed to be imposed as a formal organisational structure. The process of interest, then, was the degree to which information could flow across this

organisational network. In other words, could agents who themselves had access only to partial information collectively reach optimal outcomes through strategic knowledge sharing? While this line of work went on to form an important foundation for SNA studies in organisational and management science (Burt, 2004; Lazer & Friedman, 2007; Mason & Watts, 2012; Riedl et al., 2021), it is notable here for its early attention to the interconnected roles of human language, cognition and social ties. Networks of social relationships don't simply exist on their own, they serve as critical pathways for information sharing and meaning-making – social processes which themselves can generate cultural artefacts which in turn can be understood as having meaningful network structure.

In focusing his research on the effectiveness of different social structures, Bavelas intentionally chose to disregard 'the nature of the communication' itself (Bavelas, 1950) – aiming to design simple tasks in which the complexity of the ideas communicated or the language used to communicate would not exert undue influence on a group's ability to successfully share disparate information. For example, in one experiment described by Bavelas, each subject received a card with five geometric symbols taken from a pool of six possible symbols. By passing messages through their designated network structure, each subject had to determine which symbol was uniquely shared across all cards. Interestingly, even when the appropriate pathways for communication were available, relevant information was not always communicated flawlessly. This suggests that the process of information sharing is not only subject to the structure of social ties, but is likely also affected by the complexity of the exchanged ideas and the language used to communicate. If errors can occur in a relatively straightforward symbol-passing task, they may be more likely to occur when messages are more complex. Indeed, this is essentially the premise of the children's game Telephone, in which a starting player whispers a phrase to another player and the message is passed down the line until it is inevitably jumbled through this process of iterated exchange. Furthermore, social ties may not merely be conduits for communication, but communication itself may serve to forge and maintain social ties – again, emphasising the need to take these cultural constructs into account (Doerfel & Moore, 2016).

Given the complexity of natural language and real-world knowledge-sharing tasks, work in linguistics, computational linguistics and linguistic philosophy has long tackled the related question of what information exchange actually looks like. Consider, for example, the simple model of human

communication proposed by Cohen (1987). At minimum, communication requires a dyadic interaction between a single Speaker and a single Hearer. The Speaker sends an information signal, and the Hearer interprets that information signal. While this signal could be as simple as a written series of geometric shapes – as in the lab studies of Bavelas (1950) – a common but more complex message could take the form of an ‘argument’. That is, we might imagine the Speaker transmits some message with the goal of convincing the Hearer of some piece of information contained within that message. The Speaker’s job is to construct the message in such a way as to be interpretable, and the Hearer’s job is to interpret that message to understand the argument the Speaker is making. For such an exchange to be successful, both the Speaker and the Hearer must share a common set of rules for encoding and decoding any transmitted message: they must share a language which dictates the message’s structure.

While the ‘social network’ in this model may be trivial – two actors who are assumed to communicate freely and share information in good faith – the simplicity of the social dimension allows for closer scrutiny of the linguistic dimension. Specifically, the key insight offered here is that in order for this process of argument encoding and decoding to be successful, the argument itself must have a predictable structure (Shannon, 1948).

This is a very old idea which dates back at least as far as Aristotle, who described how a major premise (‘All men are mortal’) may be connected to a minor premise (‘Socrates is a man’) in order to deductively reach a valid conclusion (‘Socrates is mortal’) (Mill, 1882). This is an inherently networked understanding of argument structure: these statements do not stand on their own, but are connected in such a way as to justify the final conclusion. This networked approach to argument structure has continued through modern linguistic philosophy (Toulmin, 1958; Walton et al., 2008) and has most recently served as the basis for machine learning approaches for argument detection and argument mining (Feng & Hirst, 2011; Habernal & Gurevych, 2015; Mochales & Moens, 2011; Palau & Moens, 2009; Stab & Gurevych, 2014).

Importantly, the mere presence of argument structure is not enough to ensure successful communication. In order for a Hearer – either human or algorithmic – to accurately interpret the argument of a Speaker, they must be able to both interpret the individual statements made and to recognise the connections between those statements (Cohen, 1987). Observable linguistic cues may help a Hearer identify connections and detect

argument structure (Cohen, 1987; Mochales & Moens, 2011), but it takes a process of cognition or memory retrieval in order to assess the meaning or veracity of each individual statement. After all, an argument which begins with the premise that ‘all men are *immortal*’ and concludes with the claim that ‘Socrates is *immortal*’ would have the same *structure* as an argument in which all men are assumed to be mortal. A Hearer aiming to determine which of these statements to believe would therefore need some means of assessing the true relationship between the concept of ‘men’ and the concept of ‘mortality’.

