PolicyFlow: Interpreting Policy Diffusion in Context

YONGSU AHN and YU-RU LIN, University of Pittsburgh

Stability in social, technical, and financial systems, as well as the capacity of organizations to work across borders, requires consistency in public policy across jurisdictions. The diffusion of laws and regulations across political boundaries can reduce the tension that arises between innovation and consistency. Policy diffusion has been a topic of focus across the social sciences for several decades, but due to limitations of data and computational capacity, researchers have not taken a comprehensive and data-intensive look at the aggregate, cross-policy patterns of diffusion. This work combines visual analytics and text and network analyses to help understand how policies, as represented in digitized text, spread across states. As a result, our approach can quickly guide analysts to progressively gain insights into policy adoption data. We evaluate the effectiveness of our system via case studies with a real-world policy dataset and qualitative interviews with domain experts.

CCS Concepts: • **Human-centered computing → Visual Analytics**; • **Interactive systems and tools**

Additional Key Words and Phrases: Cascades, diffusion networks, network inference, policy diffusion, visual analytics

ACM Reference format:
Yongsu Ahn and Yu-Ru Lin. 2020. PolicyFlow: Interpreting Policy Diffusion in Context. *ACM Trans. Interact. Intell. Syst.* 10, 2, Article 13 (June 2020), 23 pages.
https://doi.org/10.1145/3385729

1 INTRODUCTION

Understanding the transmission of ideas, information, and resources among individuals and organizations has been a central theme in many fields, such as collective actions [18, 34], international cooperation [23], and economic development [6]. In the public policy domain, as the political actors (citizens, governments, or countries) repeatedly face similar political circumstances and uncertainty, the making and deployment of policies often involves a *learning* or *diffusion* process where political actors look to each other when making policy choices. Identifying such policy diffusion pathways is crucial for developing innovative yet consistent policies to address new societal challenges. However, as the diffusion process involves dynamic connections among political actors, observing such connections is difficult. In this work, we present a visual analytics system that enables the discovery of persistent policy diffusion patterns.

This work was funded in part by grants from the PICSO Lab, including NSF grants #1637067 and #1739413, the DARPA UGB, and AFOSR awards. Any opinions, findings, and conclusions or recommendations expressed in this material do not necessarily reflect the views of the funding sources.
The reviewing of this article was managed by associate editor Shixia Liu.
Authors' addresses: Y. Ahn and Y.-R. Lin, University of Pittsburgh, Pittsburgh, PA; emails: {yongsu.ahn, yurulin}@pitt.edu.

2160-6455/2020/06-ART13 $15.00
https://doi.org/10.1145/3385729

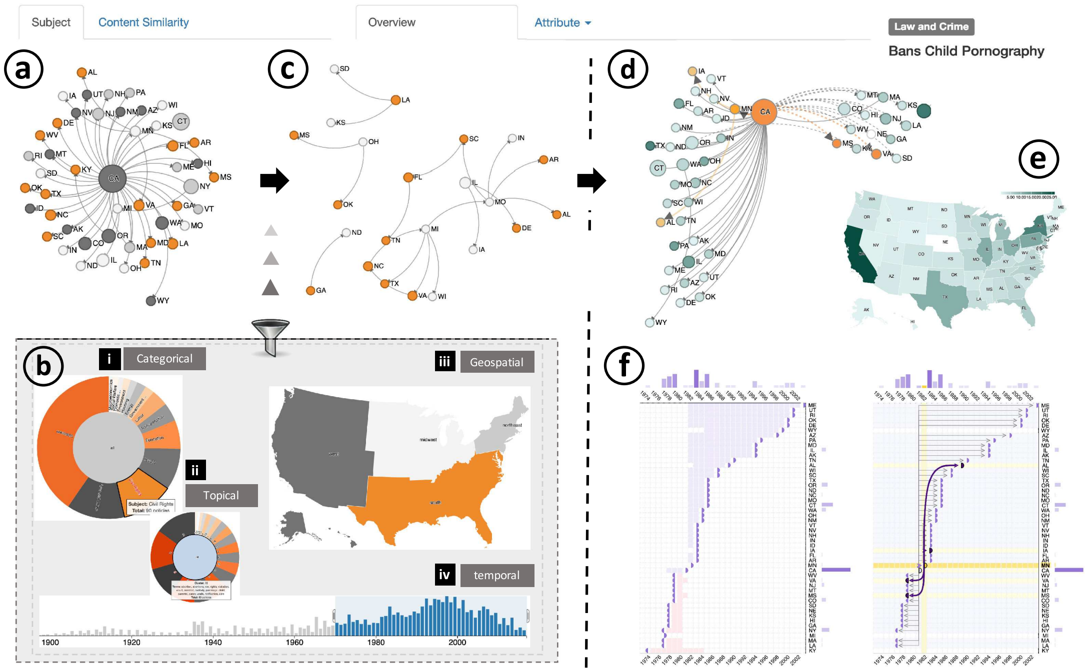

Fig. 1. PolicyFlow is a visual interactive system for exploring the time-evolving patterns of policy adoption. (a) Network-View visualizes the underlying policy diffusion network of states inferred from the trajectory of policy adoptions. Each node within the network represents a state, with the color indicating the corresponding region (Western, Southern, Mideastern, or Midwestern). (b) The system supports filtering the dataset by (i) categorical, (ii) topical, (iii) geo-spatial, and (iv) temporal dimensions. (c) This results in a smaller diffusion network corresponding to the specific political context (Midwestern (white) and Southern states (orange) appear in the network). (d) When a user selects a specific policy, a policy-specific network is rendered with the edges indicating the discrepancy between the general diffusion network and the adoption pattern of the selected policy. (CA is hovered in the current view, with similar states also colored as orange). (e) The socio-economic attributes of states (e.g., total population) helps reveal how such factors correlate to the adoption patterns (state nodes in (d) and (e) are colored green with respect to total population). (f) Policy-Inspection-View shows the adoption sequence of the selected policy and reveals the likely influence of policies from/to other states.

There is abundant political science literature that analyzes *policy diffusion*, in which the policies adopted in a given place and time are repeatedly influenced by prior policy choices made elsewhere. For example, the pioneering work by Walker [38] analyzed and theorized how policies spread from the pioneering states to the rest of the American states. Berry and Baybeck [6] incorporated geographic information systems (GIS) to analyze economic diffusion between contiguous states. Most of these works are limited to studying a single policy. Recently, Desmarais et al. [15] combined the machine learning algorithm and accumulation of policy data to characterize persistent policy diffusion patterns. They used the latent network inference algorithm, called *NetInf* [22] to infer policy diffusion networks connecting the American states over time. Although these works have made great progress in learning a policy diffusion network from multiple policies, understanding such a network across different political settings—such as to flexibly examine the diffusion patterns across different times, different regions, or across different topics—is challenging.

We propose a novel visual analytics system, called PolicyFlow, (Figure 1) that allows for interpreting and examining policy diffusion *in context*—that is, across various spatial, temporal, and multiple policy settings. We work closely with domain experts to design an interactive visualization system that helps answer the following relevant questions. What would the underlying diffusion network look like? Who are the leaders in the network? Do the network and leaders change over time and regions, and across various topics of policies? What are possible factors associated with the diffusion patterns? To what extent do the inferred patterns capture the observed data? In this work, we propose a suite of visual analytic tools to better explain and assess the results derived from the *black-box* network inference algorithm and aggregate policy adoption data.

