Received: 27 March 2024 | Revised: 6 June 2024 | Accepted: 8 July 2024

DOI: 10.1111/psj.12556

RESEARCH NOTE

PSO PSJ POLICY STUDIES JOURNAL

Syntactic measurement of governance networks from textual data, with application to water management plans

Elise Zufall[1] | **Tyler A. Scott**[1,2]

[1]Center for Environmental Policy and Behavior, University of California, Davis, California, USA

[2]Department of Environmental Science and Policy, University of California, Davis, California, USA

Correspondence
Tyler A. Scott, Center for Environmental Policy and Behavior, University of California, Davis, CA, USA.
Email: tascott@ucdavis.edu

Funding information
National Institute of Food and Agriculture, Grant/Award Number: 2021-68012-35914; National Science Foundation, Grant/Award Number: 2205239

Abstract

This paper demonstrates an automated workflow for extracting network data from policy documents. We use natural language processing tools, part-of-speech tagging, and syntactic dependency parsing, to represent relationships between real-world entities based on how they are described in text. Using a corpus of regional groundwater management plans, we demonstrate unique graph motifs created through parsing syntactic relationships and how document-level syntax can be aggregated to develop large-scale graphs. This approach complements and extends existing methods in public management and governance research by (1) expanding the feasible geographic and temporal scope of data collection and (2) allowing for customized representations of governance systems to fit different research applications, particularly by creating graphs with many different node and edge types. We conclude by reflecting on the challenges, limitations, and future directions of automated, text-based methods for governance research.

KEYWORDS

automation, governance networks, groundwater management, network analysis, NLP

INTRODUCTION

Measurement remains a major constraint on governance network scholarship. This paper demonstrates an automated workflow for producing high-resolution network data from unstructured text. We use part-of-speech (PoS) tagging and syntactic dependency parsing to extract relational events from text and build a structural representation of a governance system. In our demonstration, we represent relationships between organizations based on how they are described in relation to one another in text. But this method supports including any type of entity in the network graph, such as projects, people, regulations, or geographic features. In social-ecological network studies, entities might be forest plots, fishing grounds, or ecological system components (Sayles et al., 2019). For other policy contexts, entity types could include neighborhoods, physical infrastructure, or financial instruments.

This syntactic measurement approach generates more detailed data than is typically generated via survey because it records arbitrarily complex typologies of nodes and relationships (e.g., type, frequency, direction, negation). Public policy and management researchers can then customize this representation to address specific research questions. We demonstrate example representations supported by this method. We conclude by discussing the challenges and limitations of text-based network measurement, as well as how this approach complements and extends existing methods in public management and governance research.

MEASURING GOVERNANCE NETWORKS

Network components

In this article, we use Torfing's general definition of *governance networks*[1] as "networks of interdependent actors that contribute to the production of public governance" (Torfing, 2012).[2] Networks have four basic building blocks: *nodes* (entities like people or organizations), *edges* (things that connect nodes), *flows* (things that move along edges), and *protocols* (rules that shape node entry, edge formation, and flow behavior) (Galloway & Thacker, 2007). A simple representation of a groundwater governance network, for instance, might comprise (1) nodes such as water agencies and water users; (2) edges between these nodes defined by connections, such as co-authorship of a management plan, water purchase arrangements, or a regulator/regulate relationship; (3) flows of information, funds, and yes, water, that transit edges; and (4) protocols set forth by signed agreements, regulatory requirements, and legal responsibilities that shape, for instance, what information will be provided, how much water can be used, or where the money must come from.

Traditional measurement approaches

Generally, there are two ways in which networks have been measured in governance and public management research. First, recall-based measurement, typically conducted via survey or interview, solicits information from members of a sample population. Recall-based approaches generate subjective measures, in that each respondent formulates their own response. The formulation process is not fully observed nor is it replicable. Second, observational measurement using *found data* extracts network information from data originally created for another purpose. Many governance scholars have used automated information extraction to analyze networks based on co-mentions in newspaper articles, co-membership in groups and partnerships, hyperlinks between organizational webpages, and social media interactions (Fried et al., 2022; Hayes & Scott, 2018; Hileman & Lubell, 2018; Sayles et al., 2022; Yi & Scholz, 2016). Observational measurement produces objective measures that are fully observed and replicable.[3]

Surveys tend to compress the complexity of a governance network into a simplified format (Berardo et al., 2020). For instance, if you ask someone to name their most frequent and trusted information sources, the elicited responses reveal an edge (an information conduit) and (ideally) a little bit about flow (*good* information *frequently* transits this edge). This type of information compression is in many respects quite useful. Asking someone who their closest friends are is likely a better way to measure friendship than observing phone calls or emails, as respondents can make a holistic assessment of a complex relational concept like friendship.[4] Of course, surveys also pose numerous challenges. The way survey instruments elicit context- and domain-specific compressed information makes replication across cases difficult. More generally, longitudinal data collection requires a great deal of time and effort. It is hard to get people to answer surveys, and survey instruments are limited by recall limitations and response burden. In contrast, policy and planning documents provide a longitudinal record capable of matching the often slow timeframe of policy change and on-the-ground impacts.

Prior research using automated approaches to analyze found data have typically considered a single network component in isolation. For instance, hyperlinks between organizational webpages can be used as a measure of *edges* (Hayes & Scott, 2018; McNutt, 2008; Sayles et al., 2022; Yi & Scholz, 2016). Similarly, measurement approaches that code ties based on co-mentions of organizations in media (e.g., Arnold et al., 2017) in essence measures only *nodes* present in the network. Ties, and their meaning, are inferred based on nodes occurring in the same place. Simple found-data approaches are shown to reasonably approximate survey-elicited networks (Hayes & Scott, 2018), but they nonetheless tend to pose a tradeoff between scale and information value. For instance, webpage hyperlink connections can be collected at a massive scale, but in the absence of supporting context these data say little about what the connection does or why it exists.

