Departmental Papers (ASC)   Annenberg School for Communication

2018

Semantic Networks and Applications in Public Opinion Research

Sijia Yang
University of Pennsylvania, sijia.yang@asc.upenn.edu

Sandra González-Bailón
University of Pennsylvania, sgonzalezbailon@asc.upenn.edu

Semantic Networks and Applications in Public Opinion Research

Disciplines
Communication | Social and Behavioral Sciences

....

SEMANTIC NETWORKS AND APPLICATIONS IN PUBLIC OPINION RESEARCH

....

SIJIA YANG AND SANDRA GONZÁLEZ-BAILÓN

INTRODUCTION

On June 26, 2015, the US Supreme Court issued a landmark ruling, granting same-sex couples the constitutional right to marry across states. This court decision was made against the backdrop of already shifting public opinion: while only eleven states plus the District of Columbia recognized the right to same-sex marriage when the ruling was issued, more than half of Americans supported their homosexual peers' right to marry legally (Liptak, 2015). In 2001 the ratio of supporters versus opponents to same-sex marriage was 35 percent versus 57 percent (Pew, 2015); the difference had been as great as 12.6 percent versus 71.9 percent in 1988 (Baunach, 2012). What has happened over the past two and a half decades to reverse the state of public opinion about same-sex marriage? How do people understand the issue of marriage equality, as well as their position on the debate? What are the reasons, beliefs, and values underlying people's aggregated responses to the support/oppose question asked in polls and surveys, the instruments commonly used to measure public opinion? And how are those cognitions connected to each other in public discussions and political debates? Do those associations vary across demographic groups? The analysis of semantic networks can help answer these and related questions.

Semantic network analysis offers a representational framework and a set of modeling strategies to analyze language and the opinions expressed as a relational structure (van Atteveldt, 2008; Baden, 2010; Borge-Holthoefer and Arenas, 2010; Carley and Kaufer, 1993; Carley and Palmquist, 1992; Corman et al., 2002; Danowski, 2009; Diesner and Carley, 2010; Doerfel and Connaughton, 2009; Fisher, Leifeld, and Iwaki, 2012; Popping, 2006; Steyvers and Tenenbaum, 2005). Semantic networks offer a

quantitative approach to discourse and language that can also accommodate qualitative tools for data collection, analysis, and interpretation. Semantic networks differ from social network analysis in that the nodes are not social actors (e.g., citizens, non-governmental organizations [NGOs], political parties), but semantic concepts (e.g., names, places, organizations, policies, values); ties are not social relationships (e.g., friendships) but associations between concepts (e.g., co-occurrence). However, the same metrics and analytical procedures used to analyze social networks can be applied to the study of semantic networks, with the caveat that interpretation will vary depending on the domain of the data.

Since large-scale textual data such as political news coverage, transcripts of politicians' statements and debates, and social media posts and commentary from the public are readily available in today's digital era, the chapter offers an introduction to commonly used natural language processing (NLP) and text mining techniques that can help researchers exploit those data sources. These automated methods can be employed to extract from raw textual data semantic information that is necessary to construct a semantic network, and they complement more traditional approaches like manual coding (now also made more scalable thanks to crowdsourcing platforms; see Benoit et al., 2016). We also offer a discussion of methodological choices that are required to connect research questions with data. These include deciding on the level of data collection (individual, interpersonal, or collective level), the abstraction of the semantic concepts (e.g., words, topics, or themes), the type of association (e.g., based on co-occurrence or other types of semantic relationships, such as causal connections), and the metrics to summarize the structural properties of the network (e.g., centrality scores, community detection).

We start by discussing state-of-the-art techniques for the construction and analysis of semantic networks, paying special attention to this approach's substantive contribution to theory building. We end the chapter with a discussion of potential applications in public opinion research, emphasizing how semantic network analysis is particularly suitable to represent and explain public opinion as conceptualized by discursive and deliberative theories of democracy (Cappella, Price, and Nir, 2002; Carpini, Cook, and Jacobs, 2004; Crespi, 1997; Gutmann and Thompson, 2004; Mendelberg, 2002; Price and Neijens, 1997). Since research on semantic network analysis is still nascent, and studies applying this methodology to public opinion research are just burgeoning, we also gather insights from applications in adjacent domains. Ultimately, our goal is to introduce this methodology to political scientists (and, more generally, social scientists) to showcase its potential to develop public opinion research in new and exciting directions.

WHAT IS A SEMANTIC NETWORK?

Networks offer a general and flexible representational system that captures interdependence among entities and helps model those patterns of interconnection

(Borgatti, Martin, and Johnson, 2013; Newman, Barabási, and Watts, 2006; Scott, 2012; Wasserman and Faust, 1994). In the analysis of semantic networks, the first and most important question is: What semantic units need to be represented? The question that follows is: What relationships need to be mapped? In other words, researchers start by identifying the nodes to be analyzed and operationalizing the definition of a tie. The answer to these questions depends on the research goals and the availability of data. In general, though, there are two main possibilities: to define networks that map relationships between semantic units (e.g., co-occurrence of words in political discourse); and to define networks that map the association of actors with concepts (e.g., policymakers and their statements). Figure 13.1 illustrates these two types of networks.
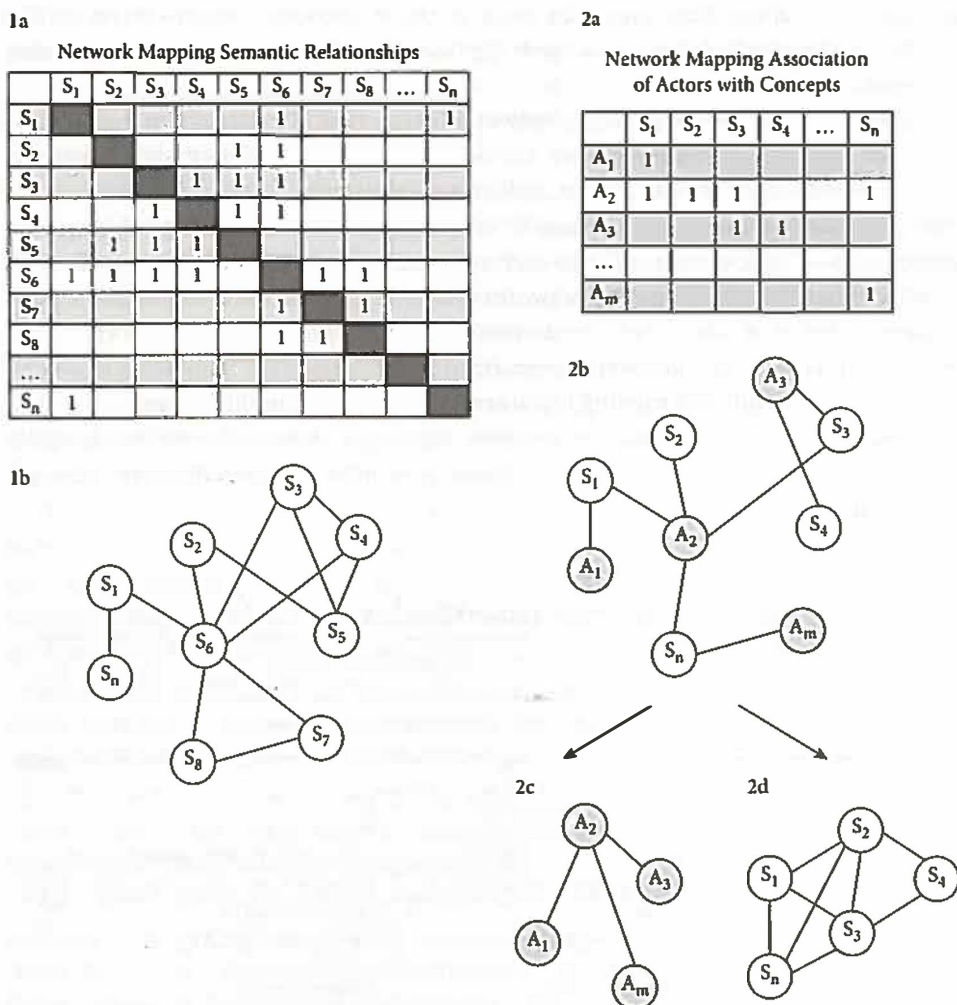
1a

Network Mapping Semantic Relationships

	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8	...	S_n
S_1				1						1
S_2				1	1					
S_3				1	1	1				
S_4			1		1	1				
S_5		1	1	1						
S_6	1	1	1	1			1	1		
S_7				1		1				
S_8				1	1					
...										
S_n	1									

2a

Network Mapping Association of Actors with Concepts

	S_1	S_2	S_3	S_4	...	S_n
A_1	1					
A_2	1	1	1			1
A_3			1	1		
...						
A_m						1

1b

2b

2c

2d

FIGURE 13.1 Types of Semantic Networks.