To summarize, our contributions include the following:

(1) *System design*: We propose a visual analytics system that offers interpretable pattern discovery of policy diffusion. We provide a suite of analytics and visualizations that facilitates the examination of policy diffusion patterns across different facets—geo-space, time, policy topics, and similarities in political contexts. Our system enables users to infer the general diffusion pattern and support its interpretation with the concept of expected and deviant patterns of cascades to evaluate the conformity of the inferred relationship with the actual policy adoption.
(2) *Visual design*: We propose a novel matrix-based diffusion representation to facilitate the understanding of policy diffusion over time and space. Given the adoption cases of a policy and the inferred relationship between states, it offers a matrix-based visual overview of how source and target states are mapped to each other, and whether each adoption case falls into the expected or deviant pattern.
(3) *Evaluation*: We provide comprehensive evaluations to study the usefulness, understandability, and utility in real-world policy settings. First, two case studies reveal how users can take advantage of the system functionality to gain insights. Second, the user study evaluates how the system is intuitive and easy to use even for non-expert users. Furthermore, we conduct in-depth interviews with three experts, and we demonstrate that `PolicyFlow` can facilitate the understanding of the potential and importance of policy diffusion.

The rest of the article is organized as follows. Section 2 presents an overview of prior studies on visual analysis of text corpus and information diffusion. Section 3 discusses the proposed design with respect to the design goals. We present the analysis approach of policy diffusion pattern in Section 4, followed by the visualization and interaction design in Section 5. We demonstrate how the system effectively helps domain experts and non-expert users to gain insights on policy adoption patterns through its usefulness and understandability via case studies, expert interviews, and a user study in Section 6. In Sections 7 and 8, we conclude our discussion with the strengths and limitations of our study, as well as future work.

2 RELATED WORK

Here we describe related work in visual analysis of the text corpus and visualization of diffusion networks.

2.1 Visual Analysis of the Text Corpus

Understanding the text corpus is becoming more challenging due to the large amounts of information involved. There have been various approaches to mitigate the cognitive load of users' exploration on the large text corpus, but one of the effective ways is to visually represent the text corpus. Since the text corpus is extracted and collected from heterogenous contexts, many studies in visual analytics focus on the evolution of the text corpus in specific contexts [20, 25, 35], where a visual analytic system for each context needs to be specialized in addressing its requirements and challenges. Using MeetingVis, Shi et al. [35] analyzed the conversation of participants in the context of group meetings. To represent three components of the meeting as a storyline including participants, topic, and timeline, each participant's spoken words were processed and represented as a trend line. Then they were grouped as a topic bubble, with the keywords used to summarize the salient topic in the timeline. However, Fu et al. [20] and Hoque and Carenini [25] focused on summarizing the online conversation. Hoque and Carenini [25] especially captured the hierarchical structure of conversation topics but also supported the interactive editing of hierarchy to deliver the user feedback. T-Cal [20] analyzed the team conversation in online chatting, where they discuss a series of tasks to be done. They visually represented a series of tasks as multiple threads, along with the calendar-based summary view. Some studies [9, 37, 41] especially focused on effectively visualizing the dynamics of spatio-temporal patterns to identify the micro-blogging behaviors associated with their temporal and spatial aspects. VAUT [41] built a visual analytic system that analyzes spatio-temporal dynamics of user-specific topics. VAiRoma [12] especially focused on visualizing Wikipedia articles about the topical, spatial, and temporal aspects of Roman historic events.

Topic visualization is another popular approach in this domain used to summarize the semantics of the text corpus [8, 10, 13, 14, 16, 17, 21, 26, 28, 30, 31, 40]. Early works proposed methods that facilitated the exploration of relationships among topics from multiple aspects [10, 30]. Several studies represented the dynamic evolution of topics along with river flow–based visualizations to address different challenges. TextFlow [13] summarized the dynamics of keywords and relevant events along with topic flows where topics are split and merged along with the timeline, whereas RoseRiver [14] visualized how the hierarchical structure of the topics evolved over time through bar-shaped visual representations of the subtopics being split and combined. Sun et al. [36] introduced EvoRiver, a visual framework dealing with topics cooperating and competing to attract opinion leaders. LeadLine [16] focused on highlighting the event-specific topic bursts, whereas HierarchicalTopics [17] aggregated the derived topics from hLDA to better understand the trend and semantics of topics.

Visualizing the distribution of text entities such as words and documents has also been given more attention recently. The embedding techniques, including word embedding, are utilized to help understand how entities are semantically distributed with some extent of similarity. Park et al. [33] help users analyze a concept (e.g., tidal flooding) explained by groups of word vector clusters and specific documents displayed. The visual analytic system ultimately aims to support lexicon building for better construction of concepts. Some studies especially focus on the scholarly communication where citation patterns are captured along with keywords and documents [5, 24].

Our system, however, focuses on capturing the contextual information of policy diffusion pathways from the text corpus. A set of policy-related documents helps users interpret as evidence, and we infer how states exert their political power onto others in the policy diffusion process using the policy adoption data that consists of adoption sequences.

2.2 Visualization of Diffusion Networks

Information diffusion over a large scale of networks today has been traced more extensively and is thereby ubiquitous in data analysis. In understanding dynamically evolving networks, due to their complex structures and diffusion patterns, visualizing the diffusion network has gained attention in different contexts, such as link sharing [1, 2], rumor spreading [3], meeting [39], posting [29], and conversation [39]. Understanding the influential actors and relationships has been of great interest in diffusion analysis [19, 39]. IdeaFlow [39] visualized how multiple groups exchange or exert their ideas onto one another in a hierarchical manner. VisForum [19] analyzed the replying behavior of online conversation as a set of interactions, where users' activeness from replies and postings are represented as glyph and analyzed. D-Map [11] proposed a visual analytic system, which is generalized to analyze the ego-centric diffusion pattern. WeSeer [29] aimed at predicting the extent to which an individual posting propagates over the network, and visualizing the detailed trends of propagation speed and volume, and information over time.

A diffusion network evolves over space, time, and topic. A visual analytic system used as a monitoring tool is capable of exhibiting those multi-aspects of the diffusion pattern with multiple panels [9, 11]. Chen et al. [11] especially emphasized an instance-based exploration that highlights the diffusion pattern centered by a user in social media with the representation of the geo-spatial, temporal, and user group structure with a hexagonal grid map. Whisper [9] is a real-time monitoring system for social media that integrated multiple contexts such as geo-spatial and topical distribution of user groups and tweets, along with their diffusion pathways. Their retweet behaviors from a large number of tweets are monitored in real time and represented in a visual framework that borrowed a metaphor from the sunflower.

Our `PolicyFlow` focuses on analyzing the adoption cases of state policies by inferring the influential relationship between states. As it is likely that such general diffusion patterns may differ from spatial, temporal, and topical context, we also support deriving the network for specific contexts. Such functionality meets users' analysis needs in the political domain in that the policy-making process requires the exploration of the context and background. Our system can support data-driven exploration on the diffusion analytics.

3 DESIGN OVERVIEW

We designed `PolicyFlow` for interpreting the latent policy diffusion network in context. Following a user-centric design process, we worked closely with a group of domain experts—a team of political scientists who specialize in the study of policy diffusion in American politics. We scheduled a series of meetings over a year-long period to aid in understanding the system requirements and to refine our prototype. Our discussions centered on what kinds of patterns need to be captured and how to reveal and interpret them in various spatio-temporal contexts and political settings. We summarize the desired system requirements as follows.

3.1 Requirements

- **R1** **overview:** The system should offer a big picture of the general diffusion patterns learned from the historic policy diffusion dataset. It should allow users to identify the underlying diffusion network and leading states (i.e., influential nodes) in such a network.
- **R2** **context:** The system should help reveal the diffusion patterns in heterogeneous contexts. It should allow users to explore various questions. For example, how do the network and leaders change over time and regions, and across policy topics? How do the diffusion patterns associate with the socio-economic context of states? What are the relationships between the diffusion and geo-proximity?
- **R3** **structure:** The system should help reveal the structural details of the diffusion patterns. Specifically, how does a particular state influence or be influenced by other states? Do policies of similar topics exhibit similar diffusion patterns? How do we identify policies with similar diffusion patterns?
- **R4** **inference assessment:** The system should enable users to interpret and assess the patterns derived from the network inference algorithm. In particular, it should clearly show the extent to which the inferred patterns apply to a particular policy, state, or diffusion pathway.