Measuring governance networks with language

Many governance processes produce written outputs, such as budgets, rules, or plans. Various administrative protocols likewise require that public entities produce meeting summaries, solicit and respond to public comments, and prepare policy analyses, such as cost–benefit analysis or environmental impact assessment. The prose of these documents contains information about the components of a governance network—who is involved, who is connected to whom, and what happens with these connections.

The method we demonstrate in this paper uses a pretrained natural language processing (NLP) model to automatically process and code network data from documents. NLP refers broadly to methods that allow computers to recognize and understand language. Automated processing of publicly available documents retains the speed and scale advantages of using found data for network measurement compared to using surveys. Analyzing the syntax and grammatical structure of these texts provides a way to turn written content into more nuanced measures of network structure than does a typical found data approach.[5]

Our method focuses on relational events. A relational event is a "discrete event generated by a social actor and directed toward one or more targets" (Butts, 2008, p. 159). In the case of a groundwater management plan, for instance, a sentence such as "Olcese GSA representatives will continue to actively participate in coordination meetings with the other Kern Subbasin GSAs" would be used to code a (future) coordination event between the subject (Olcese GSA) and objects (the other Kern Subbasin GSAs). Because events can be any type of occurrence, involving any type of entity, this open design framework offers a wealth of opportunities for multi-modal and multiplex designs, such as those pioneered in studies of social–ecological networks (Barnes et al., 2019; Bodin & Tengö, 2012; Sayles et al., 2019). In network parlance, *modes* are types of nodes, such as people, organizations, or even ecological features; *multiplex* refers to accommodating multiple types of connections between nodes.

DEMONSTRATION CASE: GROUNDWATER GOVERNANCE PLATFORMS IN CALIFORNIA

We demonstrate this method using groundwater management plans created in the state of California under the Sustainable Groundwater Management Act (SGMA), adopted in 2014. SGMA set forth state-level requirements and support for local control of groundwater management under newly formed entities called Groundwater Sustainability Agencies (GSAs). SGMA gives considerable discretion and control to local actors, who work together to fulfill the goals and requirements of the law. GSAs are meant to be at least nominally collaborative, in the sense that the process was designed to allow different actors to come together and form a GSA or participate in the subsequent development of a Groundwater Sustainability Plan (GSP). We analyzed 117 plans (148, 126 pages) ranging in length from just over 100 pages to more than 5000 pages. In specifying who is involved, what their responsibilities are, and how relevant parties will interact, each plan presents data that can be aggregated to represent local management networks.

METHODS AND MATERIALS

In what follows, we describe the series of steps involved in generating network graphs from text. We focus on methodological steps, *not* on technical implementation. In our demonstration, we perform all steps using *R* and the authors' open-source R package *textNet* (Zufall & Scott, 2024b)—but these tasks can be carried out using any object-oriented programming language and NLP engine. A fuller description of how each process is implemented, as well as all code and materials necessary for replicating the case demonstration, are available at the authors' public *github* repository (Zufall & Scott, 2024a).

Text preprocessing

As with most policy document collections, California GSPs are published in PDF format. Thus, the first step is simply to use pdf-to-text conversion software and render each document in plain text. This process is not lossless—PDFs do not contain formatting tags like html files, for instance[6]—but it is straightforward. For the remaining preprocessing steps, we use the pretrained, open-source *spaCy* NLP model (Honnibal & Montani, 2017). First, plan texts are *segmented* into sentences and tokens. Tokens are generally words but can also be a punctuation mark or symbol; each token serves a unique function in the sentence and has certain linguistic properties, such as PoS or tense. *Part-of-speech tagging* is then used to label each word token by its grammatical function, such as noun, verb, or preposition. Next, a standardized version of each word token is created that removes grammatical modifications of the same root word. This process is called *lemmatization*. Finally, *dependency parsing* takes the tagged grammatical structure and identifies the role each word token plays in a sentence (e.g., subject versus object). Figure 1 visualizes the result of this process on a sample sentence, showing tokens classified by POS and identified dependencies[7].

Named entity recognition

Entities are consistent references in text to real-world objects like people or organizations. Statistical named entity recognition (NER) is performed by the pretrained NLP model: The model identifies entities based on text features (e.g., capitalization), grammatical roles, and pretrained labels. Named entities must be *grounded*—that is, establishing the real-world object to which a named entity identified in the text corresponds. This issue is particularly relevant for policy documents wherein an agency name might be used in full and abbreviated in different places, as both names refer to the same entity.

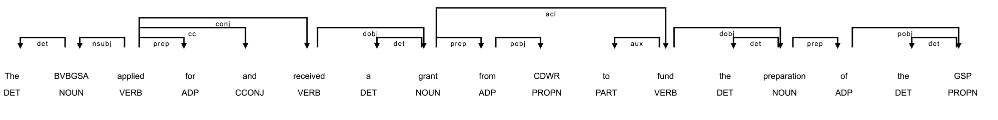

FIGURE 1 Example of a sentence that is tokenized, tagged for part-of-speech, and parsed for grammatical dependencies.

For instance, "Bear Valley Basin Groundwater Sustainability Agency" or "BVBGSA" refers to the same agency. We use rule-based classification for grounding and linking entities based on known abbreviations of real-world entities such as groundwater management agencies, local governments, and geographic features in the state of California.

Event extraction

In NLP, an event is a relationship involving an entity. The event can be punctuated—that is, something that happened at a particular moment in time—or refer to an extended relationship (e.g., "The County Public Works Department oversees flood control services."). The sentence parsed in Figure 1 contains two relational events involving organizations: a funding application from the BVBGSA to CDWR and a grant provided by CDWR to BVBGSA (Figure 2).