Panels 1a and 1b in **Figure 13.1** show a schematic representation of a network of semantic relationships, both as a square matrix $S_{n \times n}$ and as a graph. Since the matrix is symmetrical (i.e., the upper and lower triangles are identical), the network is undirected. The cell values could indicate a binary semantic relationship (e.g., two concepts co-occur or not) or a continuous count, in which case the network would be weighted and the cell values would indicate to what degree two concepts are semantically related (e.g., the number of times two concepts co-occur). Panels 2a-2b show an actor-concept network, again in its matrix form, $A_{m \times n}$, and as a bipartite graph. Here ties measure an actor's level of endorsement of a semantic unit or concept; again, cell values can either indicate endorsement as a binary relation or give a measure of strength. As with any bipartite network, there are two types of network projections: the actor-to-actor network (panel 2c), in which ties indicate how many concepts any two actors share in their discourses, and the concept-to-concept network (panel 2d), in which ties indicate the number of actors that commonly use any pair of concepts. We expand on each of these representations and illustrate their applications in political science in the following sections.

Since the literature on semantic networks spans multiple disciplines and lacks a coherent theoretical framework, we decided to structure our discussion along two dimensions: the *level of analysis* (i.e., individual, interpersonal, and collective) and the *type of semantic network* (i.e., a network mapping semantic relationships or a network mapping associations between actors and concepts, as illustrated in figure 13.1). To illustrate how the analysis of semantic networks can make a significant contribution to public opinion research, we discuss applications in four of the six possible cells created by the two dimensions we consider, summarized in table 13.1. We introduce and discuss examples drawn from the existing literature and assess their implications for scientific progress in the following domains: cognitive mapping of opinion formation, discourse analysis with concept-to-concept and actor-to-actor network projections, and research on framing and issue salience.

Table 13.1 Examples of Semantic Network Analysis

Level of Analysis	Type of Application	
	Concept-to-Concept	Actor-to-Actor
Individual	cognitive mapping (figure 13.1, panel 1)	
Interpersonal	discourse network analysis—concept congruence/conflict network (figure 13.1, panel 2d)	discourse network analysis—actor congruence/conflict network; discursive fields (figure 13.1, panel 2c)
Collective	salience and framing (figure 13.1, panel 1)	future research

Networks Mapping Sematic Relationships

Cognitive mapping, salience, and framing research relies on networks that only consider relationships among semantic units or concepts; that is, such applications do not consider actor-concept endorsement. While *cognitive mapping* treats the individual as the basic unit of analysis and focuses on the mental representations of external entities (e.g., political issues like same-sex marriage), the other two applications operate on aggregated semantic data produced by a collective of individuals or social actors. For example, by analyzing media coverage of the issue of same-sex marriage rights over time, media frames can be identified from aggregated semantic data of news coverage; similarly, frames adopted by the general public can also be extracted from social media by aggregating the opinions expressed publicly. The distinction between *salience* (e.g., network agenda setting; see Guo, 2012; Vargo et al., 2014) and *framing* (Baden, 2010; Miller, 1997) research at the collective level on the one hand, and *discourse network analysis* (Leifeld, 2013; Leifeld and Haunss, 2012) at the interpersonal level on the other, is that at the interpersonal level, information on who issued what expressions or endorsed what beliefs is retained as a key part of the analysis, but is not explicitly considered in research at the collective level.

Cognitive mapping and its closely related application *mental model mapping* are commonly used by computational linguists and cognitive psychologists. Compared with cognitive mapping, which focuses more on interdependence among lower-level concepts, *mental models* more typically deal with composite cognitions such as causal belief structures (Carley and Palmquist, 1992; Diesner and Carley, 2011; Morgan et al., 2001). However, both applications emphasize the complexities within an individual's cognitive representation and share a common methodological framework. Typically, the matrix S is formed by words directly extracted from some corpus of textual data, and the cell values are typically some type of collocation relationship (e.g., co-occurrence in a size-n moving window, in the same paragraph, sequentially in the same utterance). However, other semantic relationships such as perceived causal relations can also be encoded in the matrix (Morgan et al., 2001; Young, 1996). For computational linguists and cognitive psychologists, collocations in text are assumed to encode semantic affinities that go above and beyond syntactic and grammatical restrictions (Borge-Holthoefer and Arenas, 2010; Steyvers and Tenenbaum, 2005). Frequent collocations extracted from sliced utterances, or "free association tasks,"[1] are taken to represent semantic relatedness (Collins and Loftus, 1975; McRae and Jones, 2013); this includes, but is not limited to, comembership in a same category (e.g., *robin* and *raven* are both *bird*), concepts with similar features (e.g., the geographic contour of *China* and the physical shape of *rooster*), and thematic relations (e.g., *bread* and *butter*).

Associative models of semantic representation emphasize that meanings of words and concepts lie in patterns of word-word associations. Researchers further attempt to link network-level properties of collocation patterns with the evolution and functions of human language, as well as cognitive processes and functions at the individual level. For

example, it has been shown that large-scale human lexical semantic networks such as the WordNet and the University of South Florida English Free Association Norms display both small-world properties (high clustering coefficient and low average path length; see Newman, 2000; Watts and Strogatz, 1998) and long-tailed degree distributions (Barabási, 2009; Steyvers, 2005). One study has also shown that more creative people can be distinguished from their "cliché thinking" peers based on the *small-worldness* of free association networks (Kenett, Anaki, and Faust, 2014). Noticeably, small-worldness is a network-level property that cannot be identified by merely counting word frequencies or using conventional methods like self-reported survey; it can only be captured by analyzing the entire semantic network.

Metrics calculated on the basis of individual semantic networks offer another class of individual-level attributes that can be used to predict attitudinal and behavioral changes, which creates a point of connection with persuasion research. For example, cognitions or beliefs occupying more central positions in a mental model are likely to be more consequential for individual behavior than beliefs with similar levels of accessibility but lower centrality scores. A recent study shows that, indeed, the most central cognitions (e.g., goals and values) of farmers correlate better with their adoption of sustainable practices (Hoffman, Lubell, and Hillis, 2014). In the political domain, cognitive mapping has been used to characterize the mental models and belief structures of political leaders as revealed in their public statements; and structural features of these mental models were further found to explain policy initiatives (Kim, 2004).

Since it is usually difficult to gather data on the beliefs and opinions of political elites, their public statements become an important source of data reflecting their psychological states and processes. Crucially, in Kim's (2004) study, the mental models of two political leaders were formed by similar cognitions; it was the way in which those cognitions were structured that seemed to have made a difference in the political initiatives they advocated. Even though this paper analyzed only two mental models, one can expand this line of work by identifying a list of relevant network statistics (e.g., density, centrality, community structure) that are psychologically meaningful and analyze their association with behavioral outcomes in a larger sample of political actors.

Individual-level semantic network has also been used to reveal another dimension of political discourse: persuasiveness. For example, Doerfel and Connaughton (2009) analyzed an archive of speeches made during televised presidential debates from 1960 to 2004 and extracted co-occurrence semantic networks for every candidate. They found that a semantic network characterized by a tightly clustered group of concepts predicted election winning. In this case, semantic network analysis offers an empirical tool to test whether a discourse that is coherently structured along central themes is more persuasive than one with multiple topics. These properties are relatively difficult to unravel using frequency counts of specific words, yet are naturally foregrounded when the entire discourse is viewed as a semantic network.

In research on salience and framing, the analyst aggregates mental models across a collective of individuals but the process excludes the mapping of "who said what" from the

analysis. The construction of aggregated semantic networks offers a unique approach to content analysis—a well-established method in the study of political communication (Althaus et al., 2011; Grimmer and Stewart, 2013; Krippendorff, 2012)—and complements frequency-based techniques by emphasizing the structural properties of the corpus under analysis (van Atteveldt, 2008; Carley and Palmquist, 1992; Diesner and Carley, 2011). Researchers adopting the semantic network approach have expanded theories on public opinion formation by highlighting the importance of associations and the relational nature of belief structures. For example, research on agenda setting shows that political elites represented in the media have a persuasive impact on public attitudes, opinions, and policy support (McCombs and Shaw, 1972; Zaller, 1992). A recent study compared the semantic networks of mass media outlets with semantic networks of Obama and Romney supporters during the 2012 presidential election cycle, both constructed from a large corpus of 38 million Twitter messages (Vargo et al., 2014). The findings suggest that media shape not only the salience of issues and issue attributes in public discussions, but also how those issues are coreferenced and linked to each other (i.e., network issue agenda-setting).

The meso-level properties of semantic networks, such as their community structure, offer a different way to operationalize frames. If *frames* are defined as patterns of closely interconnected concepts, semantic networks make it possible to analyze interframe relationships such as how focal frames are associated with other elements of the political discourse (Baden, 2010). For example, for a given policy issue, multiple frames could be put forth to the public without referring to and elaborating upon each other. In a semantic network, this would result in a number of loosely connected communities with few between-community ties. Alternatively, parties could choose to structure their public statements around a few central frames and use other peripheral frames to further explicate these core frames. This would lead to a core-peripheral structure in the semantic network. Treating frames as self-contained semantic entities and analyzing frequency distribution of frames alone will miss these important structural patterns specifying how frames are connected to each other. In the case of the Dutch referendum campaign on the European Union (EU) constitution, parties structured the whole discourse around a few central frames that were often composed of dialectically opposing claims; more important, these central frames did not stand alone but were connected to a series of supportive—though more peripheral—frames (Baden, 2010). Semantic network analysis offers a flexible and systematic tool to identify and compare frames in this context-sensitive way and can reveal nuanced and interesting interframe connection patterns.