3.2 System Overview

The proposed work, `PolicyFlow`, is designed and developed based on the aforementioned system requirements. As shown in Figure 2, the system allows users to interactively explore how state-

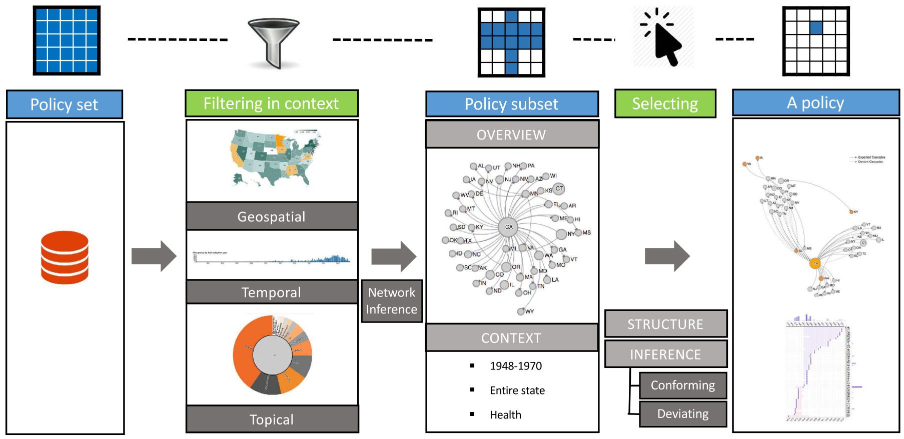

Fig. 2. The system framework of PolicyFlow. Given the dataset of policy adoptions and the metadata of states, the system supports (1) filtering the multi-dimensional political contexts and (2) selecting an individual policy and inspecting a policy-specific adoption pattern.

based policies are diffused in the United States. It allows users to browse and select a subset of policies from the entire policy database with the multi-faceted (geo-spatial, temporal, and topical) context filters (R2). Based on the whole or user-selected subset of policy texts, the system computes the underlying policy diffusion network and generates an overview of the network (R1), with multiple coordinated views showing the geo-spatial and states' socio-economic contexts (R2). Users can explore the structural details through interactions to reveal how states (nodes) are close to each other in terms of their network connectivity or policy adoption similarity (R3). Our system supports not only the exploration of aggregated patterns in context but also the examination of a single policy in detail. The contextual information of a policy is also displayed to help users analyze the diffusion patterns. The diffusion pattern of a policy is also visualized and diagnosed in terms of how the inferred diffusion patterns conform to or deviate from a particular policy's actual adoption sequence (R4).

4 ANALYSIS OF POLICY ADOPTION

4.1 Policy Data

Here we use the policy adoption data collected by Cao et al. [9]. The scope of policies in the system includes 764 state-wide policies over 300 years ranging from 1691 to 2017. The dataset consists of the metadata of states, policies, and the set of policy adoption cases. Each policy is associated with a history of adoption that indicates which states have adopted this particular policy and the timeline of adoption. The additional information on policy adoption includes policy subject (e.g., health, education), start year (the year of the first adoption), end year (the year of the last adoption), and the number of states adopting the policy. In the policy data, the underlying diffusion network (i.e., who influenced whom in the adoption decision) is not observable. The process of inferring such an underlying network will be detailed in the next section. Another type of information collected in our dataset is a set of state-related attributes that provide key socio-economic context. These include per capita income, minority diversity, legislative professionalism, citizen ideology, total population, and population density. Such attributes provide theoretically important covariates [9] and contextual information for users in analyzing the diffusion patterns. We provide analytical modules that may help hypothesize how each state's attributes correlate to its role in the propagation of political agendas as a follower or influencer.

4.2 From Adoption to Diffusion Network

As mentioned earlier, the policy adoption data allows for observing how a policy was adopted over time and across states, but the underlying influence or diffusion network—that is, who tends to lead and who tends to follow—is often unobservable [15]. The goal of network inference is thus to infer a *latent diffusion network* of political actors (i.e., states) based on observable data related to the repeated adoption choices that those actors make.

We apply the latent network inference algorithm, `NetInf` [22], to infer policy diffusion networks. A policy's adoption history can be considered as multiple network *cascades*, where each cascade consists of a sequence of adoption cases called *contagions*. A contagion c, denoted as a tuple $(p, u, v, t_v)_c$, means that the adoption of a policy p has spread from state u to state v at time t_v. A directed edge $u \rightarrow v$ connecting a pair of state nodes was used to indicate that policies diffuse from *source* node u to *follower* node v. The set of cascades for a specific policy p can be obtained by grouping the contagions by policy, denoted as $\{(p', u, v, t_v)_c | p' = p\}$.

In practice, a contagion can only be observed through $(p, v, t_v)_c$ that describes the time t_v when node v became infected by the contagion c of a policy p. The `NetInf` algorithm aims to recover the unobserved directed network G^* (i.e., the policy diffusion network over which the contagions spread). The algorithm is based on a probabilistic model that provides the probability of a contagion c between a pair of nodes u and v, and the probability that contagion c propagated in a particular cascade tree. Based on this, $P(C|G)$, the probability of a set of cascades C occurring in G, can be obtained, and the latent diffusion network G^* as a final output of the algorithm is approximated by $\hat{G} = \text{argmax} P(C|G)$ with a sparsity constraint on all possible G. Note that a latent diffusion network can include cycles between states, which indicates the possibility of having edges $u \rightarrow v$, $v \rightarrow w$, and $w \rightarrow v$ among nodes u, v, and w, for example. This indicates that the influential relationship may not be simple: when aggregating multiple policies to infer a general network that encodes the interaction between states, certain states tend to influence other states for some policies, but at the same time, those source states are likely to be influenced in turn by those other states. Such structure indicates the complex relationship among states when considering policy adoptions.

The time complexity of the algorithm depends on the structure of the network to be inferred from the cascades. According to Leskovec and Krause [27], the underlying network (the best propagation tree) can be found in time linear in the number of edges "by simply selecting an incoming edge of highest weight for each node." With the greedy algorithm and the two speeding-up improvements, the algorithm has been shown to be scalable in a large network.

5 INTERACTIVE VISUALIZATION

We implement a visual analytic system to offer a better understanding of the trajectory of policy diffusion and the relationship between states. Each component not only supports visualizing the multi-dimensional contexts but also coordinate with each other to meet the requirements described in Section 3.1. Throughout the system, the consistent system-wide color (gray and orange) scheme is chosen based on the following criteria: (1) there is no overlapping with other colors in the system, (2) they are visually distinguishable from each other, and (3) they are color blind–friendly. The following sections address each requirement in the system in Section 3.1, but at the same time, they describe how each component works (Section 5.1) and how multiple components interactively support the system requirements (Sections 5.2 and 5.3). Finally, the policy-specific patterns and detailed information are described in Section 5.4.

5.1 Overview of Spatio-Temporal Contexts

PolicyFlow's primary goal is to give users a comprehensive overview of the general diffusion pattern (R1). Several visual components in our system provide different aspects of the network (by time, by region, or by topic) independently.