In governance research, the most common approach to text-based network measurement has been to code relational events based on co-occurrence in the same piece of text. Because co-occurrence-based measurement does not use dependency parsing, it does not distinguish between the subject and object or identify the verb(s) connecting the subject and object. This means that co-occurrence graphs are undirected and do not contain information about the relationship in question. Relational event coding based on syntactic relationships establishes the direction and nature of the event. This enables more detailed analysis but also fundamentally changes graph structure. Because co-occurrence-based coding does not distinguish between subject and object, undirected ties are coded between every co-occurring entity. For instance, co-occurrence-based coding of the following sentence "The SVBGSA's membership includes the County of Monterey, Monterey County Water Resources Agency (MCWRA), City of Salinas, City of Soledad, City of Gonzales, City of King (King City), the Castroville Community Services District (CCSD), and Monterey One Water" produces a fully dense network motif with 28 ties connecting every pair of entities (Figure 3a). In contrast, syntax-based event coding establishes seven ties, one each between the SVBGSA and its constituent members (Figure 3b).[8]

The prior example sentence is grammatically simple. Parsing this syntax into edge relationships is straightforward. Policy documents, however, are not known for simple writing. To produce as much (accurate) data as possible, we developed a process for handling nested events and other sentence complications. For instance, in the sentence "The Coordinating Committee will meet and discuss the Technical Memoranda," the Coordinating Committee ought to be the source for both the verb "meet" and the verb "discuss." We address these situations by tracing each sentences' dependency tree and applying a series of decision rules for parsing and coding. The end result is that for different verbs in a sentence, a separate edge is created for every combination of source entity and target entity that are connected to that verb. Finally, potential source and target entities are filtered to leave only those designated as named entities by the NER process.

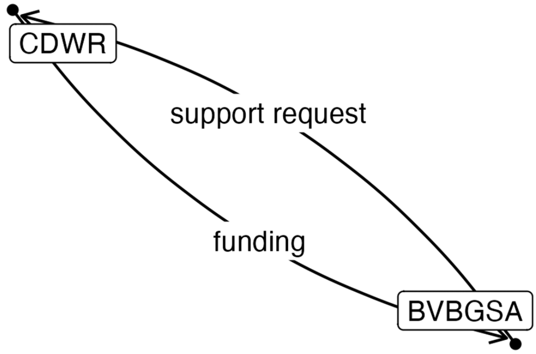

FIGURE 2 Example of relational event structure extracted from sentence parsed in Figure 1.

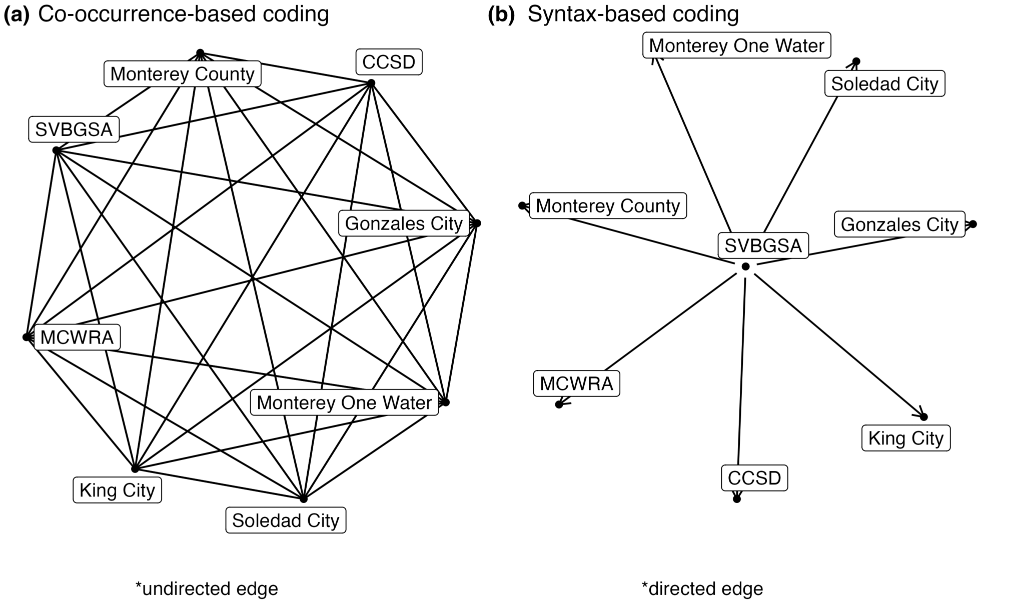

FIGURE 3 Comparison of graph structure that results from co-occurrence-based coding (a) versus event-based coding (b).

Event filtering

We include edge attributes for each verb edge based on additional properties of the verb phrase, such as verb tense, the presence of auxiliary verbs, and the presence of hedging words. Verb tense is important to distinguish past policy actions from present or ongoing policy actions and future plans. Many policy documents combine statements about network structures that currently exist and structures which will exist under certain conditions or if a particular policy alternative is adopted. The English language does not have a future verb form, so in Anglophone contexts, the identification of future actions requires inference based on additional tokens that create verb conjugation (e.g., "will fund…" or "going to implement"). *Hedging detection* was used to identify cases where events are qualified or otherwise softened to denote uncertainty. If one of these words was found in the sentence, such as "believed" in the phrase "believed to be," the edge was marked as containing a hedge. In many cases, policy documents will also spell out negative relations—what an entity will not do. Using the results of the PoS tagging and dependency parsing, we perform *negation detection* using a set of simple rules that key on the presence of negative terms used as adverbs (e.g., "will not share data") in a sentence. In our demonstration below, edges that contain a negation are removed from the data set, but researchers studying network conflict (e.g., McLaughlin et al., 2022) can keep both positive and negative edges.

Finally, some entities are mentioned in a document but not in relation to other entities of interest. This occurs, for instance, when an entity is a subject of a sentence that has an object that is not a named entity. These observations do not add edges to the social network, since there is no relational event between entities. But subject-verb relationships can be used to create nodal attributes based on the number and type of recorded actions. For instance, an entity described as "leading" could be coded as having a leadership role, or the count of times each entity is described as taking any action could be used as a relative measure of involvement or activity.