Applied to public opinion research, these ideas are useful to reveal the complexities of opinion formation and can provide new information to describe shifts in public opinion as a path-dependent process. At the individual level, semantic networks can reveal the key cognitions (e.g., beliefs and values) that are responsible for the opinions held and provide guidelines to evaluate normative implications and even possible interventions (such as when opinions are based on misconceptions). When aggregated to the group or

population level, this type of semantic network can be compared and contrasted, which can help identify groups and demographic segments that employ (dis)similar frames when discussing a political issue; population-level semantic networks (as reconstructed from, say, social media) can be further compared with the networks of political elites, and with the appropriate time resolution, can be used to test temporal models of opinion formation. We discuss these applications further below.

Networks Mapping the Association of Actors with Concepts

At the interpersonal level, discourse network analysis (Fisher, Leifeld, and Iwaki, 2012; Leifeld and Haunss, 2012) and research on *discursive fields* (Bail, 2012) explicitly include the information on how social actors endorse semantic units or concepts in their public discourse. The action of "endorsing" can be captured on a continuous scale (such as the number of times an actor uses a word) or as a binary variable, as represented in figure 13.1, panel 2a, indicating, for example, agreement or disagreement.

There are not many studies that analyze the two-mode network connecting actors with concepts—among other reasons because the range of available tools for analyzing such networks is limited. More often, researchers project the affiliation matrix A into two square matrices: the $m \times m$ matrix of actors represented as a graph in figure 13.1, panel 2c, and the $n \times n$ matrix of concepts represented as a graph in figure 13.1, panel 2d. Ties in panel 2c and 2d are weighted by the number of common concepts two actors co-endorse (panel 2c) and the number of actors a pair of concepts share (panel 2d), respectively. It is important to acknowledge that projections of two-mode networks inherently involve information loss. In projecting the affiliation matrix A into the $m \times m$ one-mode network of actors, information on the set of concepts that link actors is lost. Implicitly, the researcher is making the assumption that all the concepts are equally important and weighted in the same way when determining actor-actor relationships. This choice prevents researchers from directly incorporating nodal attributes of concepts (e.g., whether or not the concept is related to misinformation) in the analysis. Similarly, the concept-concept one-mode projection excludes potentially important information about political actors. These weaknesses can be overcome by applying metrics and statistical methods developed specifically for two-mode networks, which are beginning to appear in the literature (e.g., Agneessens and Everett, 2013); however, in this chapter we focus the discussion on one-mode projections, as they remain the main object of analysis (for an exception, see Kleinnijenhuis and de Nooy [2013]).

The most attractive feature of reconstructing semantic networks while retaining information on actor-concept endorsement is its ability to characterize and model the discourse of specific actors. Researchers can examine how affiliation relationships among actors develop in response to their public statements, which is of relevance in the study of political coalitions and polarization during policy debates. This approach also helps investigate how public discourse becomes more complex or restricted as actors interact with and respond to each other. One exemplary application is the discourse

network analysis, developed by Leifeld and colleagues (2012) to simultaneously represent the overall typology of political elites and public political expressions over time. This technique builds up the actor-concept bipartite semantic network by coding political elites' agreement or disagreement (i.e., ties are binary) with policy-related statements from textual archives (e.g., news coverage, congressional testimonials). By analyzing the concept-concept projection (in which ties represent the degree to which a concept is commonly endorsed by multiple political elites, called *concept congruence network*), the authors were able to identify rhetorical patterns of the winning side in the debate. For example, in the case of the software patents controversy in Europe, the prevailing coalition of political actors (including governments, companies, and NGOs) coherently knitted multiple concepts together, a feature missing in opponents' discourse from the other side of the controversy (Leifeld and Haunss, 2012). In the same study, the actor-actor projection whose ties encode the level of shared beliefs and positions (i.e., *actor congruence network*) empirically revealed two competing coalitions with a high degree of within-group consensus and between-group conflicts; this distinction emerged from the construction and analysis of semantic networks, that is, without imposing a grouping criterion based on subjective judgment. This discourse network analysis approach has also been used to track how consensus emerged over time on the issue of climate change in the 110th US Congress, in contrast with the bigroup structure in the 109th Congress (Fisher, Leifeld, and Iwaki, 2012). Chapter 12 in this handbook elaborates in more detail on the methodological and substantive principles behind this analytical approach.

Actor-concept semantic networks can be further complemented by incorporating other types of information. In a study on how the discourse of civil society organizations affected media coverage about Islam after the terrorist attack on September 11, 2001, researchers first positioned these organizations in the discursive field of Muslims and 911 to identify *fringe* organizations, that is, organizations occupying peripheral corners as revealed by patterns of coendorsement of Muslim-related frames (Bail, 2012). Position in the discursive field was found to significantly impact media's mirroring of frames originally adopted by these organizations. More interestingly, these fringe organizations were more likely to influence mass media coverage if their discourse displayed higher levels of emotionality. This example demonstrates that when coupled with other types of information (e.g., content attributes of the text and other characteristics of actors), semantic network analysis can significantly enrich researchers' explanatory repertoire and enable the empirical study of theories (e.g., discursive fields) otherwise difficult to operationalize. Although both the actor congruence/conflict network from discourse network analysis and Bail's (2012) study on discursive fields analyze patterns of endorsing semantic units between individual political actors, it is possible to extend this line of research to groups of actors defined by demographic status, partisanship, or other dimension that is of theoretical significance.

When applied to public opinion research, semantic network analysis is particularly suited to address questions derived from discursive and deliberative theories of opinion formation. These theories aim to analyze the generation, production, and exchange of ideas, arguments, reasons, and even sentiments—all of which can be treated as semantic

units in a network. This makes the analysis of semantic networks useful to decode how the collective discourse enriches and is enriched by individuals' cognitive representations of political issues, which is core to much public opinion research (Baek, 2011; Carpini, Cook, and Jacobs, 2004; Dijk, 1995; Gutmann and Thompson, 2004). Semantic network analysis operating at the interpersonal level have so far been applied mostly to political elites, but we foresee nothing preventing researchers from analyzing other types of political actors, especially citizens engaging in political discussion or deliberation, either offline or online.

The analysis of semantic networks differs in important ways from more conventional approaches to text, such as frequency-based content analysis (Krippendorff, 2012). First, content analysis that follows a frequency-based approach focuses exclusively on the counts and distribution of concepts in a corpus; a network approach, on the other hand, adds the additional layer of how concepts are related to each other. The frequency-based approach, in other words, extracts only a text's "fundamental building blocks" but not "the structure in which these blocks are arranged" (Carley and Palmquist, 1992, 605). Second, frequency-based approaches cannot directly incorporate the social aspect of message generation and exchange, that is, the actor-concept relationships captured by the two-mode semantic network. Semantic networks have the additional advantage of being able to represent a variety of semantic units and the ways in which those units are interconnected. This information can be extracted from textual data using NLP and text-mining techniques, as explained in the following section.

BUILDING SEMANTIC NETWORKS

How to Extract Semantic Units (Nodes) from Textual Data

Collecting raw textual data. Raw textual data to construct semantic networks can be collected from various data sources, including but not limited to news coverage, political testimony before the US Congress, social media posts on Twitter and discussion forums, experimentally generated data (e.g., offline and online deliberation studies, free association tasks), recorded and transcribed utterances in natural conversations, and open-ended survey questions. For public opinion research, semantic content generated by political actors is the primary focus.

Typically, researchers interested in user-generated messages from the Internet (e.g., comment boards on news websites, Facebook and Twitter posts, discussion forums) can gain access through scrapping websites, accessing application programming interfaces (APIs), and purchasing data from vendors. Unfortunately, to date there is no single aggregator that offers easy user interface to access the whole body of ever-growing, user-generated texts across social media sites and platforms. Moreover, both

ownership issues regarding commercial companies' claim to proprietary data and concerns over user privacy are likely to impact the scope and nature of accessible online semantic data. Which data source to use and through what means to access it need to be evaluated on a case-by-case basis. For instance, tweets posted by average citizens can provide insights into the current state of public opinion (Vargo et al., 2014); however, they are less ideal to examine how public opinion is formed via discursive interactions due to the character limit per tweet and the lack of prolonged discussions among typical Twitter users.