Centered within the system is Network-View. A diffusion network is a directed network where an influencer state u and a follower state v is connected by an edge $e = (u, v)$. By default, the system renders the general diffusion network inferred from the full set of policies (Figure 1(a)). The node size is adjusted by the influence, which includes five node centralities (outdegrees, betweenness, closeness, page rank, and hit) [32] available under the "influence" dropdown menu. Network-View represents either a network of general diffusion patterns or a specific policy. When a specific policy is selected, Network-View re-renders the network with the information of how the adoption sequence of the specific policy conforms or deviates from the general diffusion (see Section 5.4.1 for details of expected and deviant patterns). Those two networks are rendered differently, as shown Figure 1(a) and (c), where 23 state nodes in the policy network of "Mandated Coverage of Clinical Trials" (Figure 1(c)) are connected by solid lines (*expected* cascade) and dashed lines (*deviant* cascade), whereas the general network does not represent such edges.

Subject-Browser (Figure 1(b-i)) and Content-Browser (Figure 1(b-ii)) displays the categorical and topical distribution of policies. For Subject-Browser, we obtain the policy categories from our policy metadata, which includes its name, category, and first and last adopted year. To capture the topical dimension of the policies in Content-Browser, we collect and present, for each policy, the top five documents retrieved using Google search. Then we leverage latent Dirichlet allocation (LDA) [7], a topic modeling approach that obtains the latent topic distributions based on the probabilistic distribution, to analyze the semantic relationships of the collected documents. We present K clusters as a result of the topic modeling ($K = 20$ in this work), and for each cluster, the 15 most salient words and the number of policies are shown in the tooltip when a user hovers over a topic in Content-Browser.

We visualize Subject-Browser and Content-Browser as a pie chart, for the purpose of (1) providing a space-efficient visualization and filtering module, and (2) providing categories and derived topic distributions over the entire policies. Here, a pie chart is an intuitive visualization that meets our requirements in that it is not only comprises pieces of items (i.e., categories or topics), whose size represents its proportion, but also users can select one piece at a time to filter policies by category. The detailed information, such as the number of policies within each category and related keywords for each topic, is shown when a user hovers over an individual category or topic in the pie chart. Other alternative designs were considered but ultimately were excluded due to their limitations. For example, a word cloud can provide a keyword-based overview and effectively visualize the semantic aspect of each topic derived from LDA, but it is not well suited to the purpose of the per-category summary module with its proportion and functionality as a context filter. Some of the limitations in the current visualization that stem from the design of the pie chart still exist. In the pie chart, a category with a small number of relevant policies is not likely to take enough space to present the name of topic/content and its keywords. To get around this issue, we additionally provide a tooltip of each category or topic when it is hovered over. We take advantage of those spaces to display the name and relevant keywords, as those contents are likely to be verbose, thereby requiring significant space even with other alternative visualization designs.

Timeline-View linearly shows the frequency of policy adoption throughout the entire policy adoption history (Figure 1(e)). Map-View represents the spatial locations and geo-proximity of states with two modes: state-wide (Figure 1(e)) and regional (Figure 1(b-iii)).

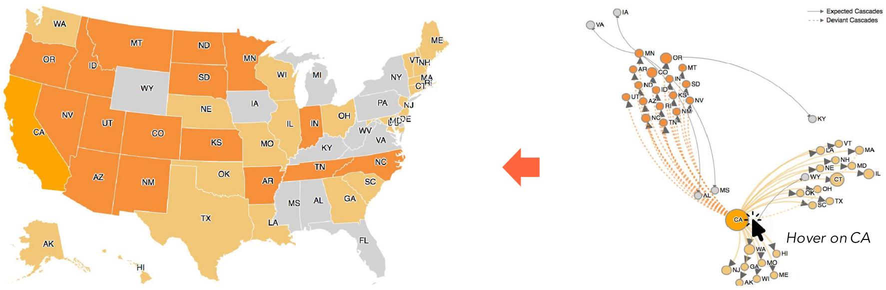

Fig. 3. Examining structural details via "Connected" mode. `Map-View` and `Network-View` pair renders the influential relationship of states regarding the policy "Framework for Donation of Organs, Other Body Parts." On hovering over CA in the network, the Connected mode highlights the states that are connected to California (with lighter orange for expected cascades and darker orange for deviant cascades).

5.2 Highlight of Diffusion Structural Details

We support the exploration of the similarity and connectivity of states and policies in a network with interactive visualization (R3). Oftentimes, the network layout comes with visual clutter such that users may not easily understand the structural details as a network involves several nodes and edges. Exploring a node and its influencer and follower states is an important task in the system. Therefore, we set it up as one of our requirements (R3), which is for the system to decompose the structural details of relationships between states or policies. Specifically, we attempt to help users answer these questions: How are *states* similar to each other in terms of their connections and similarities from the network? What are the similar *policies* in the aspect of content and cascade?

To answer the first question, we provide two hover modes to examine the similar *states* in the policy network view: Connected and Similar (Figure 3). When users hover over a specific state, the system highlights the connected states (i.e., states that are connected to the hovered state via edge) or the five most similar states in the network (i.e., states that share the most connected nodes in common). Users can select one of two modes by adjusting the slide bar.

Then, to answer the second question, `PolicyFlow` derives the similarity of *policies* from two perspectives: (1) content similarity by calculating the pairwise TF-IDF score of relevant articles in our dataset and (2) cascade similarity derived by the pairwise similarity of adoption sequences. The cascade similarity is specifically measured by (Jaccard score × Kendall score) of adoption sequences from two policies. For two policies p_A and p_B of our interest, and the adopting states S_{p_A} and S_{p_B}, we first identify the common states within two sequences and calculate how those common states are proportional to the union of all states within two sequences using the Jaccard score where $J(S_{p_A}, S_{p_B}) = |S_{p_A} \cap S_{p_B}|/|S_{p_A} \cup S_{p_B}|$. After obtaining two sequences that only preserve the common states $S_c = \{s | s \in S_{p_A}, s \in S_{p_B}\}$ from the first step, we calculate the ranking correlation of the same length of two sequences by calculating the Kendall distance. The final similarity score is calculated by the multiplication of these two scores. The higher the score is, the more similar two adoption sequences are. `Policy-Detailed-View` is dedicated to providing this information along with the detailed policy information. If a policy is selected, similar policies are listed with their similarity score. Users can use the tab interface to toggle between the two lists.

5.3 Filter on Geo-Political Contexts

The visual components in `PolicyFlow` not only serve to explore the overview but also to filter the dataset by context. In our system, the interaction of three filters allows users to narrow down to a subset of policies (Figure 4) with the coordination of multiple views and user interactions as

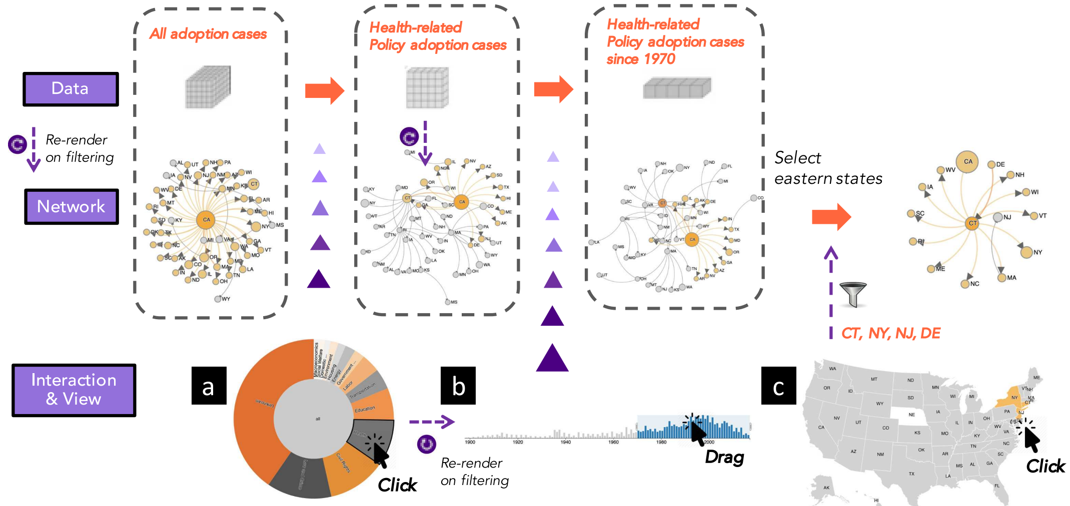

Fig. 4. An interaction scenario involving three visual components for the purpose of filtering the multi-dimensional contexts (R2). By selecting a topic (or a category) (a), or a certain period of time (b), a subset of policy adoptions is selected, which triggers updating the inferred diffusion, followed by re-rendering `Network-View`. (c) When users select some states in `Map-View`, the corresponding nodes are visually filtered from the existing inferred network.