Assembly

The extracted events for each plan are compiled into an edgelist, with each row representing a unique combination of subject node, verb phrase, and object node. This list is then merged with a database of verb attributes for the respective primary verb, including tense, lemma (canonical verb form), and VerbNet classification. A nodelist is also generated, consisting of all entities in the plan that are organizations or people, regardless of their presence in the edgelist. Nodes in these resulting lists are then processed and collapsed through disambiguation of abbreviations and acronyms. These two lists are then combined to create a network with edges linked to nodes by node identifiers.

The resulting network is a semantic network. A semantic network is a graphical representation of how entities are related. There are many different subtypes of semantic networks, depending on the type(s) of entities being studied (e.g., objects versus concepts) and the relationships of interest (e.g., social relationships versus causal linkages) (Sowa, 2014). The type of semantic relationships syntactic parsing is most suited to identify are those that involve an action, since actions are the edges connecting nodes in our network model (Storey, 1993). Syntactic parsing thus embeds semantic relationships between entities, such as member collection relationships ("A participates in B"), action–recipient relationships ("project receives funding"), and agent–action relationships ("agency monitors progress"). Depending on the use case, researchers might choose to leave extracted relational events as-is (e.g., when using meeting minutes to code participation actions) or perform filtering and aggregation to develop theorized representations.

DEMONSTRATION

The methodology presented in this paper is meant to be a proof of concept for the use of NLP to extract governance network information from policy documents. Our primary motivation for developing an approach for text-based measurement of governance networks is to enable large-scale studies that can better answer questions about how and why different network structures lead to different policy outcomes. Text-based governance network coding can record arbitrarily complex representations of node forms, edge types, and flow states. Coupled with advances in generalized ERGMs that support valued edges (Krivitsky, 2016), multiplex relationships, and multilayer graphs (Krivitsky, 2022), this means that researchers' imaginations can be less constrained by data limitations and package functionality.

The graph database we produced contains many different types of nodes and edges. We choose to keep people and organizations, including federal agencies, state agencies, and local Groundwater Sustainability Agencies, in our nodelist for this demonstration, but other researchers may also wish to include node types such as locations, geopolitical entities, dates and times, or subject-specific node categories. In principle, any list of entities can be preserved in the nodelist, such as water bodies, chemical contaminants, species, diseases, or references to legal code, depending on the type of network relationships the researcher wishes to analyze. In the example figures below, we color-code the type of node to visualize where different types of nodes appear in the network. Federal actors are peripheral to the system, whereas state agencies are more highly connected and have more neighboring nodes. The network shown in Figure 4a, which represents a plan designed by multiple GSAs, shows GSAs as more peripheral to the system compared to 4b, a plan designed by a single GSA. We also demonstrate the tagging of edges based on their edge attributes; the verb tense most commonly found in sentences linking two nodes is colored as past, present, or future. This enables inspection of which parts of the system are active at different points in the development of the plan. For instance, federal actors in the plots below always have past-tense or present-tense verbs most commonly linking them to their neighbors. This indicates that federal agencies within this system do not play a large role in carrying out future plan actions.

Using PoS tagging and dependency parsing also unlocks a wide array of opportunities for building multilayer and multiplex network graphs that differentiate edges and flows based on verbiage and

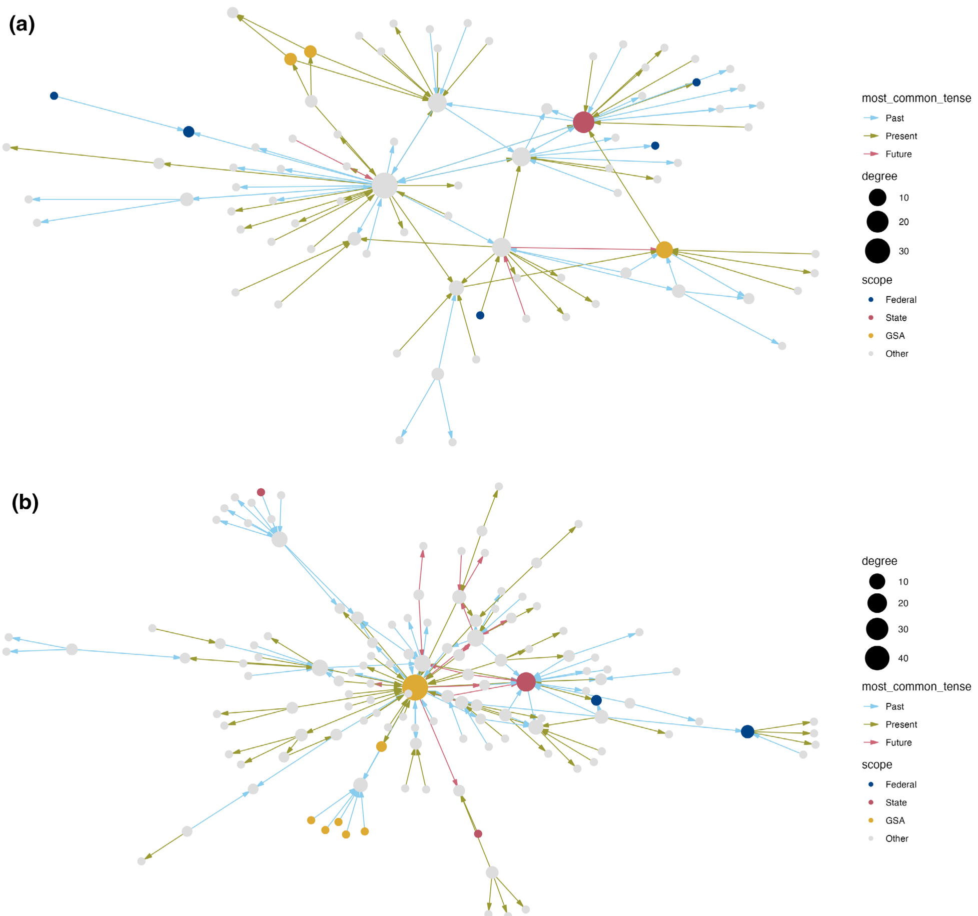

FIGURE 4 (a) The largest component of the entity network of a plan jointly developed by multiple Groundwater Sustainability Agencies (GSAs). (b) The largest component of the entity network of a plan developed by a single GSA.

syntax. In the next example, we start with the entity network of a GSP, then keep only the subgraph including the verbs "submit" and "supply." This allows us to visualize the flow of objects such as reports and documents throughout the network. In Figure 5, the local GSA responsible for creating the plan "submits" four items to the state Department of Water Resources and one item to the local advisory committee. Meanwhile, Tule Subbasin supplies one item to the advisory committee and one item to the Department of Water Resources.