Once the data have been accessed, researchers often need to develop a valid list of keywords to retrieve a corpus that is relevant to the issue being investigated. Compared with news coverage, user-generated content is likely to include nonstandard language, slang terms, and misspellings. The keyword list for data retrieval is thus issue-specific and needs to be developed through trial and error. The performance of the keyword list can be assessed by the criteria of *precision* (i.e., low false positives, the fraction of true fits over the size of the corpus extracted and labeled as relevant) and *recall* (i.e., low false negatives, the proportion of true hits in the extracted corpus over the total set of matched texts, including those not retrieved or mistakenly labeled as irrelevant; see Manning, Raghavan, and Schütze, 2008). For example, assume a researcher wants to retrieve tweets posted in 2015 related to same-sex marriage. Using a test version of a keyword list, the researcher retrieves m tweets in total, out of which j tweets are validated as relevant. Then a point estimate for precision is the simple fraction j/m. To assess recall, the researcher needs to pull out a random sample of tweets of any kind posted during the same time window (say 1 percent of all tweets under Twitter's Firehose retrieval method). In this test sample, the researcher then removes tweets ($n_1 = k$) that are (1) already in the corpus retrieved via keyword matching and (2) classified as true hits. Next, the researcher identifies tweets related to same-sex marriage using a combination of human coding and machine classification ($n_2 = p$). The point estimate of the recall of this retrieving strategy using the particular keyword list will be $\dfrac{k}{k+p}$. Both precision and recall can be improved by adding or deleting keywords to/from the search term list.

Defining the semantic ontology. The first and most important decision prior to analyzing sematic networks is what units to extract from textual data. The primary goal of this step is to define the dimensions of the column space of the two essential matrices introduced in session 2 (i.e., S_n in figure 13.1, panels 1a and 2a). This decision should be guided by the nature and scope of the research. In general, though, two issues should be taken into consideration: first, deciding whether to take a deductive or an inductive approach to the identification of semantic units, that is, derive them from the data or identify them in text according to preconceived categories (Carley and Kaufer, 1993; Carley and Palmquist, 1992); and second, determining the level of abstraction of the semantic units to be analyzed.

The answers will largely determine the methodological techniques to employ, such as whether NLP tools and automatic text processing algorithms are needed. In this section

we define an issue-specific collection of texts (e.g., newspaper coverage, public discourse transcribed into texts, testimonials from politicians) as the *corpus*. For the sake of simplicity, we assume *text$_i$* is the basic unit of analysis after slicing, binning, or bracketing the corpus based on rules applied to the original documents (e.g., aggregating all the tweets collected on the same day as a single *text*). We focus our discussion primarily on applications related to individual- and collective-level concept-concept semantic networks (e.g., cognitive mapping, salience, and framing research in table 13.1), because procedures to construct political discourse networks at the interpersonal level are discussed in detail in Chapter 12 of this handbook. Methodological considerations discussed below also apply, however, to discourse network analysis, although it typically requires additional layers of complexity—that is, to identify actor-concept mappings from textual data. This additional step poses some unique challenges to the application of automated text-mining techniques, as political actors, semantic units, and pairwise mapping between the elements in these two categories are difficult to automatically extract from unstructured textual data all at once. As a result, current applications of discourse network analysis require human coding to fill in the $A_{m \times n}$ matrix in figure 13.1, panel 2a.

The total set of unique semantic units to be represented in the network creates the *semantic ontology* of a public issue. Often researchers have a predefined set of beliefs, values, themes, or positions they want to analyze. For example, Vargo et al. (2014) identified eight specific policy domains related to the 2012 presidential election judged relevant by researchers (e.g., *economy, foreign policy*). Since categories in a predefined semantic ontology typically represent higher-order concepts or themes that are more abstract than words in raw texts, the deductive approach implicitly poses a double-classification problem that involves determining (1) whether *text$_i$* matches or mismatches with each of the categories and (2) if so, in what way (e.g., merely mentioning, positively or negatively endorsing). Both problems can be solved using human coding following standard procedures of content analysis (Krippendorff, 2012); however, if the researcher intends to employ machine learning algorithms with limited human input to achieve more efficiency, the second problem will be much harder to address than the first, especially when the dimensionality of the semantic ontology and ways relating *text$_i$* to semantic units grow large. These technical difficulties are likely to be ameliorated in the future, given rapid developments in machine learning and text mining research.

The inductive approach to defining the semantic ontology is exploratory in nature. Researchers identify relevant semantic units after interacting with the raw data. For example, in the study on cognitive networks of farmers around the notion of "sustainable agriculture" and their relationship with sustainable practices, Hoffman, Lubell, and Hillis (2014) elicited unique concepts related to this issue only after inspecting collected data from an open-ended survey question. The inductive approach has the advantage of being more likely to identify novel semantic units and to portray a comprehensive picture of the totality of concepts used by the public when expressing their opinions. For these reasons, this approach is more suitable for early stages of a research project or for more descriptive and exploratory studies.

Lower-order semantic units refer to exact words or phrases used in the text after pre-processing (e.g., stemming and lemmatization; see more details below). On the other hand, higher-order semantic units lack exact one-to-one mapping to the exact words or phrases used and hence require procedures for *inference*, that is to assign meaning (i.e., the word-concept mapping) to lexical symbols used in the text (Corman et al., 2002). Most lexical semantic networks from computational linguistics, networks of association norms from cognitive psychology, and networks constructed using the "moving window" approach (Danowski, 2009; Yuan, Feng, and Danowski, 2013) choose to represent lower-order words, while political discourse networks (Leifeld, 2013; Vargo et al., 2014) tend to use higher-order semantic units. Methodologically, if the researcher is satisfied with lower-order semantic units, directly inspecting the raw text after tokenization and other preprocessing steps will reveal unique semantic units needed to define the dimensionality of the two essential matrices. However, if the researcher chooses to employ automatic textual analytical tools and analyze higher-order semantic units, the text requires more sophisticated algorithms to process, such as topic modeling, or scalable human coding using crowdsourcing platforms (Benoit et al., 2016). To reiterate, the purpose of defining the semantic ontology is to determine the column and row dimensions in matrix S (figure 13.1, panel 1a) and the columns in matrix A (figure 13.1, panel 2a).

Preprocessing of raw textual data. Once raw textual data are collected, they are usually stored as long strings of characters that need to be *tokenized* before further processing. Tokenization breaks the long string down to the smallest unit of semantic content, typically at the word level (Jurafsky and Martin, 2008). Tokenization discards the order by which words appear in the text, although using *n-grams* ($n > 1$) will partially preserve the information on word order. *N*-grams are contiguous sequences of words of length *n*. For example, bigrams for the short phrase "I love you" would be "I love" and "love you," each treated as a unique token. Discarding word order, known as the bag-of-word approach, is typically found to pose little harm to performance for common NLP tasks such as classification, sentiment analysis, and topic modeling (Grimmer and Stewart, 2013). However, for specific types of one-mode network of semantic relationships, such as causal mapping in the mental models tradition (Carley and Palmquist, 1992; Morgan et al., 2001), word orders convey information about the direction of the causal relationship. For this type of application, more sophisticated NLP techniques such as part-of-speech tagging and semantic role labeling, will be useful to tag tokens with additional attributes (e.g., roles as in agent-predate-patient structure; Jurafsky and Martin, 2008).

A recent study shows that using *n*-grams of different size might reveal different causal mechanisms that lead to similarity in statements produced by US Congressman. While addressing similar topics (i.e., topic similarity) was positively correlated with authors' patterns of using similar words and phrases only when the raw corpus was preprocessed with shorter *n*-grams ($n < 3$), working in the same chamber was positively correlated with author-author similarity in language use when longer *n*-grams were incorporated as tokens ($n > 16$) (Lin, Margolin, and Lazer, 2015). This suggests that independently constructing messages is likely to pose a limit on the length of phrases used, but directly

copying others' statements poses no such restraint. Therefore, using n-grams of different sizes can be beneficial when the researcher has specific causal mechanisms in mind to test for public opinion formation processes, such as exposure to different media outlets with independent opinion formation (shorter n-gram) versus social learning and persuasion (longer n-gram).

From a data processing perspective, human language is "noisy." The same word can have different forms that are not immediately recognized by a computer (e.g., "U.S."/ "United States"/"U.S.A"/"America"). This issue is addressed by *stemming* and *lemmatizing*, text manipulations that aim to reduce inflectional forms of a word to a common base (Jurafsky and Martin, 2008). Typically preprocessing will also remove punctuation, capitalization, and common functional words that are used to preserve grammatical integrity rather than convey specific meanings (known as "stop-words," typically including prepositions and common verbs like "be").

After preprocessing, the set of unique tokens in the corpus becomes the semantic ontology for the political issue in question (for inductive and lower-order definition of the semantic ontology). Some specific research questions may require extracting a subset of those tokens. For example, Corman et al.'s (2002) *centering resonance analysis* only pays attention to nouns and noun phrases, while discarding verbs, subjectives, and adverbs that might also encode important semantic information about the public issue being analyzed. Decisions like this should be based on the specific application of the semantic network approach and the substantive theoretical question.

Using machine learning to identify higher-order semantic units. For *higher-order* definitions of semantic ontology, the mapping between tokens and categories (defined in the preconceived set or inferred from the data) or latent semantic dimensions (for inductive approach to higher-order semantic units) must be specified. A machine-learning approach would require as input data the term-document matrix with word frequencies after pre-processing, as well as the category labels provided by human coders. For the deductive approach, mapping poses a classification problem to which supervised machine learning algorithms can be applied, such as Naïve Bayesian, Support Vector Machine, Random Forests, Neutral Networks, and the ensemble approach (Aggarwal and Zhai, 2012). However, given the inherent complexities in human language and discourse, good performance of automatic text classification algorithms is not ensured and needs to be evaluated and validated on a case-by-case basis. Current applications (see Vargo et al., 2014) are restricted to higher-order themes (e.g., climate change, terrorist attack) rather than complicated policy statements or positions (e.g., "climate change is real and anthropologenic"). For the latter type of application, human coding remains the primary method (see Fisher, Leifeld, and Iwaki, 2012), but automatic methods are rapidly being developed.

