Semantic Networks

Semantic networks are structured representations of knowledge that are used for reasoning and inference. A large variety of theories, models, methods, and practical applications for creating and utilizing semantic networks have emerged from various branches of academia and industry. Back in 1975, Woods stated that semantic networks are an attractive notion, but lack theoretical grounding and rigor in representational conventions. Since then, technical advances, especially the advent of the web, have lead to an exponential increase in the availability of data that need to be converted into knowledge and appropriately managed to allow for versatile utilization. Today, many of the approaches to semantic networks still have little in common beyond calling their object of study a semantic network. This entry reviews the main application domains for semantic networks, and describes the corresponding methods for creating semantic networks that can serve these purposes.

Representation of Semantic Networks

Semantic networks require three constituent parts:

- A syntax that specifies the types of nodes and edges that can be considered.
- Specification of the meaning or semantics that the nodes, links, and entire network can represent.
- Inference rules.

Data structures representing semantic networks comprise, at a minimum, nodes that are referred to as concepts, and edges between the concepts. Concepts are abstract representations of the ideas, thoughts, and units of knowledge and meaning that people conceive in their minds. Concepts can also be referents of objects. Examples for concepts are “social network analysis”, “collaboration” and “community of practice”. When concepts have a representation in natural language, the respective word or sequence of words are used as labels for nodes in semantic networks. If links in semantic networks are typed, the link type explicates the nature of the relationship between the connected nodes. Otherwise, links represent or establish some meaningful association between pairs of concepts. Link formation in semantic networks heavily depends on the type of input data and intended use of the network as described later in this entry. Beyond these constraints, representations of semantic networks can be anywhere between strictly formalized to informal.

Across many different approaches, the smallest structured unit of a semantic network is typically a triplet that describes the source or subject, predicate, and object or target of an action. This basic structure is used, for instance, to describe who did what to whom for representing event data, and who said what to whom for communication data. Such triplets can further be enhanced with background information, such as spatial and temporal data, and attributes of nodes and links.

Semantic networks result from a translation or transformation of the data that the networks were constructed from. Translations convert natural language input into isomorphic, structured representations that are used as input to inference mechanisms. This approach has been pursued in Artificial Intelligence as one attempt to understand natural language, and has shown to be usable for small amounts of input that. Transformations are abstraction processes that are used to preserve and reveal the entities and relations that are explicitly or implicitly represented in the input data. The goal with transformations is the reduction of the dimensionality of the input data in order to capture the relevant structural interrelations and to make network data available for inference.

Utilization and corresponding methods for creating semantic networks

Semantic networks are used for the individual and collective acquisition, organization, management and utilization of knowledge. Knowledge is gained from semantic networks by performing reasoning and inference on the network data. The knowledge stored and inferred from semantic networks does not have to be factually correct or logical. This is due to two reasons: First, semantic networks were designed to represent what is meant by a piece of information. This meaning can differ from the truth conditional content or the most likely interpretation of some information. Second, depending on the data from which the network was constructed, a semantic network can represent universal, culturally dependent, domain-specific, and individual knowledge.

Meaning Extraction from Semantic Networks

Given a semantic network, the meaning of a concept is the network that is activated when the node of interest is triggered. An early formalization of this notion is the theory of spreading activation suggested by Collins and Loftus in 1975. This theory was designed to model how humans process and memorize information in their brains. Other application domains for the same methodology for meaning extraction are the identification of the meaning intended by the author of some information and the recipient's interpretation of the information. Retrieving the network of concepts that are associated with a node offers an extension to the content analysis methodology: in content analysis, the relative prevalence of information is captured by associating text terms with concepts, and comparing the cumulative frequency of concepts. Content analysis assumes words and concepts to be conditionally independent from each other. Thus, differences in prominence and meaning that are not due to the identity and frequency of individual concepts, but that are due to the embedding of concepts in networks, cannot be identified. These differences can be considered by building a semantic network of the concepts. However, working with structured representations of knowledge is only useful when the respective research questions or tasks at hand can be expressed in form of structured variables and abstract concepts as opposed to atomic words. If these conditions are met, information contextualized via associations between concepts can be further used as input to various types of computational analyses as suggested by Griffiths, Steyvers, and Tenenbaum in 2007:

- Prediction: semantic network data can be used to forecast the set of concepts that will be evoked when a certain note is activated.
- Disambiguation: clustering the network activated when prompting a word can result in groups of concepts that each represent a different aspect of the meaning of a concept.
- Summarization: retrieving a concept's ego-network can serve as a technique for distilling the essence of some input data in a concise and structured form.