described in Figure 4 (R2). When a user is interested in a specific category or topic of policies, they can start from `Subject-Browser` or `Content-Browser` by selecting a static pie chart that shows the overall distribution of policies in terms of their categories or topics. Then the policy adoption cases in the current analysis are limited within the scope of the selected category or topic, which makes the system update the general diffusion network in real time given the set of policy adoptions. `Network-View` and `Timeline-View` are re-rendered based on the new inferred network and dataset. In the same manner, selecting a time period by dragging a part of `Timeline-View` also updates the selected policy adoption cases (i.e., limited to ones from 1970 to 2000), which triggers inferring a new diffusion network and re-rendering `Network-View`. `Map-View`, however, serves to select some nodes (i.e., states) of users' interest, to visually filter the nodes within the current inferred network (without updating it) and help users focus on analyzing the selected nodes and their relationship. We also note that when the network is being re-calculated and updated in real time, by being triggered upon selecting a new set of policies, such as when users adjust `Content/Subject Browser or Timeline View`, running the algorithm in real time takes a few seconds at most (given the moderate size and number of cascades (764 policies as the total number of cascades; 50 states in a cascade at maximum), and users are provided with visual feedback for the status update.

`PolicyFlow` also supports exploring and identifying some state-related factors and their correlation with the inferred relationship. Although inferring a network from the policy adoption cases that occurred in history gives us an understanding of the post hoc relationship between states, the reasoning behind such relationship is another dimension of understanding the policy diffusion. Our system contains a set of the socio-economic status of states as metadata. It consists of per capita income, minor diversity, legislative professionalism, citizen ideology, total population, and population density in the dataset. Our system computes the rank correlation between the node centrality of the policy network and socio-economic attributes in real time. When a policy is selected, the correlation measure is updated along with the attribute names in the dropdown menu.

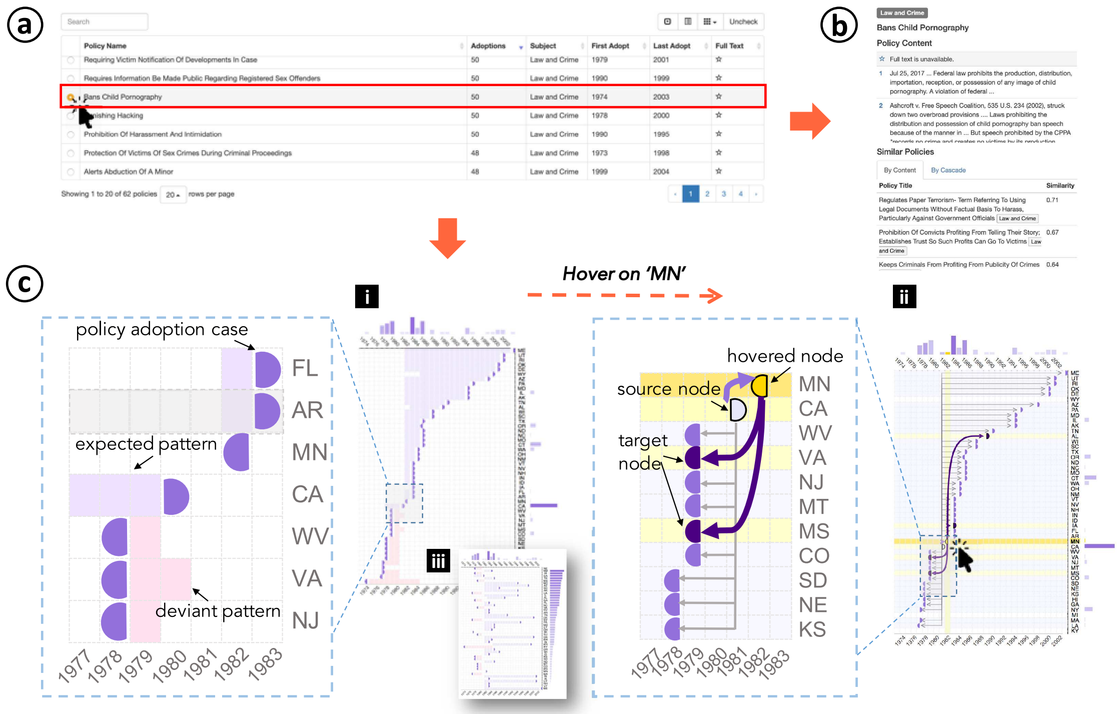

Fig. 5. The interaction between components on the selection of a policy "Bans Child Pornography." Policy-List-View displays a set of policies with their additional information (a). When a user selects the policy, Policy-Detailed-View shows detailed information on the selected policy (b), and Policy-Inspection-View conveys the diffusion pathways of the policy (c-i). When a user hovers over a semi-circle (e.g., an adoption case where MN adopted the policy in 1982), all source (colored as light purple) and target nodes (colored as light purple) are specified with arrows (c-ii). The matrix view is capable of being ordered by state attributes (e.g., minor diversity in the screenshot) in the y-axis (c-iii).

5.4 Inspecting on Policy-Specific Adoption Pattern

Separate from filtering over a set of policies, our system supports inspecting the policy-specific adoption pattern (R4). It starts from `Policy-List-View` (Figure 5(a)), which lists all policies within the dataset in a table view on default, where users are not only allowed to select an individual policy but also browse through policy metadata. The list of policies within `Policy-List-View` is coordinated with context filtering introduced in Section 5.3 such that it displays a subset of policies selected from the context filters by user. Once a policy is selected, `Policy-Detailed-View` (Figure 5(b)) displays the policy-specific information, including the policy name and category, full text data, and the list of similar policies. `Policy-Inspection-View` is also enabled for users to inspect the policy-specific diffusion pattern, which will be introduced in detail in the following section.

5.4.1 Policy Inspection View. `Policy-Inspection-View` visualizes how individual policy diffuses through states. Different from other components serving as a context filter, `Policy-Inspection-View` associates with a single policy and provides the summary of policy adoptions when users select a specific policy.

`Policy-Inspection-View`, shown in Figure 5(c), is essentially a two-dimensional array where each cell is a spatio-temporal incidence of policy adoption (s, t) with a state s from the vertical axis and a year t from the horizontal axis. Each row is then a timeline of a state's policy adoption history. Here we choose a matrix view among the possible alternatives of dynamic graph visualization [4] for two reasons. *First, overview*: a form of matrix can provide a visual overview of multi-dimensional aspects of policy adoption. The temporal dimension as a horizontal axis can line up from the left to right to indicate the intuitive temporal orientation. *Second, ordering*: a matrix is capable of ordering the axes. In our visual component, the spatial dimension of the states can be ordered by their properties (e.g., ordering states by minor diversity in Figure 5(c-iii)). *Third, focus*: the matrix view can easily encode individual adoption cases as an individual cell that marks when the adoption case is adopted by which state, then the target and source state can be highlighted with its arrows. However, the alternative designs to effectively represent a dynamic graph, such as a network view, may not able to clearly visualize such spatio-temporal patterns and the adoption incidents visually encoded in one view.