On the other hand, the inductive approach poses a latent dimension extraction problem, in which latent semantic dimensions or topics need to be inferred from the original term-document matrix. One approach is to view this as a matrix factorization problem, in which lower-order n-grams in the original term-document matrix need to be projected

into a lower-rank latent semantic space (Dumais, 2004). This is also known as *latent semantic analysis*, and the original term-document matrix can be factorized using singular value decomposition and its variants (Dumais, 2004). The underlying logic is that the latent space is a more parsimonious summarization of observed text-word co-occurrence patterns, and dimensions of this space represent the higher-order semantic themes that the researcher is attempting to uncover. Another approach assumes a probability-based generative model, in which the text is generated from a distribution of topics, which in turn can be derived from distributions of words. Popular probabilistic models such as latent dirichlet allocation (LDA; see Blei, Ng, and Jordan, 2003) is available and can help uncover latent higher-order semantic themes. After iteratively identifying the optimal solution for the topic generative model, the texts' loadings on the latent semantic dimensions can be used to define topic-by-topic relationships—that is, a semantic relationship at a higher level of abstraction than the original tokens. A potential problem with topic modeling, however, is the ambiguity in interpreting the substantive meaning of uncovered latent topics. As a result, once the semantic network is constructed, researchers might benefit from focusing on network-level properties rather than nodal structural attributes, given the node's potentially unclear substantive meaning.

How to Extract Relationships among Semantic Units (Ties) from Textual Data

The strength, sign, and directionality of ties differ in importance as the goal of research changes (Carley and Palmquist, 1992). For example, if the purpose is capturing semantic associations in general, word co-occurrence can be used to assess their strength; directionality does not matter. However, when representing causal relationships, determining the origin and destination of a directed tie is crucial to preserve causal order (Morgan et al., 2001). Ties can also be signed and map positive and negative associations, although in practice researchers tend to favor positive signs and convert semantically negative relationships into a positive number. For example, politicians' *disagreement* with a policy position can be treated as a separate semantic unit in addition to *agreement* (i.e., see Fisher, Leifeld, and Iwaki, 2012).

Ties based on co-occurrence of semantic units. In concept-concept semantic networks at the individual and collective levels (e.g., cognitive mapping/mental models, salience and framing research), *ties* typically refer to affinity in a general sense, without a concrete meaning: two concepts can be related because they belong to the same higher-order category; they indicate a causal association; or they are connected in any other way, as long as they "make sense" to message producers such as in free association experiments (Borge-Holthoefer and Arenas, 2010; Marupaka, Iyer, and Minai, 2012; McRae and Jones, 2013; Steyvers and Tenenbaum, 2005). Ties representing this general form of semantic affinity do not impose restrictions on the type of concepts to work with; they can be higher- or lower-order semantic units defined deductively or inductively. In constructing this type of tie from raw textual data, researchers typically rely

on counting the frequency of co-occurrence within specified boundaries: words within a moving window of size n (Danowski, 2009, Yuan, Feng, and Danowski, 2013), words within a sentence in transcribed oral speeches from top policymakers (Shim, Park, and Wilding, 2015), predefined issues and names of organizations within the same paragraph in news coverage (Kleinnijenhuis and de Nooy, 2013), predefined issues within the collection of Tweets aggregated by day (Vargo et al., 2014), or inductively derived concepts within the same person's response to open-ended survey questions (Hoffman, Lubell, and Hillis, 2014; Smith and Parrott, 2012). Some texts collected in their original form possess a natural boundary (e.g., Twitter's 140-character limit), while others lack such structure and require researchers to be sensitive to the context in which the text is produced. For example, for orally produced messages such as transcribed political debates, breaks could be placed between shorter passages—turn-takes would serve as the natural break in this context—while the opposite is expected for longer written text such as news coverage.

The obtained co-occurrence matrix based on raw frequencies may require normalization, especially for inductively derived lower-order semantic units, as the raw frequencies are likely to be dominated by a few commonly occurring words that span a great number of texts but convey little substantial meaning (e.g., the pronoun "you"). Normalization can mitigate this problem by comparing the obtained raw frequencies to a null model of expected co-occurrences when no systematic relations exist between pairs of semantic units (Baden, 2010; Griffiths and Steyvers, 2002). Researchers can then threshold or filter out raw frequencies against the distribution derived from the null model.

When applying semantic network analysis to cognitive psychology, the assumption is that mapping concept associations helps uncover how semantic memory is organized in people's minds (Collins and Loftus, 1975; Kenett, Anaki, and Faust, 2014). In this area of research, ties are generated in a lab experiment setting in which participants are asked to provide responses to predefined target words. Responses across participants are then aggregated to define to what extent a pair of target words triggers similar response patterns—the more similar the set of elicited words is, the higher the strength of the tie between two subjects. Rather than using raw counts of frequencies, correlations between pairs of target words are used to adjust for the general tendency for one particular target word to elicit, on average, more responses. In this way, all the target words are placed on the same multidimensional semantic space, spanned by the aggregated set of response words; their pairwise similarity in this space is then used to define tie strength (Kenett, Anaki, and Faust, 2014; McRae et al., 2005).

Some versions of these experiments do not rely on free associations; instead, they supply a full set of predefined concepts to participants and ask them to directly indicate level of perceived association between every pair of concepts (Guo, 2012). This method is more suitable for a smaller set of concepts, as n concepts will require $n(n - 1)/2$ judgments from each participant, which is very demanding for participants when n gets large. The idea of viewing semantic units as situated in a

multidimensional space defined by the term-document matrix is particularly useful to define ties between higher-order semantic units uncovered by topic modeling. Metrics such as Euclidean distance or cosine similarity can be used to scale strength of ties.

Ties beyond co-occurrence. Although ties based on co-occurrence are popular in the early stages of semantic network research, scholars are increasingly interested in capturing more specific and substantive semantic meanings, such as cause-effect relationship in research on mental models (Morgan et al., 2001), or actors' agreement/disagreement with policy statements in discourse network analysis (Leifeld and Haunss, 2012). Young (1996) offered a summary of commonly seen categories of semantic relationships, including cause-effect, if-then, entity-attribute, warrant-for, and many more; however, what relationship to extract again depends on specific research goals. For example, researchers can use open-ended interviews and confirmatory questionnaires in an iterative fashion to assess patterns of perceived causal relationships among a set of concepts related to the issue under investigation (Morgan et al., 2001).

When the level of specificity increases, it becomes more difficult to rely on automatic machine learning algorithms. Even extending simple co-occurrence to classify relationships into positive or negative ties is not trivial. Researchers are likely to need more sophisticated NLP techniques, such as named entity recognition and disambiguation, part-of-speech tagging, and syntactic analysis to denote grammatical roles of semantic units (e.g., subject-object pairs) to facilitate the classification task (Atteveldt et al., 2008; Atteveldt, Kleinnijenhuis, and Ruigrok, 2008). This is especially true for actor-concept semantic networks when both (1) the ontology of actors and semantic units and (2) actor-concept association need to be directly extracted from the same raw textual data.

When ties map the association of actors with concepts, their meaning differs depending on the projection analyzed. In the actor-to-actor network (figure 13.1, panel 2c), a tie means that two social actors endorse the same set of semantic units. In the concept-to-concept network (panel 2d), a tie means that two semantic units are endorsed by the same set of social actors. Importantly, ties constructed in this way denote quite distinctive information in contrast to ties in networks mapping semantic relationships (figure 13.1, panel 1). While the former emphasize coendorsement patterns across the population of political actors analyzed, the latter focus on either general semantic affinity or more specific semantic meanings independent of the actor-concept association patterns.

Political scientists have applied this projection method to analyze patterns of political coalition and conflicts among political decision makers and stakeholders (Fisher, Leifeld, and Iwaki, 2012; Leifeld, 2013). When some actors are overrepresented in the data, tie weights in the network need to be normalized. For example, in the actor-policy position network, some actors may contribute disproportionately to the volume of public statements, hence covering a much wider spectrum of policy positions that define the semantic ontology (this is termed "institutional bias" by Leifeld [2013]). If

no normalization is done, the analysis of the network and the resulting metrics might reflect this bias and hide meaningful relationships among actors. When original entries are binary, normalization methods such as taking the pairwise correlation coefficient are no longer valid; an alternative is using the family of Jaccard coefficients, which in its simplest form takes the fraction of the intersection between two binary vectors over their union. This coefficient also returns a value that ranges between 0 and 1. More details on this approach, and an expanded discussion of its application, can be found in Chapter 12.