In a similar fashion, we can follow the flow of resources and authority by exploring a subgraph of a plan containing the verbs "grant" and "submit" and their adjacent edges. In this motif, the largest connected component of one of the GSP subnetworks containing the verbs "grant" and "submit," we see an example of two departments, one at the state level and another at the federal level, granting resources or authority to the local groundwater sustainability agency. In return, the agency is expected to "submit" two items in return to the state department, but not to its federal counterpart (Figure 6).

In the prior examples, we showed specific local motifs. However, the text-based governance network extraction methodology can also be used to scale up across documents into a comprehensive representation. For instance, the graphs extracted from each GSP can be joined to develop an aggregate representation of groundwater governance across California. The primary means of merging different graphs is to merge common nodes. For instance, every plan mentions the California Department of Water Resources, so it is

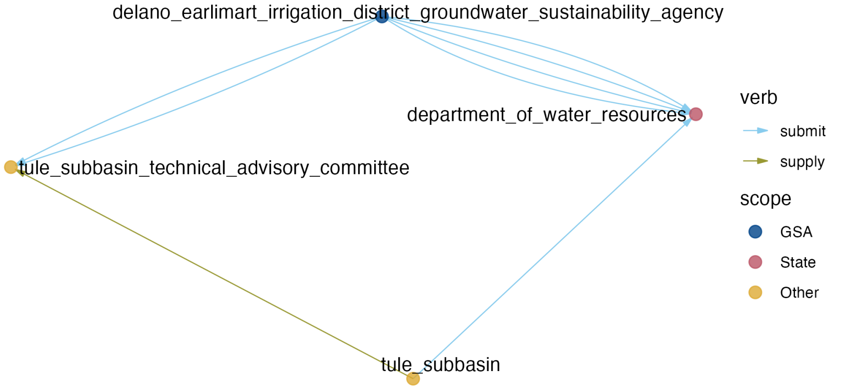

FIGURE 5 Report/document supply and submission information flow motif.

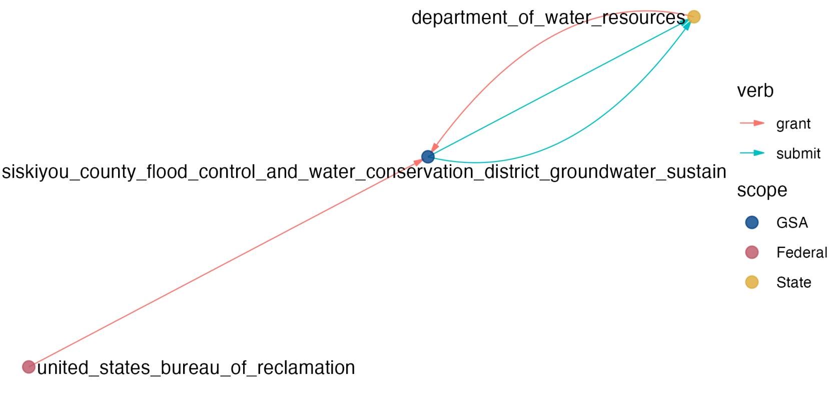

FIGURE 6 Local/state/federal authority and reporting motif.

easy enough to join all edges incident on the agency. However, most policy documents also use location specific pronouns like "the County" (to refer to a county government) or in this case "the GSA" to refer to the plan authoring agency. In case-specific analysis, these entities do have to be fully disambiguated.[9] For GSAs, full disambiguation is relatively easy, since we know which GSA authored which plan. But more generally, the frequent use of entity names that are concrete within a document but ambiguous across documents means that the analyst must figure out how to appropriately resolve interdocument ambiguity on a domain-specific basis.

In the example plot below, we first combine the networks resulting from all 117 GSP networks into a single network. Next, we filter the network to keep only federal agencies, state agencies, and local GSAs, to observe the structure of connections between these entities. Most of the GSAs have only one connection to other agencies in the plot: the California Department of Water Resources. Most state agencies and GSAs are not directly connected to one another. This type of graph gives insight into the different stages of plan development, and who is involved in each stage. For instance, the federal nodes appear to be more commonly referred to with past tense verbs. Statistical analysis to establish this is beyond the scope of the current paper but is a visual indication that in the groundwater governance system in California, federal actors are involved in the early stages of setting up the plan context, not as actively engaged in plan activities or engaged in future tasks (Figure 7).