Along with `Policy-Inspection-View`, we introduce the concepts for interpreting the policy adoptions against the general diffusion pattern: *expected* and *deviant* pattern. Once the system infers a latent diffusion network from multiple policies, each policy adoption is associated with an edge $e(u, v)$, where u is an inferred source state and v is the state that adopted the policy regarding the policy adoption case. These two types of patterns $\mathcal{E} = \{e(u, v)\}$, *expected* and *deviant* patterns, serve to indicate conformity between the model's inference from a set of policies (i.e., u influences v in the general network) and the time dependency (i.e., which state (u or v) actually adopted the policy earlier), and visualize such patterns in this view. Specifically, an *expected* pattern $e_{expected} = \{e(u, v) \mid t_u < t_v\}$ indicates that the policy was adopted to the influential state u and the follower state v identified in the inferred network in chronological order. However, a *deviant* pattern $e_{deviant} = \{e(u, v) \mid t_u > t_v\}$ is where the follower state u in the inferred network actually adopted the policy earlier than the influencer state v (i.e., the adoption sequence does not match the influential relationship inferred based on a set of policy adoption cases).

An example of the `Policy-Inspection-View`, with a detailed adoption pattern of the "Bans Child Pornography" policy, is presented Figure 5. In Figure 5(c-i), a semi-circle indicates a single event of adopting this policy by Florida in the corresponding year. Two different types of the strips coming to a semi-circle icon represents the *expected* or *deviant* pattern of policy adoption. We visually represent a *expected* pattern as a purple-colored strip with a semi-circle heading forward (i.e., the expected direction where the time goes) and a *deviant* pattern of an adoption as a pink-colored strip with a semi-circle heading backward. As a summary of overviewing the expected and deviant pattern with respect to the selected policy, we provide a summary statistic of those patterns called the *conforming score*. We calculate the conforming score by $|\mathcal{E}_{expected}|/|\mathcal{E}|$ of each policy. In other words, the score is the proportion of expected edges among all edges, indicating how the adoption sequence of selected policy conforms to the general pattern. We provide the conforming score of each policy within `Policy-Inspection-View`.

On hover of a semi-circle component, `Policy-Inspection-View` provides the connectivity between nodes within the general diffusion pattern. In other words, the layout reveals the egocentric network of the state in a way that all edges connected to it are represented. In Figure 5(c-ii), the hovered node is represented as a yellow node. The incoming and outgoing edges connect the node to its source and target nodes, which are colored as light and dark purple. We also specify which other nodes are influenced by its source node. The gray edges coming from the source node indicate in Figure 5 that CA is a great influencer impacting on most of the other states.

6 EVALUATION

This section presents our evaluation of `PolicyFlow` through three studies. The first is the *case study*. We present two usage scenarios and demonstrate how the system meets its design

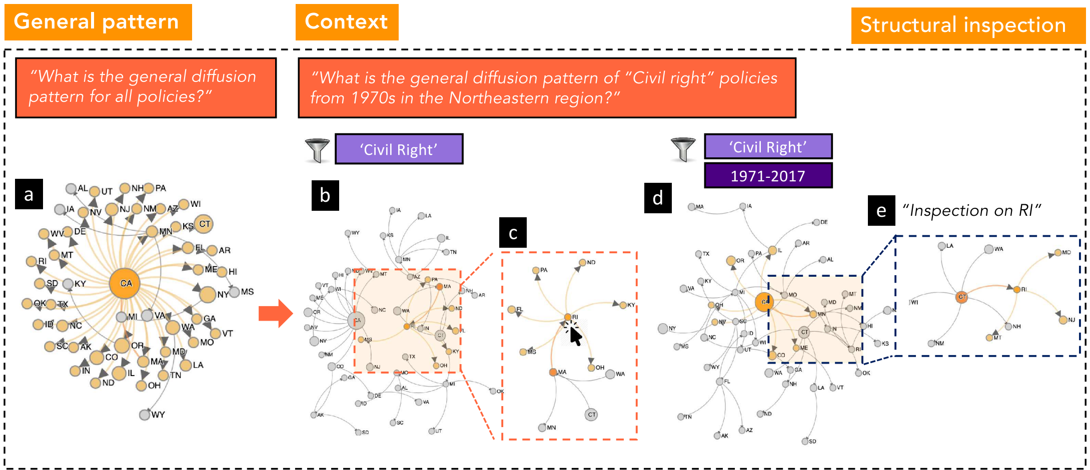

Fig. 6. Case study 1. (a) The general diffusion pattern derived from the entire policy dataset. It shows that CA is a dominant state in the diffusion network. (b) The diffusion network after selecting "Civil Rights" policies and the highlight of the sub-network connecting to RI in (c). (d) The diffusion network after further selecting a time period and the sub-network connecting to RI in (e). Colors indicate the focal nodes—orange: the mouse-hovered node (focal state), light orange: followers of the focal state, and dark orange: influencers of the focal state.

requirement and how its functionality facilitates pattern discovery and in-depth analyses on policy diffusion. The second is the *expert interview*. We conduct interviews with three domain experts who have background knowledge in policy diffusion. We present how the system helps domain experts examine their questions of interest or explore new research questions. The third is the *user study*. We conduct a user study on non-expert users to evaluate the system's usability and effectiveness to meet its design requirements.

6.1 Case Study

In this section, we demonstrate the effectiveness and utility of `PolicyFlow` in helping users gain a better understanding of policy diffusion. We consider political analysts such as political scientists, policymakers, and lobbyists, and members of the NGO, as the main users of our system. For those experts, understanding the trajectory of policy adoptions over states is a crucial task for planning and creating a new policy, and moreover, predicting how different aspects of a new policy may unfold in the near future. Here we provide two usage scenarios to show how users can use `PolicyFlow` to gain insights from the policy dataset.

Case 1: Abortion ban. We use policies related to "abortion ban" to demonstrate how `PolicyFlow` can be used to obtain a big picture of the policy diffusion on a given topic, and acquire a deeper understanding about specific policies through the system's functionality.

Anna is a policy analyst interested in understanding the diffusion of "abortion ban" policies that can be traced back to the Northeast region in the late 1900s. Substantive questions related to the topic include the following: *In general, what would the diffusion pattern on this topic look like? How can a specific role in the network or a specific period of the network be examined? How can policies with different diffusion dynamics be characterized? How can correlated factors (e.g., socio-economic attributes of the states) for patterns of diffusion be analyzed?*

Anna tried to answer these questions using `PolicyFlow`. By default, the system showed the overall diffusion pattern in `Network-View` (Figure 6(a)), where the diffusion network was derived from the entire set of policy data. By looking at this network, Anna saw an overall picture of