APPLICATIONS IN PUBLIC OPINION RESEARCH

Classic models treat public opinion as an "emergent product" of interpersonal discussions and social relationships, free and uncensored in procedures, and supported by rational and well-informed arguments (Blumer, 1946; Lazarsfeld, 1957; Price, 1992). Later empirical research found this approach difficult to operationalize, and over time surveys and polls have become the preferred tools to represent the opinion of the public (Converse, 1987). Reflecting on the deficiencies of using surveys to represent public opinion, noticeably because of their disconnection with public discussions, scholars have recently made attempts to refine the measurement tools, including deliberative opinion polls (Fishkin, 1991, 1997). Recent theorists of deliberative democracy place a large weight on citizens' discursive interactions and activities in the political decision-making process (Carpini, Cook, and Jacobs, 2004; Dryzek and Niemeyer, 2008; Fishkin, 1997; Gutmann and Thompson, 2004; Mendelberg, 2002). This involves paying attention to responsiveness to arguments, the extensiveness of the information on which opinions are based, their consistency with values, and the overall stability of opinion—all important dimensions to evaluate the quality of public opinion. The empirical operationalization of these criteria requires a way of representing the full spectrum of arguments, reasons, beliefs, and concepts that are used during discursive interactions; the flexibility to switch between levels of analysis (i.e., from the individual to the collective and population levels); and tools to model changes in the structure of the opinion-related semantic ontology over time. The analysis of semantic networks offers a framework to empirically address these requirements and can serve as a useful tool to study public opinion in its complexity.

In the existing literature, typical applications of semantic network analysis focus on elite political discourse (Baden, 2010; Doerfel and Connaughton, 2009; Fisher, Leifeld, and Iwaki, 2012; Kleinnijenhuis and de Nooy, 2013; Leifeld, 2013; Lin, Margolin, and Lazer, 2015); very few studies have taken advantage of this methodology to study the

public's political discourse (for an exception, see Vargo et al. [2014]). Part of the reason is the difficulty of gathering data on political discussions among average citizens in their everyday lives. However, given the explosion of social media and online interactions, in which discursive political exchanges happen naturally, the problem of data access is now considerably mitigated (although issues of representativeness remain). Moreover, semantic network analysis can also handle discussion data generated and collected in more controlled settings, such as online deliberation experiments (Cappella, Price, and Nir, 2002; Price, Nir, and Cappella, 2006) and structured online town hall meetings (Minozzi et al., 2015). Following these developments, we summarize existing and potential applications of semantic networks to the study of public opinion, following the categorization scheme outlined in table 13.1.

A citizen's cognitive representation of a public issue can be mapped out as a semantic network that contains both "cognitions" and "how these cognitions are connected" (Carley and Palmquist, 1992; Johnson-Laird, 2010; Morgan et al., 2001; Popping, 2006), with the nature of ties varying from general associations to specific types like perceived causal relationships. Most public opinion research is concerned with evaluative judgments (e.g., support or oppose); mapping the full range of concepts underlying those judgments can reveal more information about individual issue positions. Carley et al. (1992, 1993) pioneered the methodology to automatically reconstruct mental semantic networks from texts generated by individuals, paving the way for empirical researchers to process large-scale textual data. One potential challenge of this application is the possible discrepancy between written text and genuine thoughts: individuals might generate messages in a biased way for impression management, as a result of social pressure, or to achieve other types of strategic goals.

Another possibility at this level of analysis is to use polls and surveys to reconstruct associations between belief systems and attitudes. This creates a more complex relational account of what underlies political opinions and helps identify types of electorate on the basis of their similarities and differences in the entire system of beliefs (Baldassarri and Goldberg, 2014). More research is needed, however, to appropriately link mental representations with actor networks.

Once a network is constructed, it can be analyzed as any other network object, and researchers can focus on different structural properties in line with the research questions posed. At the node level, a popular measure of a semantic unit's structural importance is its centrality score. Based on betweenness and degree centralities, Shim, Park, and Wilding (2015) offered a typology that classifies nodes into (1) hub for the entire network (high on both centrality measures), (2) hub for the local meaning community (high on degree but low on betweenness centrality), (3) bridges (high on betweenness but low on degree centrality), and (4) peripherals (low on both centrality measures). Interpreting and categorizing nodes in the mental semantic network provides information about *structural* properties, which can be analyzed later along with other nodal attributes the researcher might have data on already. For example, if the researcher knows which cognitions are misconceptions, he or she can identify when

misconceptions occupy the hub position in networks and when they are peripheral, then use this structural variation to explain different responses to persuasive attempts, beyond the fact that everyone might hold the same type of misconception.

At the meso level, the community structure of semantic networks can help operationalize concepts such as *schemata* or *frame* popular in political psychology (Lau, Smith, and Fiske, 1991; Scheufele and Tewksbury, 2007). Communities are parts of a network in which connections are denser internally than externally, as compared to what to expect in a random benchmark (e.g., a randomly generated network with the same degree sequence; see Girvan and Newman, 2002). By interpreting what these communities consist of and correspond to, researchers can gain insights into the number of dimensions an individual thinks about when considering a public issue. Methodologically, these community structures can be identified by employing community detection algorithms (Fortunato, 2010; see also Kolaczyk and Csárdi, 2014).

At the global level, even simple metrics such as the size of the semantic network and the overall density tell us how sophisticated an individual is when asked to reflect on an issue. Both the size of the entire semantic network and the number of distinctive communities serve as measures for the degree of cognitive differentiation, though with different resolutions. Compared with size, the number of communities is less likely to be biased by factors like individuals' literacy levels. Larger and more unique communities correspond to a better ability to consider unique dimensions of an issue (i.e., they are proxies for higher levels of cognitive differentiation; for a discussion on cognitive differentiation, see Conway et al. [2014] and Tetlock et al. [2014]). On the other hand, the density of the network measures cognitive integration (i.e., the ability to make connections among unique dimensions), often considered the hallmark for political sophistication (Conway et al., 2014; Tetlock et al., 2014).

When analyzing semantic networks formed by collective actors, such as organizations or groups of people, researchers can focus on identifying rhetorical attractors or disentangling the dynamics of coalition formation. A node with high centrality in the projected concept-to-concept network means that the concept is semantically related to the majority of other concepts also mentioned in the text; these concepts serve as rhetorical attractors around which the discourse converges. A node with high centrality in the projected actor-to-actor network, on the other hand, may play the role of consensus builder and help create a common denominator. In addition to centrality, all the other metrics assessing the local, meso, and macro structure of the network can be applied to characterize discourse and group dynamics. The difference is that at this level, the metrics will reflect properties of the aggregated "group mind" rather than individual mental models. For example, on the issue of energy policies among six countries, a study analyzing semantic networks of speeches made by top-level policymakers found that while Germany emphasizes *clean energy*, the discourse of the United States and Japan was structured by *nuclear safety* and *energy security*, as measured by betweenness centrality of nodes (Shim, Park, and Wilding, 2015). These networks, and their differences, help characterize national policy frames.

Finally, the analysis of semantic networks can zoom out to the population level and model opinion dynamics on a societal scale. Yuan, Feng, and Danowski (2013), for example, applied a modularity-maximization algorithm to identify communities of words in China's Twitter-like social media platform Sina Weibo related to the notion of *privacy*. The analyses helped identify different dimensions underlying the notion privacy, closely related to specific cultural roles. In addition to community detection algorithms, researchers have applied hierarchical clustering analysis to identify groups of words, which offers another way to operationalize frames. After the extraction of clusters of words, Baden (2010) analyzed framing in the interpretation of the EU constitution and concluded that the coherence of the discourse as a whole was determined by how well the core concepts are interconnected.

One application of these ideas attempts to explain public opinion formation from the structure of media coverage. This theoretical approach, dubbed network agenda setting (NAS), argues that the agenda set by news organizations not only concerns how often a specific issue is covered, but also the relationship between different news items (Guo, 2012; Vargo et al., 2014). It follows that overall connection patterns in the semantic network of news coverage should predict connection patterns in the aggregated public mind. Using the quadratic assignment procedure (QAP), the authors found that in the 2012 election, the semantic network of Obama supporters was predicted by vertical media's network agenda (i.e., newspapers and broadcast news networks); the network of Romney supporters, on the other hand, was better accounted for by horizontal media's network agenda (i.e., cable news networks and talk shows).

The longitudinal analysis of those networks can further enrich our understanding of public opinion formation; it provides valuable information on the process by which individuals' cognitive representations evolve as they interact with collective discourse, and it offers metrics to track polarization or convergence over time. So far, however, research using semantic networks to study public opinion formation is rare. Existing literature tends to focus on policy-related elite discourse and on how political coalitions are formed based on actors' responses to each other (Fisher, Leifeld, and Iwaki, 2012; Kleinnijenhuis and de Nooy, 2013; Leifeld, 2013). In the quest to develop this line of research across levels of analysis, researchers can take advantage of powerful techniques to model network dynamics. Although we have not yet seen many studies applying likelihood-based inferential techniques designed specifically for networks, such as the family of exponential random graph modeling (ERGM) methods (for details see Cranmer and Desmarais [2011]; Desmarais and Cranmer [2012]; Ingold and Leifeld [2014]; and Chapter 8), the lack of studies should not prevent researchers from employing such methods—as long as their application complies with the main assumptions made by those techniques. What sets semantic networks apart from social networks is that nodes are not actors having the agency to decide how to form their ties, which is a basic assumption in generative models like ERGMs. Researchers intending to apply ERGMs need to reflect carefully on whether the specific type of semantic network to be analyzed satisfies the analytical assumptions of ERGMs.