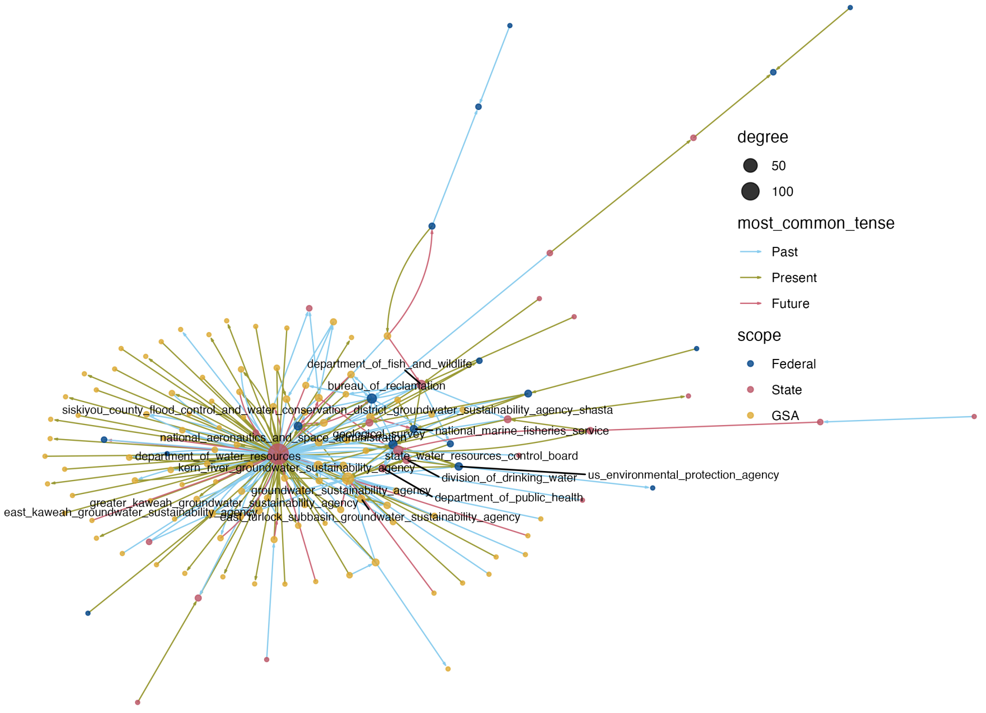

FIGURE 7 Aggregated network; all nodes with a degree of over five are labeled with their name.

REFLECTION

How and why network structures change policies and affect policy outcomes are fundamental questions for governance network scholarship. Answering these questions requires data that support cross-case comparison and offer sufficient longitudinal scale to detect change. Text generated by longitudinal policy processes like planning, rulemaking, permitting, and grantmaking is a source of network data that can overcome these research design limitations. The ability to create multiplex and multilevel networks via syntactic parsing also presents opportunities to better test and extend policy network theory. For instance, the Risk Hypothesis (Berardo & Scholz, 2010) is arguably the dominant theory explaining how and why policy actors form collaborative ties. However, policy actors are involved in multiple types of relationships that have differing risks. Establishing the *types* of relationships for which actors form bridging and bonding ties, and how those layers correlate with one another, can demonstrate the influence of risk in a way that is not possible when networks are compressed to encompass low- and high-risk activities.

Using policy texts to measure governance networks also presents its own methodological concerns that merit discussion. The general idea of automated extraction of governance network structures from the text of policy documents is fairly simple. But our experience is also that wielding automated methods in this fashion requires a "human-in-the-loop." Considerable manual effort is required to filter out the noise and produce meaningful data.[10] For policy researchers, it is also important to recognize that semantic networks are an interactive product not just of the content of the document(s) but also of the underlying data-generating process—who wrote the document and for what purpose. Measuring networks using text mining requires consideration of how the purpose and authorship of text shapes what is actually being measured and how that serves to test research questions and theory. In using how entities are described in relation to one another in text to create a semantic network graph, one assumes that this knowledge graph is a valid representation of the "real" governance network. By doing so, we are not measuring network structure directly but rather inferring structure indirectly from document content. This means that the purpose and provenance of a given document, and corpus, shapes the content of the document and thus how a given network is represented. Most importantly, to our knowledge, no type of policy document (study, report, plan, review, comment, legislation, contract, etc.) has the express purpose of *describing a governance network*. Many documents, such as the GSPs we analyze above, are intended at least in part to describe affected parties and how a system or resource will be managed.

For similar reasons, this approach is less useful for some lines of theoretical inquiry—like studying adversarial networks. For instance, the GSPs we analyze describe collective monitoring, reporting, and decision-making intended by GSA members. GSP do not explicitly document actors in conflict. One advantage of the semantic network approach is that actor/organizational entities can also be linked to policy ideas and preferences as stated in text, and therefore, adversarial network structures might be inferred from connections to different competing preferences. Nonetheless, the purpose of GSP creation generates content that more closely aligns with measuring collaborative structures.

Further, policy documents have a performative element. One purpose of a GSP, for instance, is to signal to the state regulators that the local entity is a capable manager and state intervention is not needed (SGMA has such provisions for basins that fail to develop a viable plan). Divergence between stated and actual processes is not an issue limited to the language of governance. But to the extent that this approach is meant to measure empirical networks, other data sources from the same cases might be needed to unpack how much document content is performative versus representative of the state of affairs on the ground (or measure divergence between policy *intentions* and what happens in practice).

For constructivist theories of public policy and management such as Social Construction (Ingram et al., 2007) or the Narrative Policy Framework (NPF) (Jones et al., 2014; Schlaufer et al., 2022), the subjective and performative elements of policy document authorship represent an opportunity. For instance, existing NPF research holds that policy actors cast themselves as protagonists when describing a policy environment (e.g., Gottlieb et al., 2018; Merry, 2015). Measuring the same underlying network, but described by different actors or the same actors over time, presents a way to quantify and model social construction and the relational stories that different actors tell.

Nonetheless, we believe this methodology presents a pathway for transforming network governance scholarship by enabling research designs that have long been effectively out of reach—in particular, large-scale comparative studies and multilevel aggregation. Ongoing digitalization of government (Mergel et al., 2021), and the emergence (and emergent properties) of LLMs (Wei et al., 2022) mean that we will have ever more data and better tools for implementing this approach moving forward.

ACKNOWLEDGMENTS

This work was supported by the Sustainable Agricultural Systems program, project award no. 2021-68012-35914, from the U.S. Department of Agriculture's National Institute of Food and Agriculture and the National Science Foundation's Dynamics of Integrated Socio-Environmental Systems program, under grant no. 2205239. The authors wish to thank participants at the 2023 Methodological Pluralism in Public Management symposium held at Indiana University Bloomington for their feedback on an early version of this manuscript.

ORCID

Elise Zufall https://orcid.org/0000-0003-4179-1163
Tyler A. Scott https://orcid.org/0000-0002-4214-5895

ENDNOTES

[1]This loose definition encompasses actors from all sectors and many different types of network relationships. The terms *policy network* and *collaborative network* are also used frequently in the public management literature (Ingold & Leifeld, 2016;

Isett et al., 2011; Raab, 2002; Rethemeyer & Hatmaker, 2008). While these terms are not equivalent (Kapucu et al., 2017; Rethemeyer & Hatmaker, 2008), our approach works for each.