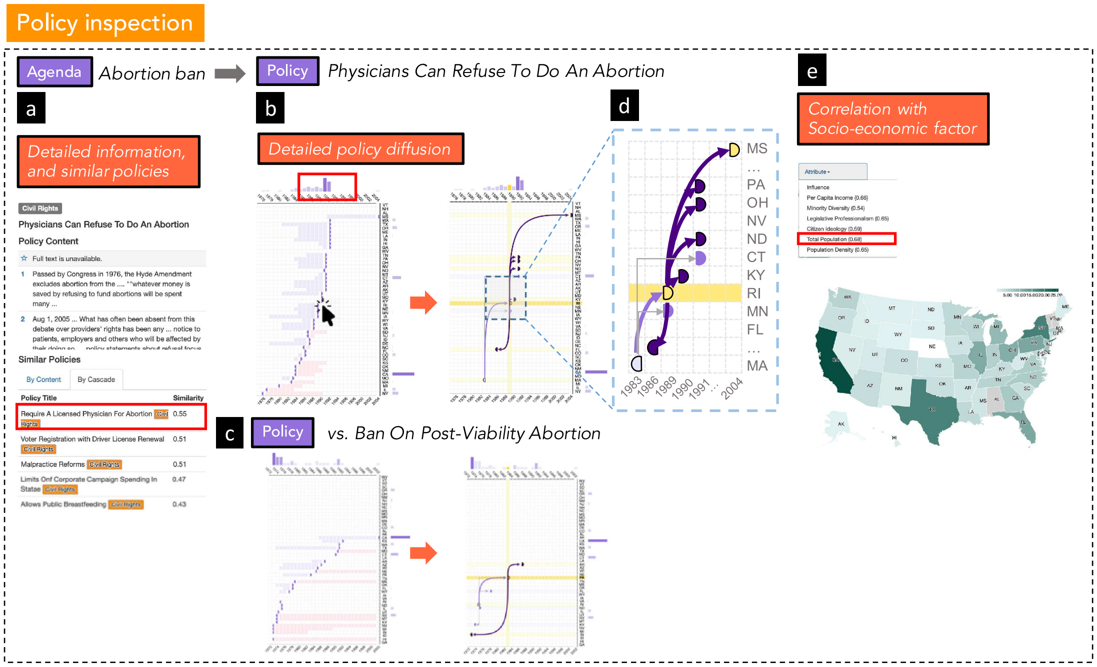

Fig. 7. Policy inspection of Case 1. After selecting the policy "Physicians Can Refuse to Do an Abortion," the detailed information and similar policies were displayed (a), and the adoption sequence was shown in `Policy-Inspection-View` (b), with the state RI highlighted (d). The user can enumerate the socio-economic attributes that are highly correlated with the adoption sequence (e) or compare the sequence with that of another policy (c).

how policies diffuse throughout history (R1). She also looked at the influential states, such as California, a state with great policy influence to other states, which was positioned in the center of the network. Anna also checked the policy adoption frequencies across time from `Timeline-View` and across subjects from the `Subject-Browser`. By clicking on the "Civil Rights" category on the `Subject-Browser`, she narrowed down the set of policies to be rendered in `Network-View` (Figure 6(b)) (R2). This view allowed her to further look at the key players in this category (R3). She then selected RI and its follower node MA in `Network-View`. The two states, as well as their directly connected states (influencers and followers), were highlighted with various colors (Figure 6(c)). She further narrowed down the policy selection using `Timeline-View` to select policies adopted after the 1970s (R2). The network was immediately updated (Figure 6(d)), through which Anna found a slightly different pattern, where RI was likely to be influenced by CT, and only influenced NM within this period. She also found that CT had a greater influence than other Northeastern states as the node with a larger size (reflecting its "outdegree"). When Anna chose to render node size with "closeness" (using the dropdown menu of network centrality measure), CT became the second largest influential state, indicating CT's influence on many states (large "outdegree"), although the influence may not be evenly distributed.

Anna decided to take a closer look at the policy "Physicians Can Refuse to Do an Abortion," as the system shows this is the most adopted policy (47 states) in this category. When a policy was selected, the system showed `Policy-Inspection-View` and `Policy-Detailed-View` (Figure 7). `Policy-Detailed-View` displayed the policy contents and the list of similar policies as shown in Figure 7(a). Anna found that, interestingly, the most similar policy in terms of policy cascade similarity was "Require a Licensed Physician for Abortion," which is also related to the abortion issue—suggesting that the two policies on similar issues have great similarity in terms of how they were adopted over time. She also discovered that the five most similar policies in terms of policy cascade similarity all fell into the "Civil Rights" category, which provided her a sense of how relevant policies diffused in a relatively similar way. Anna checked the detailed diffusion pattern of the policy "Physicians Can Refuse to Do an Abortion" in `Policy-Inspection-View` (Figure 7(b)) (R4). The pioneer states regarding this policy were NY and IL, which adopted the policy in 1976 and 1978, respectively. Starting in 1984, the policy was quickly adopted by the other 45 states within 10 years. She checked another policy of her interest, "Ban on Post-Viability Abortion" in `Policy-Inspection-View` (Figure 7(c)) to compare the diffusion pattern and found that the two policies exhibited different adoption dynamics in terms of adoption speed. The former policy was abruptly adopted within 10 years over 40 states (Figure 7(b)), whereas the latter policy had diffused over almost 30 years, with its initial 11 adoptions within the first 3 years (Figure 10(c)). When she came back to the analysis of the policy "Physicians Can Refuse to Do an Abortion," Anna found that by looking at `Policy-Inspection-View`, most adoptions were expected by the model (the general diffusion pattern derived from the set of "Civil Rights" policies) when she looked at the ratio of expected and deviant patterns encoded by purple and pink strips. Such visual encoding also allowed her to investigate how many years each policy adoption took between the source and target state's adoption year. Finally, Anna checked what socio-economic attributes of the states best explained (correlated with) the adoption sequence. By enumerating the "Attribute" dropdown menu (Figure 7(e)), she found that "Total Population" had the highest correlation score (0.68) with the adoption sequence of the selected policy. She also further checked `Map-View` and `Network-View` to find states with large "Total Population." This particular functionality allowed her to see what socio-economic statuses were common in those adopted states (R2).

Case 2: Marijuana use. We use policies related to "marijuana use" to demonstrate how to use `PolicyFlow` to discover similar policies with contrasting diffusion patterns.

Bob is a policy researcher who has been tasked with comparing policies related to marijuana use and smoking—both are health-related policies. Important questions in this comparison include the following: *Did the two kinds of health-related agenda have different policy diffusion characteristics? If so, how were they different in terms of adopted states (spatial) and adoption speed (temporal)?*

First, Bob looked at `Policy-List-View` and searched relevant policies with keywords "marijuana" and "smoking." He retrieved six and three policies from each of the keywords, and all of them belonged to the health category shown in `Subject-Browser`. To obtain a finer-grained subset, he checked the `Content-Browser`, which clusters the policies based on their content. Here he found that the two health-related agendas fell into two clusters (Figure 8)—cluster 2 had 69 policies and was associated with keywords "marijuana, gun, bully, etc.," and cluster 4 had 15 policies relating to keywords "smoking, drinking, place, restaurant, etc." He checked the policies in the two clusters. As shown in Figure 8(a) and (c), the two clusters appeared to be different in terms of pursuing values: cluster 2 consisted of issues and agendas including gun control, cybercrime, and bullying crime, which were related to the debates between rights and regularization (one of which was marijuana use), whereas in cluster 4, policies were related to public health.

6.2 Expert Interview

We conducted expert interviews to better understand whether the proposed system achieves its design goals, as well as its strengths and limitations. We invited three domain experts—three professors in different political science and public policy departments in two American universities. They had sufficient background knowledge about the policy-making process and policy adoption. One expert's research specialty was even in policy diffusion.