4. COGNITIVE, CONCEPTUAL AND BELIEF NETWORKS

In this sense, knowledge itself can be understood as having a network structure. In making sense of the world, humans do more than simply accrue long lists of declarative facts (Dorsey et al., 1999). We *organise* our knowledge, understanding concepts through their relationships with other concepts and leveraging those connections in order to efficiently store and retrieve information (Collins & Loftus, 1975; Collins & Quillian, 1969; Dorsey et al., 1999). Importantly, these connections reflect a diversity of types of knowledge, including natural facts (i.e., men are mortal) and cultural norms (i.e., wedding dresses are white). People may not always be explicitly aware that they are making or using network structures in their cognitive processes, but humans are remarkably adept at learning these network structures and identifying the resulting patterns (Lynn & Bassett, 2020). This suggests that networked-based models can be valuable for studying a range of cognitive processes. The chapter ‘Cognition and Social Networks’ (Brashears & Money, this volume) offers a more in-depth analysis of how networks help us uncover the logic of cognitive processes. Here, we focus on those processes as they intersect with the generation of cultural meaning and communicative practices.

Scholars have therefore taken different approaches to conceptualising and operationalising networks related to diverse cognitive processes including reasoning, remembering, arguing and learning. One line of work in public opinion, for example, has leveraged survey methods to infer mass ideology – that is, belief systems of connected political attitudes (Baldassarri & Goldberg, 2014; DellaPosta et al., 2015; Boutyline & Vaisey 2017; DellaPosta, 2020; Fishman & Davis, 2021). In these networks, the nodes are beliefs, measured

as responses to survey policy questions (i.e., attitude towards gay marriage, affirmative action, or defence spending), and the edges are undirected ties measuring correlation, proximity, or covariance (i.e., clustering in those beliefs). Analysis of these structures allows identifying longitudinal trends in public opinion. For instance, increasing density in these networks or changes in the centrality of nodes suggest shifts in public opinion that cannot be captured by just looking at changes in support around individual policy items. The analysis of belief networks also helps identify sub-populations – it helps separate groups of people depending on how similar their belief structures are. It can also help identify organising heuristics that guide how people filter the information they encounter. Emerging work in this area has further built upon survey methods to not only measure the belief nodes, but to also measure the edges – directly asking respondents about the connections they see, or don't see, between their policy stances and their social ideals.

One of the benefits of using survey instruments to elicit beliefs is that they make tangible what exists only in the mind of a subject; they also help extrapolate those beliefs to entire populations. Belief networks help track shifts in public opinion and mass belief systems, which in turn can contribute to the study of political rhetoric, policy framing and issue salience (Yang & González-Bailón, 2016). Opinion surveys, however, are only useful for the items included in the survey, which may not represent the whole range of beliefs deemed important to subjects themselves. The explosion of digital data has given rise to new computational opportunities to be creative with the measurement of these networks. For instance, we can now analyse large-scale textual data that includes political news coverage, transcripts of political statements and debates, and social media posts and commentary from the public (Yang & González-Bailón, 2016). This ever-expanding corpus of written expression is enlarging the territory we can map with belief networks. But digital technologies also allow us to go beyond written communication.

Communication is essential to share information, ideas, or beliefs, and it makes use of all sorts of cultural artefacts capable of conveying information from one person to another. When asked for their opinion, a person might use words to describe their views (Shugars, 2020) but there are other forms of expressing that opinion. For instance, when aiming to educate, a person may use a combination of written, audio, or visual information to teach (Shavelson, 1974). When working in groups, individuals may take cues from each other to converge to a common way of speaking (Saint-Charles & Mongeau, 2017). In other words, as we

aim to complete a task, seek to educate, or simply communicate with one another, we generate observable trace data in the form of interconnected cultural artefacts that manifest themselves beyond words. In the digital age, these artefacts also leave a digital trail that can be more easily parsed for measurement and analysis.

These observable trace data may then be used to meaningfully infer a subject's underlying conceptual network structure – or, in other words, their mental models (Carley & Palmquist, 1992). For instance, while we can't directly observe someone's political reasoning, we can observe the language they use when explaining or justifying their views. This opens up opportunities for inferring conceptual network structure from free response text or other documents in which people share political opinions (Atteveldt, 2008; Axelrod, 1976; Shugars, 2020).