Semantic Networks for Knowledge Representation

Semantic networks are often used to represent the knowledge that a person, group, or mankind has about topics, domains or the world at a given time. An early attempt to provide a structured and comprehensive representation of the things that exist on earth is a tree created by the Greek philosopher Porphyry of Tyre, who lived 234–305 A.D.. This tree is an example of declarative or definitional semantic networks, which basically are collections of the knowledge about something. Typically, declarative networks are domain specific, but can be used for a large variety purposes. Many declarative networks are based on ontologies. Ontologies specify the set of elements and relations between elements that are possible or permissible in a given domain. For example, when creating structured representations of event data, the ontology might accept concepts that represent the who, what, when, where, why and how of an event. Mappings from data points, typically words, to such concept classes can be pre-specified in dictionaries or thesauri, or can be identified from the data. The relations are often derived from:

- The syntax or grammar of the underlying data, such as the aforementioned triplets of subjects, predicates, and objects.
- Relevant relations between concepts as perceived by the analyst, which is for example used in Grounded Theory methodology.

Ontologies can furthermore be hierarchically structured, and might allow for inheritance of features from parent nodes to child nodes. An example for a hierarchical ontology with inheritance are phylogenetic trees, also known as evolutionary trees. The relations between the elements of an ontology can furthermore be:

- Structural, such as references from a table of content or index term to a chunk of text, or pointers between webpages.
- Logical, such as equivalence relationships (“is a”), and subtype relationships (“is part of”). The domain of Artificial Intelligence has provided logics that can be used to model certain relationships between concepts.

Overall, the issue with ontologies is that they can be obsolete, incomplete and do not allow for deviations such as variations in spelling and nuances in meaning. An alternative and more flexible approach are formalized description languages that people can use to annotate data objects with description. An example for this approach is the Semantic Web. Semantic Web data is generated by people who use the Resource Description Framework language to mark up words, concepts, relations and other data objects on the web. The goal with this collective effort is to produce machine-readable definitions of data that can be automatically processed and used for inference and reasoning. Another approach to overcome the rigidness of ontologies is to drop the exclusivity criterion that allows for only one definition per concept. Free and open collaboration infrastructures such the web facilitate the collection of large numbers of definitions for concepts that obviously prompt different associations in different people, such as “happiness” or “friendship”. This approach also enables the collection of context of concepts that are seemingly synonymous, but might have subtle differences in their meaning, such as “pleased” and “delighted”. Such fine-grained distinctions in meaning are relevant for concept disambiguation tasks and for social computing applications, such as opinion mining and sentiment analysis. Finally, the latter approach crowd-sources the dictionary generation process, contributes to the expressiveness of semantic networks, and allows for computing probability distributions over concept definitions.

Once semantic networks are constructed and stored in a relational database, they can be used for various data management task, such as search and retrieval, structured queries, question answering and definition acquisition. Semantic networks can also be visualized. Since most network visualization techniques have no syntax, the resulting images are appropriate as heuristic tools and for in-depth, qualitative analysis of networks of moderate size.

Mind mapping and concept mapping are methods that facilitate individual and collective learning processes, brainstorming and the organization of knowledge. These methods are performed in a manual and computer-supported fashion as follows: starting from a central governing concept, people arrange all pertinent information according to their understanding of a concept. This understanding can for instance aim to resemble a lecture or emerge from a group discussion. The representational framework for such networks is highly informal in order to motivate creativity and stimulate the acquisition of a body of knowledge. In the resulting networks, nodes and links can be typed or not, and can be enhanced with attributes or not.