[2]The term *governance* is itself defined in many ways. In this paper, our use of the term most closely aligns with Fukuyama's third definition of governance: "the regulation of social behavior through networks and other nonhierarchical mechanisms" (Fukuyama, 2016, p. 89).

[3]We use the terms "subjective" and "objective" here, and elsewhere in the paper, in an instrumental sense to refer to observability and replicability of measurement—not in the normative sense of being wholly unbiased or value-free. Moreover, while measurement of found data is observable and replicable, the production of found data is generally not. In the Discussion section below, we consider how the purpose of the data or the interests of the producer influences content and what the implications might be for text-based measurement of governance networks.

[4]Note that this sort of framing blurs the line between relational events and semantic relationships. Often, surveys use event extraction to establish semantic relationships. For instance, a survey might ask "Who have been your most important collaborators this past year?" and then treat reported events (collaboration in the past year) as semantic relationships (who is a collaborator with whom).

[5]Note that this approach is similar in practice to the Institutional Grammar Tool (IGT) (Lien 2020) since both rely on syntactic measurement. Whereas the IGT has a predefined set of objects and labels for determining inclusion and classification, our approach is open: It does not apply a fixed set of categories but rather provides a structural framework for coding network data from text while enabling researchers to define and filter entity and verb categories based upon the specific research task and the domain and purpose of the text itself. Syntactic network measurement could be directed specifically at measuring institutional grammar by focusing on specific entity types and verbs described within a specific piece of legislation or legal agreement. But more broadly, this is a general framework for measuring *any* relational phenomenon described in written or spoken words.

[6]To improve conversion, we remove the first and last few rows of each page set apart by at least one empty line, since these often mark headers and footers that contain repetitive information that would artificially inflate certain nodes in the network. We also remove pages that have an unreasonably large number of characters per page, likely due to multiple layers of text. This often indicates the presence of a map, table, figure, or other nonprose page, which should not be treated as a series of sentences.

[7]This figure follows the naming conventions used in spaCy, which we utilize for our preprocessing protocol. The dependency labels shown in this figure in lowercase letters are defined according to the ClearNLP system. A complete, alphabetical list of each label and its meaning can be found at https://github.com/clir/clearnlp-guidelines/blob/master/md/specifications/dependency_labels.md. The part-of-speech (POS) tags shown in Figure 1 in all capital letters follow the Universal Dependencies v2 POS category definitions. A complete, alphabetical list of each abbreviation and its meaning can be found at https://universaldependencies.org/u/pos/.

[8]Our demonstration below uses the approach shown in panel B (Figure 3)—but these data can readily be converted into the structure shown in panel A simply by taking the cross-product of the underlying matrix.

[9]Also, removing these semi-ambiguous entities is undesirable because we should expect that they are used with highest frequency for the most important entities in a text. That is, a plan is most likely to use this form of reference when the entity in question is obvious to the reader (e.g., "the County" in reference to the county in which the GSA is located).

[10]Our experience thus far of incorporating large language models (LLMs) (Wei et al., 2022) into this workflow indicates that LLMs can support improvements in many respects, but not—at least in the foreseeable future—eliminate the need for customization and support. In other words, if you ask ChatGPT to read a text and produce a network graph showing all the collaborative relationships between organizations described in the text, the result will likely be far from research-ready. In fact, our approach to integrating LLMs going forward is to focus on addressing specific subtasks, such as disambiguating pronouns based on entities referenced nearby in the text, rather than prompt engineering for a simple input/output model.

REFERENCES

Arnold, G., L. A. Nguyen Long, and M. Gottlieb. 2017. "Social Networks and Policy Entrepreneurship: How Relationships Shape Municipal Decision Making about High-Volume Hydraulic Fracturing." *Policy Studies Journal* 45(3): 414–41.

Barnes, M. L., Ö. Bodin, T. R. McClanahan, J. N. Kittinger, A. S. Hoey, O. G. Gaoue, and N. A. J. Graham. 2019. "Social-Ecological Alignment and Ecological Conditions in Coral Reefs." *Nature Communications* 10(1): 2039.

Berardo, R., M. Fischer, and M. Hamilton. 2020. "Collaborative Governance and the Challenges of Network-Based Research." *The American Review of Public Administration* 50(8): 898–913.

Berardo, R., and J. T. Scholz. 2010. "Self-Organizing Policy Networks: Risk, Partner Selection, and Cooperation in Estuaries." *American Journal of Political Science* 54(3): 632–49.

Bodin, Ö., and M. Tengö. 2012. "Disentangling Intangible Social–Ecological Systems." *Global Environmental Change: Human and Policy Dimensions* 22(2): 430–9.

Butts, C. T. 2008. "4. A Relational Event Framework for Social Action." *Sociological Methodology* 38(1): 155–200.

Fried, H., M. Hamilton, and R. Berardo. 2022. "Theorizing Multilevel Closure Structures Guiding Forum Participation." *Journal of Public Administration Research and Theory* 33(4): 633–46.

Fukuyama, F. 2016. "Governance: What Do we Know, and how Do we Know it?" *Annual Review of Political Science* 19: 89–105.

Galloway, A. R., and E. Thacker. 2007. *The Exploit: A Theory of Networks*. Minneapolis, MN: University of Minnesota Press.

Gottlieb, M., E. Bertone Oehninger, and G. Arnold. 2018. "'No Fracking Way' Vs. 'Drill Baby Drill': A Restructuring of Who Is Pitted against Whom in the Narrative Policy Framework." *Policy Studies Journal* 46(4): 798–827.

Hayes, A. L., and T. A. Scott. 2018. "Multiplex Network Analysis for Complex Governance Systems Using Surveys and Online Behavior." *Policy Studies Journal: The Journal of the Policy Studies Organization* 46(2): 327–53.