Words themselves are cultural artefacts which represent or communicate some underlying concept. In this sense, observed words can be interpreted as representations of the latent concepts they seek to express – concepts which can then be treated as nodes in a conceptual network structure. Edges between these concepts can be understood based on observed co-occurrence within some designated span, or based on the grammatical structure of a text itself – after all, grammar is fundamentally a cultural tool aimed at signalling to a Hearer how expressed concepts are logically connected (Cohen, 1987; Shugars, 2021).

Once these conceptual networks are constructed, researchers can determine the degree to which an individual's network exhibits certain structural properties – such as connectivity, density, or clustering. Existing work suggests these structural properties may be meaningfully correlated with personality measures – providing a means to examine *how* people reason about political topics separately from the *content* of those reasons (Shugars, 2021). Changes in the structural properties of conceptual networks also signal shifts in mass opinion and the boundaries in political conflict (Yang & González-Bailón, 2016). The substantive meaning of those changes, of course, is contingent on measurement.

Inferring network structure from text requires making first several important methodological choices. In the context of conceptual networks, for example, not every word necessarily represents a unique or meaningful concept – some words may serve to represent the same concept while others contain no real meaning. Identifying this second type of word can be accomplished through the use of stopword lists or by simply excluding the parts of speech typically associated with stopwords (i.e., articles, conjunctions). Determining which

words refer to the same concept, on the other hand, requires a measure of word similarity – a measure which would then allow ‘similar’ words to be clustered together and interpreted as the same concept or node (Shugars, 2021).

One popular means of determining word similarity comes in the form of word embeddings, which themselves are based on a networked understanding of linguistic structure. When linguist Robert Firth observed that ‘you shall know a word by the company it keeps’ (Firth, 1957), he was implicitly making an argument for inferring semantic networks from text: a word (node) is connected (edge) to nearby words, and those connections convey meaningful information about the word itself. We might imagine some Hearer – either human or computer – ‘learning’ a word’s meaning by seeing it used in context many times. Semantic knowledge bases (Navigli & Ponzetto, 2012; Speer & Havasi, 2012) are premised on this idea – collecting large corpora of text in order to build near-complete, multilingual semantic networks. These knowledge bases can then serve as an ‘encyclopaedic dictionary’ (Navigli & Ponzetto, 2012) in which a user can determine a word’s varied meanings by examining that word’s (node’s) connected words and concepts. More recently, this semantic network idea has served as the foundation for word embeddings (Devlin et al., 2019; Levy & Goldberg, 2014; Mikolov et al., 2013). This *natural language processing* (NLP) technique ‘embeds’ words in high-dimensional space, representing each word as a vector and calculating those vectors in such a way that words which appear in similar contexts are represented by similar vectors. Using extremely large corpora such as Wikipedia (Devlin et al., 2019) or Google News (Mikolov et al., 2013) these mathematical representations of words allow for meaningful insight into word meanings, connections and biases.

5. NETWORKS OF CULTURAL ARTEFACTS

Text, images and other forms of cultural expression are therefore another key focus of research in the analysis of socially constructed networks. Instead of trying to gauge intangible constructs or cognitive concepts, this stream of research centres on communication and what is revealed through communication processes. For instance, the combination of NLP and network analysis tools has allowed researchers to identify cultural bridges in how advocacy organisations engage with the public in social media (Bail, 2016). The nodes in this network are advocacy groups, and the

weighted edges between these organisations measure the overlap of nouns and noun phrases in the posts they publish on social media. A measure of ‘cultural betweenness’ is then calculated for each organisation (using betweenness centrality), which allows testing the hypothesis that organisations that build more cultural bridges (i.e., organisations with high centrality scores because they connect typically unrelated themes) get more engagement from their audiences.

In another recent study, semantic networks are used to analyse the evolution of American politics as reflected in the presidents’ State of the Union addresses over the 1790–2014 period (Rule et al., 2015). Methodologically, this research builds and analyses networks of words co-occurring over time, which reveals evolving semantic neighbourhoods for specific terms (i.e., ‘ideals’ or ‘constitution’). These terms are, again, obtained using NLP techniques. Proximity scores are then computed to measure the relatedness of each pair of those terms, based on their co-occurrence in the same paragraph in a document published at a specific time. A community detection algorithm is then used to identify cohesive subsets of words; the clusters that emerge are treated as discursive categories. The temporal analysis of these networks and the discursive categories they reveal (and how they change over the centuries) allows for the identification of historical transitions in the evolution of American political thought.