CONCLUSIONS

Research applying semantic network analysis to the study of politics in general and public opinion in particular is only in its nascent stage. In this chapter we reviewed existing literature analyzing semantic networks on the individual, interpersonal, and collective levels. Semantic networks offer a flexible representational and analytical framework that is particularly well suited to studying public opinion as it forms and evolves, potentially offering new empirical insights to discursive and deliberative theories of democracy. Studies employing this tool in public opinion research are still rare, although significant methodological advances have already been made in other areas, such as computational linguistics, cognitive psychology, political sociology, and policy studies. All of these developments have paved the way for public opinion researchers to apply semantic networks to study issues that have long been neglected due to overreliance on mass opinion polling.

Future research should consider ways to combine and triangulate data sources. For example, the structural information should be combined with other attributes available for actors (e.g., individual's ideological stance, political information consumption, general political knowledge) and semantic units (e.g., values versus factual belief, scientifically valid versus misconceptions). Second, although in this chapter we primarily focus on static semantic networks (i.e., snapshots of evolving structures), their temporal dimension encodes important information that should also be considered. For the networks mapping semantic relationships, a longitudinal view can offer a long-term perspective of public opinion formation. For the networks mapping associations of actors with concepts, the analysis of the processes that lead to differentiation, polarization, and consensus can offer empirically testable hypotheses of discursive and deliberative models of public opinion. Achieving those goals will require solving many challenges that are still open, but we foresee much exciting research rising to the task.

NOTE

1. In a free association task, subjects are asked to report the first word coming to mind that is related in a specific way (e.g., meaning, rhyme, make words) to the prompted cues. The University of South Florida Word Association Norms is the largest collected free association data set in the United States, covering more than six thousand participants and encompassing nearly three-quarters of a million responses to more than five thousand stimulus words (see Nelson, McEvoy, and Schreiber, 2004).

REFERENCES

Aggarwal, C. C., and Zhai, C. X. (2012). "A Survey of Text Classification Algorithms." In *Mining Text Data*, edited by C. Aggarwal and C. X. Zhai, pp. 163–222. Springer. http://link.springer. com/chapter/10.1007/978-1-4614-3223-4_6.

Agneessens, F., and Everett, M. G. (2013). "Introduction to the Special Issue on Advances in Two-Mode Social Networks." *Social Networks* 35(2): 145–147.

Althaus, S. L., et al. (2011). "Assumed Transmission in Political Science: A Call for Bringing Description Back In." *Journal of Politics* 73(4): 1065–1080.

Atteveldt, W. v. (2008). *Semantic Network Analysis: Techniques for Extracting, Representing, and Querying Media Content*. Charleston, SC: BookSurge.

Atteveldt, W. v., Kleinnijenhuis, J., and Ruigrok, N. (2008). "Parsing, Semantic Networks, and Political Authority Using Syntactic Analysis to Extract Semantic Relations from Dutch Newspaper Articles." *Political Analysis* 16(4): 428–446.

Atteveldt, W. v., Kleinnijenhuis, J., Ruigrok, N., and Schlobach, S. (2008). "Good News or Bad News? Conducting Sentiment Analysis on Dutch Text to Distinguish Between Positive and Negative Relations." *Journal of Information Technology & Politics* 5(1): 73–94.

Baden, C. (2010). *Communication, Contextualization, & Cognition: Patterns & Processes of Frames' Influence on People's Interpretations of the EU Constitution*. Delft: Eburon Academic Publishers.

Baek, Y. M. (2011). "The Impact of Deliberation on Social and Semantic Networks: Citizens' Mental Models of Bioethical Issues in Genetics Using Automated Textual Analysis." http://repository.upenn.edu/dissertations/AAI3475893.

Bail, C. A. (2012). "The Fringe Effect Civil Society Organizations and the Evolution of Media Discourse about Islam since the September 11th Attacks." *American Sociological Review* 77(6): 855–879.

Baldassarri, D., and Goldberg, A. (2014). "Neither Ideologues nor Agnostics: Alternative Voters' Belief System in an Age of Partisan Politics." *American Journal of Sociology* 120 (1): 45–95. doi:10.1086/676042.

Barabási, A.-L. (2009). "Scale-Free Networks: A Decade and Beyond." *Science* 325(5939): 412–413.

Baunach, D. M. (2012). "Changing Same-Sex Marriage Attitudes in America from 1988 Through 2010." *Public Opinion Quarterly* 76(2): 364–378.

Benoit, K., et al. (2016). "Crowd-Sourced Text Analysis: Reproducible and Agile Production of Political Data." *American Political Science Review: forthcoming*. http://eprints.lse.ac.uk/62242/

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3: 993–1022.

Blumer, H. (1946). "Elementary Collective Behavior." *New Outline of the Principles of Sociology*, 170–177.

Borgatti, S., Martin, E., and Johnson, J. (2013). *Analyzing Social Networks*. Los Angeles, Thousand Oaks, CA, London: Sage Publications.

Borge-Holthoefer, J., and Arenas, A. (2010). "Semantic Networks: Structure and Dynamics." *Entropy* 12(5): 1264–1302.

Cappella, J. N., Price, V., and Nir, L. (2002). "Argument Repertoire as a Reliable and Valid Measure of Opinion Quality: Electronic Dialogue During Campaign 2000." *Political Communication* 19(1): 73–93.

Carley, K., and Kaufer, D. S. (1993). "Semantic Connectivity: An Approach for Analyzing Symbols in Semantic Networks." *Communication Theory* 3(3): 183–213.

Carley, K., and Palmquist, M. (1992). "Extracting, Representing, and Analyzing Mental Models." *Social Forces* 70(3): 601–636.

Carpini, M. X. D., Cook, F. L., and Jacobs, L. R. (2004). "Public Deliberation, Discursive Participation, and Citizen Engagement: A Review of the Empirical Literature." *Annual Review of Political Science* 7(1): 315–344.

Collins, A. M., and Loftus, E. F. (1975). "A Spreading-Activation Theory of Semantic Processing." *Psychological Review* 82(6): 407.

Converse, P. E. (1987). "Changing Conceptions of Public Opinion in the Political Process." *Public Opinion Quarterly* 51: S12–24.

Conway, L. G., Conway, K. R., Gornick, L. J., and Houck, S. C. (2014). "Automated Integrative Complexity: Automated Integrative Complexity." *Political Psychology* 35(5): 603–624.

Corman, S. R., Kuhn, T., Mcphee, R. D., and Dooley, K. J. (2002). "Studying Complex Discursive Systems." *Human Communication Research* 28(2): 157–206.

Cranmer, S. J., and Desmarais, B. A. (2011). "Inferential Network Analysis with Exponential Random Graph Models." *Political Analysis* 19(1): 66–86.

Crespi, I. (1997). *The Public Opinion Process: How the People Speak*. Mahwah, NJ: Routledge.

Danowski, J. A. (2009). "Inferences from Word Networks in Messages." In *The Content Analysis Reader*, ed. Krippendorff, K. and Bock, M., 421–429. Thousand Oaks, CA: Sage publications.

Desmarais, B. A., and Cranmer, S. J. (2012). "Statistical Inference for Valued-Edge Networks: The Generalized Exponential Random Graph Model." *PLOS ONE* 7(1): e30136.

Diesner, J., and Carley, K. M. (2010). "A Methodology for Integrating Network Theory and Topic Modeling and Its Application to Innovation Diffusion." In *2010 IEEE Second International Conference on Social Computing*, pp. 687–692. http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5591514.

Diesner, J., and Carley, K. M. (2011). "Semantic Networks." In *Encyclopedia of Social Networking*, ed. Barnett, G., 766–769. Thousand Oaks, CA: Sage publications.

Dijk, T. A. v. (1995). "Discourse Semantics and Ideology." *Discourse & Society* 6(2): 243–289.

Doerfel, M. L., and Connaughton, S. L. (2009). "Semantic Networks and Competition: Election Year Winners and Losers in U.S. Televised Presidential Debates, 1960–2004." *Journal of the American Society for Information Science and Technology* 60(1): 201–218.

Dryzek, J. S., and Niemeyer, S. (2008). "Discursive Representation." *American Political Science Review* 102(4): 481–493.

Dumais, S. T. (2004). "Latent Semantic Analysis." *Annual Review of Information Science and Technology* 38(1): 188–230.

Fisher, D. R., Leifeld, P., and Iwaki, Y. (2012). "Mapping the Ideological Networks of American Climate Politics." *Climatic Change* 116(3–4): 523–545.

Fishkin, J. S. (1991). *Democracy and Deliberation: New Directions for Democratic Reform*. New Haven, CT: Yale University Press.

Fishkin, J. S. (1997). *The Voice of the People: Public Opinion and Democracy*. New Haven, CT: Yale University Press.

Fortunato, S. (2010). "Community Detection in Graphs." *Physics Reports* 486(3–5): 75–174.