Hypertext is a dated, computer-supported method for personal text and data management. Respective tools, which used to be open-source and extensible, enable users to annotate chunks of text data and other electronically available information, such as multi-media content, with codes and user-defined links among and between data elements and codes. By constructing such structured information, users make implicit relations in the data explicit. The goal with this approach was to move from unrelated words, concepts and pieces of data to integrated interpretations. User studies on hypertext showed that people find it hard to predefine backbone structures of the knowledge they might want to consider in the future, and that they find it cumbersome to update or modified obsolete network structures. Hypertext is one instance of various qualitative text coding methods that are typically used to elicit implicit relations between concepts from texts data, explore the data in-depth, and build structural models of concepts, which can be used to formulate hypotheses that are subject to further testing.

Semantic Networks for Modeling the Mind

Semantic networks are also used as proxies for cognitive models and mental models. When used to this end, semantic networks aim to represent:

- The knowledge of individuals and groups about a certain topic. The knowledge of groups is also referred to as social knowledge, and can represent culture.
- The perceptions that people have in their minds and use to make sense of their environment.

Common methods for constructing and extracting mental models are pile sorting and analysis of natural language text data such as interviews. Once created, individual mental models are examined, compared and combined with network analytical techniques in order to identify the meaning and evolution of the knowledge of individual and groups. This technique has been used, for instance, to assess the situational awareness and shared understanding of events and tasks in remote work teams whose members use collaboration technologies to coordinate their decisions and actions. One caveat with this approach is that the relationship between natural language, mental processes and network representations is insufficiently understood.

Semantic Networks for Bayesian Inference

Semantic networks constructed by subject matter experts to model a certain phenomenon or domain are also used for Bayesian inference. This application domain of semantic networks combines structured data with statistical learning. In such networks, the concepts are random variables, and the edges indicate conditional probabilities, uncertainties and causalities. The resulting networks are also referred to as belief networks, influence diagram, procedural semantic networks and probabilistic graphical model. Regardless of the name, these models are used by assuming that the latent structure that the semantic network represents has generated the observed data through a probabilistic process. Given the observed data, the latent structure is then inferred through Bayesian inference. For example, one could assume that the words that a person expresses (random variables) were generated by certain beliefs and emotions that this person has. These beliefs and emotions are not directly observable (hidden variables), but can be inferred from the text data by applying Bayesian statistics.

References and further readings:

Berners-Lee, Tim, Hendler, James & Lassila, Ora (2001). The Semantic Web. Scientific American, 284 (5), 34-43.

Collins, Allan M., & Loftus, Elizabeth F. (1975). A spreading-activation theory of semantic processing. Psychological Review, 82 (6), 407-428.

Carley, Kathleen M., & Palmquist, Mike (1991). Extracting, Representing, and Analyzing Mental Models. Social Forces, 70 (3), 601 - 636.

Doerfel, Marya L. (1998). What Constitutes Semantic Network Analysis? A Comparison of Research and Methodologies. Connections, 21 (2), 16-26.

Griffiths, Thomas L., Steyvers, Mark, & Tenenbaum, Joshua B. (2007). Topics in semantic representation. Psychological Review, 114 (2), 211-244.

Howard, Ronald A. (1989). Knowledge maps. Management Science, 35 (8), 903-922.

Mohr, John W. (1998). Measuring Meaning Structures. Annual Reviews in Sociology, 24 (1), 345-370.

Sowa, John F. (1992). Semantic Networks. In Stuart. C. Shapiro (Ed.), Encyclopedia of Artificial Intelligence (2nd ed.), pp. 1493 - 1511. New York, NY, USA: Wiley and Sons.

Woods, William A. (1975). What's in a link: Foundations for semantic networks. In Daniel G. Bobrow & Allen Collins (Eds.), Representation and Understanding: Studies in Cognitive Science (pp. 35-82). New York, NY: Academic Press.

SEE ALSO: Communication Networks, Knowledge Networks, Word Networks, Methods of Data Collection

Jana Diesner, Carnegie Mellon University

Kathleen M. Carley, Carnegie Mellon University