More generally, this research is an example of how semantic networks allow us to extract not just mental models (as discussed in the previous section) but also culture from texts (Carley, 1994; Carley & Kaufer, 1993). Research looking at narrative structures offers another example of how to use network analysis to reveal symbolic representations across domains – autobiographical discourse (Bearman & Stovel, 2000), literature (Franzosi, 1998, 2004), or historical analysis (Bearman et al., 1999). Networks, in this case, encapsulate a sequence of events, as connected by people reflecting on their personal growth (or devolution); by writers in crafting their stories; or by historians in making sense of the past. Nodes in these networks represent events, and the ties help map how these events are encased chronologically or in narrative time.

Most of this research relies on text, or symbolic categories derived from text, but there is no reason why other forms of cultural expression, like images, could not be used. While this is an incipient area of research, advances in computer vision techniques have allowed the development of new measurement tools for large-scale analysis of images (Williams et al., 2020; Chen et al., 2021). Visual communication has become ever

more central in digitally mediated interactions, and the automated analysis of images allows scaling up thematic coding and the extraction of units of meaning (i.e., objects, people, affect) that constitute the building blocks of networks. Although there is not a lot of research applying network tools to the analysis of images (i.e., how they co-appear, or how they cluster in terms of their visual features), we anticipate this will become a rising area of interest.

6. THE SOCIAL CONSTRUCTION OF SEMANTIC AND CULTURAL NETWORKS

The analysis of cultural networks aims to uncover symbolic connections intrinsic to culture and mental models of the world; but the research process is itself a social construction to the extent that it depends on choices and subjective decisions. There are two key questions in this process: (1) how to delimit the research domain (i.e., how to delimit the corpora and sampling frame); and (2) how to operationalise key constructs (i.e., what counts as a tie in the network). In this section, we outline the steps involved in the analysis of

semantic and cultural networks, offering both a summary of the research discussed so far and an entry-point guide to research design in this area.

Figure 14.1 offers a schematic representation of the basic steps involved in the analysis of semantic networks. At the outset, there is the data collection stage, where the corpora of documents or artefacts to be analysed is defined; sampling procedures are executed; and the identification and extraction of units of meaning are performed. These units of meaning can be words, categories, or actors (or any other social construct), and they constitute the building blocks of the networks that are then assembled in step 2. The key decision in this step is how to define the edges in the network. Connections can signal affiliation ties (as when certain actors are linked to certain concepts); they can signal co-occurrence (with stronger ties connecting, say, words that co-appear more frequently); or they can signal a sequence in a larger narrative (as when concepts evolve in historical discourse). Once the network is assembled, the key decision is how to extract the most relevant information from that structure. Past research has used statistics like density, centrality, or structural constraint to identify meaningful changes in public discourse over time or the actors in more advantageous positions to influence the public or enact