Girvan, M., and Newman, M. E. J. (2002). "Community Structure in Social and Biological Networks." *Proceedings of the National Academy of Sciences* 99(12): 7821–7826.

Griffiths, T., and Steyvers, M. (2002). "A Probabilistic Approach to Semantic Representation." In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, pp. 381–386. Citeseer. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.9.5961&rep=rep1&type=pdf.

Grimmer, J., and Stewart, B. M. (2013). "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3): 1–31. doi:10.1093/pan/mps028.

Guo, L. (2012). "The Application of Social Network Analysis in Agenda Setting Research: A Methodological Exploration." *Journal of Broadcasting & Electronic Media* 56(4): 616–631.

Gutmann, A., and Thompson, D. (2004). *Why Deliberative Democracy?* Princeton, NJ: Princeton University Press.

Hoffman, M., Lubell, M., and Hillis, V. (2014). "Linking Knowledge and Action through Mental Models of Sustainable Agriculture." *Proceedings of the National Academy of Sciences* 111(36): 13016–13021.

Ingold, K., and Leifeld, P. (2014). "Structural and Institutional Determinants of Influence Reputation: A Comparison of Collaborative and Adversarial Policy Networks in Decision Making and Implementation." *Journal of Public Administration Research and Theory* 26 (1): 1–18. doi:10.1093/jopart/muuo43.

Johnson-Laird, P. N. (2010). "Mental Models and Human Reasoning." *Proceedings of the National Academy of Sciences* 107(43): 18243–18250.

Jurafsky, D., and Martin, J. H. (2008). *Speech and Language Processing.* 2d ed. Upper Saddle River, NJ: Prentice Hall.

Kenett, Y. N., Anaki, D., and Faust, M. (2014). "Investigating the Structure of Semantic Networks in Low and High Creative Persons." *Frontiers in Human Neuroscience* 8. http:// www.ncbi.nlm.nih.gov/pmc/articles/PMC4051268/.

Kim, D.-H. (2004). "Cognitive Maps of Policy Makers on Financial Crises of South Korea and Malaysia: A Comparative Study." *International Review of Public Administration* 9(2): 31–38.

Kleinnijenhuis, J., and de Nooy, W. (2013). "Adjustment of Issue Positions Based on Network Strategies in an Election Campaign: A Two-Mode Network Autoregression Model with Cross-Nested Random Effects." *Social Networks* 35(2): 168–177.

Kolaczyk, E. D., and Csárdi, G. (2014). *65 Statistical Analysis of Network Data with R.* New York: Springer New York. http://link.springer.com/10.1007/978-1-4939-0983-4.

Krippendorff, K. H. (2012). *Content Analysis: An Introduction to Its Methodology.* 3d ed. Los Angeles, London: Sage Publications.

Lau, R. R., Smith, R. A. and Fiske, S. T. (1991). "Political Beliefs, Policy Interpretations, and Political Persuasion." *Journal of Politics* 53(3): 646–675.

Lazarsfeld, P. F. (1957). "Public Opinion and the Classical Tradition." *Public Opinion Quarterly* 21(1): 39–53.

Leifeld, P. (2013). "Reconceptualizing Major Policy Change in the Advocacy Coalition Framework: A Discourse Network Analysis of German Pension Politics." *Policy Studies Journal* 41(1): 169–198.

Leifeld, P., and Haunss, S. (2012). "Political Discourse Networks and the Conflict over Software Patents in Europe." *European Journal of Political Research* 51(3): 382–409.

Lin, Y.-R., Margolin, D., and Lazer, D. (2015). "Uncovering Social Semantics from Textual Traces: A Theory-Driven Approach and Evidence from Public Statements of U.S. Members of Congress." *Journal of the Association for Information Science and Technology* June 2. doi:10.1002/asi.23540.

Liptak, A. (2015). "Supreme Court Ruling Makes Same-Sex Marriage a Right Nationwide." *New York Times,* June 26. http://www.nytimes.com/2015/06/27/us/supreme-court-same-sex-marriage.html.

Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to Information Retrieval.* New York: Cambridge University Press.

Marupaka, N., Iyer, L. R., and Minai, A. A. (2012). "Connectivity and Thought: The Influence of Semantic Network Structure in a Neurodynamical Model of Thinking." *Neural Networks* 32: 147–158.

McCombs, M. E., and Shaw, D. L. (1972). "The Agenda-Setting Function of Mass Media." *Public Opinion Quarterly* 36(2): 176–187.

McRae, K., Cree, G. S., Seidenberg, M. S., and Mcnorgan, C. (2005). "Semantic Feature Production Norms for a Large Set of Living and Nonliving Things." *Behavior Research Methods* 37(4): 547–559.

McRae, K., and Jones, M. N. (2013). "Semantic Memory." In *The Oxford Handbook of Cognitive Psychology*, pp. 206–219. Oxford University Press.

Mendelberg, T. (2002). "The Deliberative Citizen: Theory and Evidence." *Political Decision Making, Deliberation and Participation* 6(1): 151–193.

Miller, M. M. (1997). "Frame Mapping and Analysis of News Coverage of Contentious Issues." *Social Science Computer Review* 15(4): 367–378.

Minozzi, W., Neblo, M. A., Esterling, K. M., and Lazer, D. M. J. (2015). "Field Experiment Evidence of Substantive, Attributional, and Behavioral Persuasion by Members of Congress in Online Town Halls." *Proceedings of the National Academy of Sciences* 112(13): 3937–3942.

Morgan, M. G., Fischhoff, B., Bostrom, A., and Atman, C. J. (2001). *Risk Communication: A Mental Models Approach*. Cambridge, UK, and New York: Cambridge University Press.

Nelson, D. L., McEvoy, C. L., and Schreiber, T. A. (2004). "The University of South Florida Free Association, Rhyme, and Word Fragment Norms." *Behavior Research Methods, Instruments & Computers* 36(3): 402–407.

Newman, M., Barabási, A.-L., and Watts, D. J. (2006). *The Structure and Dynamics of Networks*. Princeton, NJ: Princeton University Press.

Newman, M. E. J. (2000). "Models of the Small World." *Journal of Statistical Physics* 101(3–4): 819–841.

Pew Research Center. (2015). "*Changing Attitudes on Gay Marriage*." http://www.pewforum. org/2015/06/08/graphics-slideshow-changing-attitudes-on-gay-marriage/.

Popping, R. (2006). "Text Analysis for Knowledge Graphs." *Quality & Quantity* 41(5): 691–709.

Price, V. (1992). *4 Public Opinion*. Newbury Park, CA: Sage Publications.

Price, V., and Neijens, P. (1997). "Opinion Quality in Public Opinion Research." *International Journal of Public Opinion Research* 9(4): 336–360.

Price, V., Nir, L., and Cappella, J. N. (2006). "Normative and Informational Influences in Online Political Discussions." *Communication Theory* 16(1): 47–74.

Scheufele, D. A., and Tewksbury, D. (2007). "Framing, Agenda Setting, and Priming: The Evolution of Three Media Effects Models." *Journal of Communication* 57(1): 9–20.

Scott, J. (2012). *Social Network Analysis*. 3d ed. Los Angeles: Sage Publications.

Shim, J., Park, C., and Wilding, M. (2015). "Identifying Policy Frames through Semantic Network Analysis: An Examination of Nuclear Energy Policy across Six Countries." *Policy Sciences* 48(1): 51–83.

Smith, R. A., and Parrott, R. L. (2012). "Mental Representations of HPV in Appalachia: Gender, Semantic Network Analysis, and Knowledge Gaps." *Journal of Health Psychology* 17(6): 917–928.

Steyvers, M., and Tenenbaum, J. B. (2005). "The Large-Scale Structure of Semantic Networks: Statistical Analyses and a Model of Semantic Growth." *Cognitive Science* 29(1): 41–78.

Tetlock, P. E., Metz, S. E., Scott, S. E., and Suedfeld, P. (2014). "Integrative Complexity Coding Raises Integratively Complex Issues: Complexities of Integrative Complexity Coding." *Political Psychology* 35(5): 625–634.

Vargo, C. J., Guo, L., McCombs, M., and Shaw, D. L. (2014). "Network Issue Agendas on Twitter During the 2012 U.S. Presidential Election: Network Issue Agendas on Twitter." *Journal of Communication* 64(2): 296–316.

Wasserman, S., and Faust, K. (1994). *Social Network Analysis: Methods and Applications.* Cambridge, UK, and New York: Cambridge University Press.

Watts, D. J., and Strogatz, S. H. (1998). "Collective Dynamics of 'Small-World' Networks." *Nature* 393(6684): 440–442.

Young, M. D. (1996). "Cognitive Mapping Meets Semantic Networks." *Journal of Conflict Resolution* 40(3): 395–414.

Yuan, E. J., Feng, M., and Danowski, J. A. (2013). "'Privacy' in Semantic Networks on Chinese Social Media: The Case of Sina Weibo: Privacy in Semantic Networks on Sina Weibo." *Journal of Communication* 63(6): 1011–1031.

Zaller, J R. (1992). *The Nature and Origins of Mass Opinion.* Cambridge, UK, and New York: Cambridge University Press.