Hileman, J., and M. Lubell. 2018. "The Network Structure of Multilevel Water Resources Governance in Central America." *Ecology and Society* 23(2): 48. https://doi.org/10.5751/ES-10282-230248.

Honnibal, M., and I. Montani. 2017. *spaCy 2: Natural Language Understanding with BLoom Embeddings, Convolutional Neural Networks and Incremental Parsing*. https://spacy.io/

Ingold, K., and P. Leifeld. 2016. "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.

Ingram, H., A. L. Schneider, and P. DeLeon. 2007. "Social Construction and Policy Design." In *Theories of the Policy Process*, edited by P. A. Sabatier, 169–89. Boulder, CO: Westview Press.

Isett, K. R., I. A. Mergel, K. LeRoux, P. A. Mischen, and R. K. Rethemeyer. 2011. "Networks in Public Administration Scholarship: Understanding where we Are and where we Need to Go." *Journal of Public Administration Research and Theory* 21(s1): i157–i173.

Jones, M., E. Shanahan, and M. McBeth. 2014. *The Science of Stories: Applications of the Narrative Policy Framework in Public Policy Analysis*. New York, NY: Palgrave Macmillan.

Kapucu, N., Q. Hu, and S. Khosa. 2017. "The State of Network Research in Public Administration." *Administration and Society* 49(8): 1087–1120.

Krivitsky, P. N.. 2016. *Ergm.Count: Fit, Simulate and Diagnose Exponential-Family Models for Networks with Count Edges*. The Statnet Project (http://www.statnet.org). http://CRAN.R-project.org/package=ergm.count

Krivitsky, P. N.. 2022. *Ergm.Multi: Fit, Simulate and Diagnose Exponential-Family Models for Multiple or Multilayer Networks*. The Statnet Project (https://statnet.org). https://CRAN.R-project.org/package=ergm.multi

McLaughlin, D. M., J. M. Mewhirter, and M. Lubell. 2022. "Conflict Contagion: How Interdependence Shapes Patterns of Conflict and Cooperation in Polycentric Systems." *Journal of Public Administration Research and Theory* 32(3): 543–60.

McNutt, K. M. 2008. "Policy and Politics on the Web: Virtual Policy Networks and Climate Change." *Canadian Political Science Review* 2(1): 1–15 accessed 3 August 2022.

Mergel, I., S. Ganapati, and A. B. Whitford. 2021. "Agile: A New Way of Governing." *Public Administration Review* 81(1): 161–5.

Merry, M. K. 2015. "Constructing Policy Narratives in 140 Characters or Less: The Case of Gun Policy Organizations." *Policy Studies Journal* 44: 373–95. https://doi.org/10.1111/psj.12142.

Raab, J. 2002. "Where Do Policy Networks Come from?" *Journal of Public Administration Research and Theory* 12(4): 581–622.

Rethemeyer, R. K., and D. M. Hatmaker. 2008. "Network Management Reconsidered: An Inquiry into Management of Network Structures in Public Sector Service Provision." *Journal of Public Administration Research and Theory* 18(4): 617–46.

Sayles, J. S., M. Mancilla Garcia, M. Hamilton, S. M. Alexander, J. A. Baggio, A. P. Fischer, K. Ingold, G. R. Meredith, and J. Pittman. 2019. "Social-Ecological Network Analysis for Sustainability Sciences: A Systematic Review and Innovative Research Agenda for the Future." *Environmental Research Letters* 14(9): 1–18 accessed 11 November 2019.

Sayles, J. S., R. P. Furey, and M. R. ten Brink. 2022. "How Deep to Dig: Effects of Web-Scraping Search Depth on Hyperlink Network Analysis of Environmental Stewardship Organizations." *Applied Network Science* 7(1): 36.

Schlaufer, C., J. Kuenzler, M. D. Jones, and E. A. Shanahan. 2022. "The Narrative Policy Framework: A Traveler's Guide to Policy Stories." *Politische Vierteljahresschrift* 63(2): 249–73.

Sowa, J. F. 2014. "Semantic Networks." In *Encyclopedia of Social Network Analysis and Mining*, edited by R. Alhajj and J. Rokne, 1654. New York, NY: Springer New York.

Storey, V. C. 1993. "Understanding Semantic Relationships." *The VLDB Journal: Very Large Data Bases: A Publication of the VLDB Endowment* 2(4): 455–88.

Torfing, J. 2012. "Governance Networks." In *The Oxford Handbook of Governance*, edited by D. Levi-Faur. Oxford, UK: Oxford University Press.

Wei, J., Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, et al. 2022. Emergent Abilities of Large Language Models *arXiv* [*cs.CL*]. arXiv. http://arxiv.org/abs/2206.07682

Yi, H., and J. T. Scholz. 2016. "Policy Networks in Complex Governance Subsystems: Observing and Comparing Hyperlink, Media, and Partnership Networks." *Policy Studies Journal* 44(3): 248–79.

Zufall, E. J., and T. A. Scott. 2024a. "Kings Repository: Network Methods." GitHub, June 14, 2024. https://github.com/ucd-cepb/kings/tree/main/Text_Gov_Network_Methods_Paper/Code

Zufall, E. J., and T. A. Scott. 2024b. "textNet." R, January 2024. https://github.com/ucd-cepb/textnet

AUTHOR BIOGRAPHY

Elise Zufall is a PhD candidate in the Geography Graduate Group at the University of California, Davis. Zufall studies the role of actor relationships, beliefs, and knowledge in environmental resource management and decision-making. She received her M.S. and B.S. in Earth Systems from Stanford University.

Tyler A. Scott is an Associate Professor of Environmental Science & Policy at the University of California, Davis. He co-directs the UC Davis Center for Environmental Policy & Behavior.

How to cite this article: Zufall, Elise and Tyler A. Scott 2024. "Syntactic Measurement of Governance Networks From Textual Data, With Application to Water Management Plans." *Policy Studies Journal* 00(0): 1–14. https://doi.org/10.1111/psj.12556.