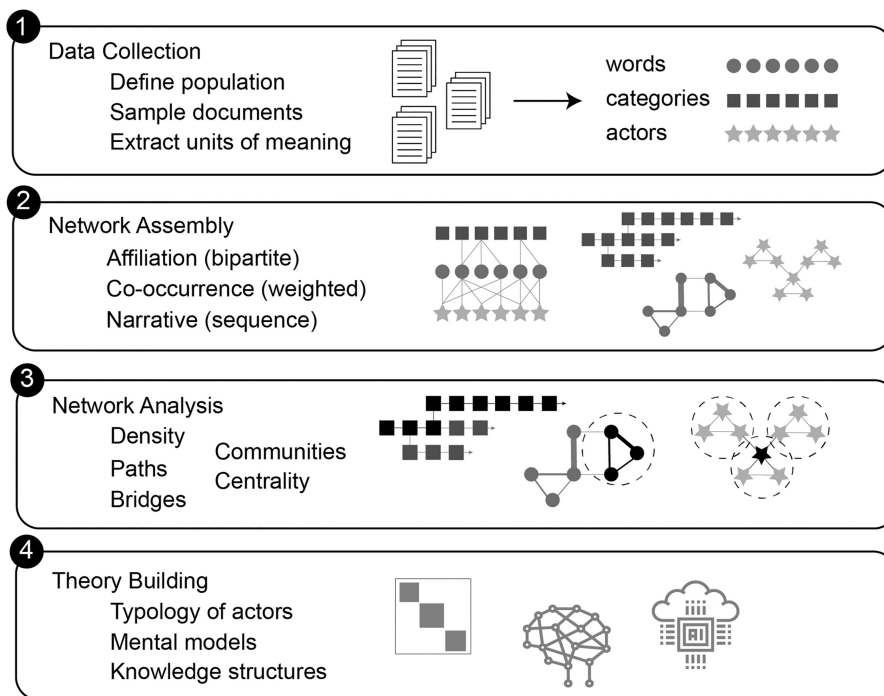


Figure 14.1 Schema of four analytical steps in semantic network research

change. Ultimately the purpose of these analyses is to improve our theoretical understanding of belief and cultural systems, and map knowledge structures as they relate to actions and behaviour – and, increasingly, to automate decisions and AI technologies.

The idea that humans store and retrieve information in a relational way (i.e., using network structures) has served as an important theoretical model for work in artificial intelligence (Collins & Loftus, 1975; Shapiro & Rapaport, 1987). A researcher can teach a computer to ‘reason’ by supplying or iteratively ‘teaching’ it appropriately structured knowledge, as inferred from human logic. For example, both a human and a computer can accurately determine that the statement ‘a canary can fly’ is true – if they have the additional, connected knowledge that (1) a canary is a bird and (2) birds can fly (Collins & Quillian, 1969). While observing a human’s process for determining the truth of this statement can be challenging, a ‘conceptual network’ coded into a computer *can* be directly observed. Furthermore, such semantic network retrieval has been shown to be an efficient way to allow computers to make meaningful inference given new information (Navigli & Ponzetto, 2012; Shapiro & Rapaport, 1987; Speer & Havasi, 2012). For example, given the additional piece of information that a hawk is a bird, a computer with a networked knowledge base could correctly determine that hawks can fly, despite not being directly taught that flight is a property of hawks.

The intimate connection between human reasoning structures and the observable artefacts of human communication has fuelled further advances in machine learning and artificial intelligence. Indeed, the ‘data mining’ of modern machine learning is premised on the idea that meaningful patterns can be found within human-generated data – even within so-called ‘unstructured’ data such as text, images, audio and video. While the data of these cultural artefacts is ‘unstructured’ insofar as it does not come ready-made into a tidy data table, the fact that algorithms are able to detect and identify meaningful patterns suggests that these data are, in fact, structured in some sense of the word. Of course, the impact that cultural stereotypes and biases have on the structuring of the data can only be analysed through the data itself and a critical inspection of the data generation mechanisms (O’Neil, 2016; Buolamwini & Gebri, 2018; Kearns & Roth, 2020). When an algorithm learns that ‘white’ and ‘wedding dress’ are associated concepts, the algorithm is really learning is a social construction, not an ontological property of the world. And while this may be an innocuous association, many others can be more harmful and generate downstream consequences if

not properly interrogated prior to the deployment of AI tools. This is an area of theoretical development that will only become more relevant as AI technologies expand their reach and their societal impact becomes more salient.

To sum up, the process of analysing semantic and cultural networks requires answering four interrelated questions: What is the process to be studied? What is the simplest way of modelling this process? What can be observed and what must be inferred? And what is the theoretical motivation driving the work?

7. CONCLUSION

The analysis of semantic and cultural networks adds an important layer to the analysis of interpersonal communication by capturing cognitive and symbolic interdependencies that shape meaning – and therefore behaviour. Beyond advancing fundamental research, uncovering the information contained in these structures also facilitates the development of machine learning and AI tools deployed in a range of social settings (from search engines to targeted recommendations). The theoretical value of research in this area is, consequently, extending to all the domains in which these technologies are being applied. An already burgeoning area of research, we anticipate the analysis of symbolic and conceptual networks will become even more popular in the next few years as it absorbs a wider range of cultural artefacts created and stored in digital form.

REFERENCES

- Atteveldt, W. van (2008). *Semantic network analysis: techniques for extracting, representing, and querying media content*. BookSurge
- Axelrod, R. (1976). *Structure of decision: the cognitive maps of political elites*. Princeton University Press.
- Bail, C.A. (2016). Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*, 113(42), 11823–11828. doi.org/10.1073/pnas.1607151113
- Baldassarri, D., & Goldberg, A. (2014). Neither ideologies nor agnostics: alternative voters’ belief system in an age of partisan politics. *American Journal of Sociology*, 120(1), 45–95. doi.org/10.1086/676042

- Bavelas, A. (1950). Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America*, 22, 725.
- Bearman, P., Faris, R., & Moody, J. (1999). Blocking the future: new solutions for old problems in historical social science. *Social Science History*, 23(4), 501–533.
- Bearman, P.S., & Stovel, K. (2000). Becoming a Nazi: a model for narrative networks. *Poetics*, 27(2), 69–90. doi.org/10.1016/S0304-422X(99)00022-4
- Bolukbasi, T., Chang, K.-W., Zou, J.Y., Saligrama, V., & Kalai, A.T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29. proceedings.neurips.cc/paper/2016/hash/a486cd07e4ac3d270571622f-4f316ec5-Abstract.html
- Boutyline, A., & Vaisey, S. (2017). Belief network analysis: a relational approach to understanding the structure of attitudes. *American Journal of Sociology*, 122(5), 1371–1447. www.journals.uchicago.edu/doi/abs/10.1086/691274
- Buolamwini, J., & Gebru, T. (2018). Gender shades: intersectional accuracy disparities in commercial gender classification. Paper presented at the Conference on Fairness, Accountability and Transparency.
- Burt, R.S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2), 349–399.
- Carley, K. (1994). Extracting culture through textual analysis. *Poetics*, 22(4), 291–312. doi.org/10.1016/0304-422X(94)90011-6
- Carley, K.M., & Kaufer, D.S. (1993). Semantic connectivity: an approach for analyzing symbols in semantic networks. *Communication Theory*, 3(3), 183–213. doi.org/10.1111/j.1468-2885.1993.tb00070.x
- Carley, K., & Palmquist, M. (1992). Extracting, representing, and analyzing mental models. *Social Forces*, 70(3), 601–636. doi.org/10.2307/2579746
- Chen, Y., Sherren, K., Smit, M., & Lee, K.Y. (2021) Using social media images as data in social science research. *New Media and Society*, 14614448211038761. doi.org/10.1177/14614448211038761
- Cohen, R. (1987). Analyzing the structure of argumentative discourse. *Computational Linguistics*, 13(1–2), 11–24.
- Collins, A.M., & Loftus, E.F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407.
- Collins, A.M., & Quillian, M.R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8(2), 240–247. doi.org/10.1016/S0022-5371(69)80069-1
- DellaPosta, D. (2020). Pluralistic collapse: the ‘oil spill’ model of mass opinion polarization. *American Sociological Review*, 85(3), 507–536. journals.sagepub.com/doi/abs/10.1177/0003122420922989
- DellaPosta, D., Shi, Y., & Macy, M. (2015). Why do liberals drink lattes? *American Journal of Sociology*, 120(5), 1473–1511. www.journals.uchicago.edu/doi/abs/10.1086/681254
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. *ArXiv:1810.04805 [Cs]*. arxiv.org/abs/1810.04805
- Doerfel, M.L., & Moore, P.J. (2016). Digitizing strength of weak ties: understanding social network relationships through online discourse analysis. *Annals of the International Communication Association*, 40(1), 127–148. doi.org/10.1080/23808985.2015.11735258
- Dorsey, D.W., Campbell, G.E., Foster, L.L., & Miles, D.E. (1999). Assessing knowledge structures: relations with experience and posttraining performance. *Human Performance*, 12(1), 31–57.
- Feng, V.W., & Hirst, G. (2011). Classifying arguments by scheme. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 987–996. www.aclweb.org/anthology/P11-1099
- Firth, J.R. (1957). *Studies in linguistic analysis*. Wiley-Blackwell.
- Fishman, N., & Davis, N.T. (2021). Change we can believe in: structural and content dynamics within belief networks. *American Journal of Political Science*. doi.org/10.1111/ajps.12626
- Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12(2), 219–245. doi.org/10.1177/1077800405284363
- Franzosi, R. (1998). Narrative analysis-or why (and how) sociologists should be interested in narrative. *Annual Review of Sociology*, 24, 517–554.
- Franzosi, R. (2004). *From words to numbers narrative, data, and social science*. Cambridge University Press. www.cambridge.org/us/academic/subjects/sociology/sociology-general-interest/words-numbers-narrative-data-and-social-science, www.cambridge.org/us/academic/subjects/sociology/sociology-general-interest
- Freeman, L. (2004). *The development of social network analysis*. Empirical Press.
- Guetzkow, H., & Simon, H.A. (1955). The impact of certain communication nets upon organization and performance in task-oriented groups. *Management Science*, 1(3–4), 233–250.
- Habernal, I., & Gurevych, I. (2015). Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2127–2137.
- Kearns, M., & Roth, A. (2020). *The ethical algorithm*. Oxford University Press.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
- Levy, O., & Goldberg, Y. (2014). Dependency-based word embeddings. *Proceedings of the 52nd*

- Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 302–308. doi.org/10.3115/v1/P14-2050
- Lynn, C.W., & Bassett, D.S. (2020). How humans learn and represent networks. *Proceedings of the National Academy of Sciences*, 117(47), 29407–29415. doi.org/10.1073/pnas.1912328117
- Mason, W., & Watts, D.J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3), 764–769. doi.org/10.1073/pnas.1110069108
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), Article 1151, 1–35. doi.org/10.1145/3457607
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *ArXiv Preprint ArXiv:1301.3781*.
- Mill, J.S. (1882). *A system of logic, ratiocinative and inductive* (8th ed.). Harper & Brothers, Project Gutenberg. www.gutenberg.org/files/27942/27942-h/27942-h.htm
- Mochales, R., & Moens, M.-F. (2011). Argumentation mining. *Artificial Intelligence and Law*, 19(1), 1–22.
- Navigli, R., & Ponzetto, S.P. (2012). BabelNet: the automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193, 217–250.
- O'Neil, C. (2016). *Weapons of math destruction: how big data increases inequality and threatens democracy*. Penguin/Random House. books.google.com/books?id=60n0DAAAQBAJ
- Palau, R.M., & Moens, M.-F. (2009). Argumentation mining: the detection, classification and structure of arguments in text. *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, 98–107.
- Riedl, C., Kim, Y.J., Gupta, P., Malone, T.W., & Woolley, A.W. (2021). Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21), e2005737118. doi.org/10.1073/pnas.2005737118
- Rule, A., Cointet, J.-P., & Bearman, P.S. (2015). Lexical shifts, substantive changes, and continuity in State of the Union discourse, 1790–2014. *Proceedings of the National Academy of Sciences*, 112(35), 10837–10844. doi.org/10.1073/pnas.1512221112
- Saint-Charles, J., & Mongeau, P. (2017). Social influence and discourse similarity networks in workgroups. *Social Networks*. doi.org/10.1016/j.socnet.2017.09.001
- Shaffer, D.W., Hatfield, D., Svarovsky, G.N., Nash, P., Nulty, A., Bagley, E., Frank, K., Rupp, A.A., & Mislavy, R. (2009). Epistemic network analysis: a prototype for 21st-century assessment of learning. *International Journal of Learning and Media*, 1(2), 33–53.
- Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
- Shapiro, S.C., & Rapaport, W.J. (1987). SNePS considered as a fully intensional propositional semantic network. In N. Cercone & G. McCalla (Eds.), *The knowledge frontier* (pp. 262–315). Springer.
- Shavelson, R.J. (1974). Methods for examining representations of a subject-matter structure in a student's memory. *Journal of Research in Science Teaching*, 11(3), 231–249.
- Shugars, S. (2020). Reasoning together: network methods for political talk and normative reasoning. PhD thesis. Northeastern University.
- Shugars, S. (2021). *The structure of reasoning: inferring conceptual networks from short text*. doi.org/10.17605/OSF.IO/PNWD8
- Speer, R., & Havasi, C. (2012). Representing general relational knowledge in ConceptNet 5. *LREC*, 3679–3686.
- Stab, C., & Gurevych, I. (2014). Identifying argumentative discourse structures in persuasive essays. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 46–56. www.aclweb.org/anthology/D14-1006
- Toulmin, S.E. (1958). *The uses of argument*. Cambridge University Press.
- Walton, D.N., Reed, C., & Macagno, F. (2008). *Argumentation schemes*. Cambridge University Press.
- Williams, N.W., Casas, A., & Wilkerson, J.D. (2020). Images as data for social science research: an introduction to convolutional neural nets for image classification. Cambridge University Press.
- Yang, S., & González-Bailón, S. (2016). Semantic networks and applications in public opinion research. In M. Lubell, A.H. Montgomery & J.N. Victor (Eds.), *Handbook of political networks*. Oxford University Press.