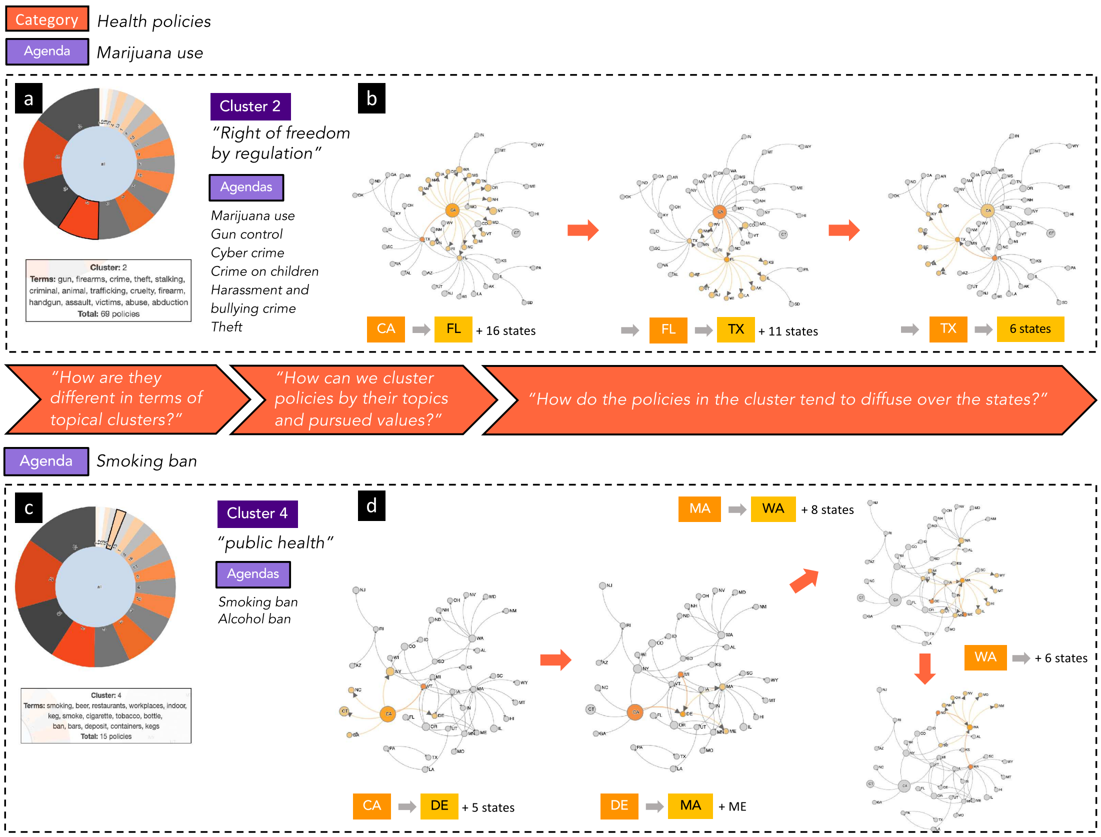

Fig. 8. Comparative analysis of Marijuana Use and Smoking Ban agendas in Case 2. (a) Marijuana use policies belong to cluster 2 with other agendas that are related to the right of freedom by regulation. (b) The system allows the user to decompose the detailed structure of policy adoption traces over states (driven by a few states including CA, FL, and TX). (c) Smoking ban policies belong to cluster 4 with alcohol ban policies. (d) These policies tend to be adopted more gradually than those in cluster 2.

Procedure. We conducted semi-structured interviews with the three experts. In each of the interviews, we began with an introduction of the system functionality, followed by a 10-minute warm-up session for experts to explore and get familiar with the system. We then had a brainstorming session where we conversed with the experts to identify a policy topic of their greatest interest. With the identified topic, we asked the experts to list some research questions and explore how our system can facilitate the analysis of these questions. From the three interviews, we examined how `PolicyFlow` helps gain insights into policy analyses on significant social/political topics (with experts A and B) and how `PolicyFlow` helps explore new research questions (with expert C).

Interview with Expert A: Policies on consumer protection. Expert A is a professor of public policy who specializes in privacy issues in cyberspace. In the brainstorming session, she identified that an interesting policy related to her research was the "California Consumer Privacy Act 2018," which is centered on the debate of users' right to control the data sold or sent to third-party providers. Related aspects include California's leading role in disseminating the legislation, other states' related effort and participation, and the role of federal legislation (Figure 9(a)). From the system design perspective, we sought to understand (1) who would benefit from the visual analytics and (2) how the system's spatial, temporal, and topical exploration functionality aid in conducting the policy analysis.

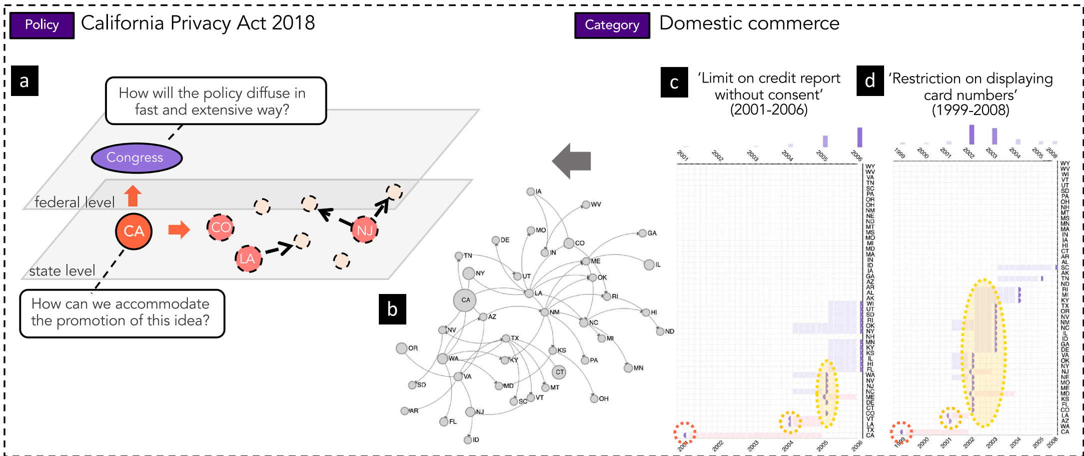

Fig. 9. The political setting and system exploration on the California Privacy Act 2018. (a) The current political status of this legislation is summarized in multiple layers of political settings. California plays a leading role as nationwide initiatives spread. Political actors with dashed lines are the projection by exploring past policies relevant to consumer information shown in (b), (c), and (d). (d) The visual summary of `Policy-Inspection-View` indicates that both of the two policies were initiated by CA, and a series of adoptions by multiple states after a few spreader states adopted the policies.

In the system, Expert A found two related consumer protection policies ("Limits Credit Agencies from Issuing a Credit Report without Consumer Consent" and "Restrictions on Displaying Credit Card Numbers on Sales Receipts") and discovered that, in their general diffusion network, CA played a leading role as a major influencer to LA, WA, and TX, and LA and WA were the next major influencers, reaching out to other the 22 states (Figure 9(b)). By further checking the two policies in `Policy-Inspection-View` (Figure 9(c)), she noticed that the adoptions were characterized by immediate diffusion over 31 and 25 states for each policy within 10 years. Expert A commented that `PolicyFlow` can at least meet two different purposes: "The leading states, like California, can be observed from the historic diffusion pattern and learn what would be the next possible state to best exert their influence. The policymakers across states can also learn from the historical patterns to predict how a policy would be accepted widely for the near future." She also noted that the system can help policymakers obtain a bigger picture and motivate them to investigate the circumstances of other states that were connected through related policies. She also pointed out a limitation of our system: "While the state-wise interaction is important, it does not consider the more complex interaction between state governments and the federal system, which could significantly affect the policy adoption at the state level."

Interview with Expert B: Policies on education. Expert B is a professor of political science. He mentioned his interest in policies about "school choice" and the political dynamics before and after the Trump administration surrounding this topic in the brainstorming session. Related questions included the role of Midwestern states and other advocates in changing the educational system (Figure 10(a)). From the system design perspective, we were interested in (1) how the system would help examine the role of Midwestern states and (2) how the policy advocates would use the system to best plan their strategy in promoting certain policies.

In the system, Expert B found that both School Choice and Charter School legislation were included in the dataset, and he was able to see the diffusion of educational policies in general (Figure 10(b)). Using the system, Expert B confirmed the vital role of Midwestern states as he mentioned earlier but with more concrete evidence, where OH was the most influential state,

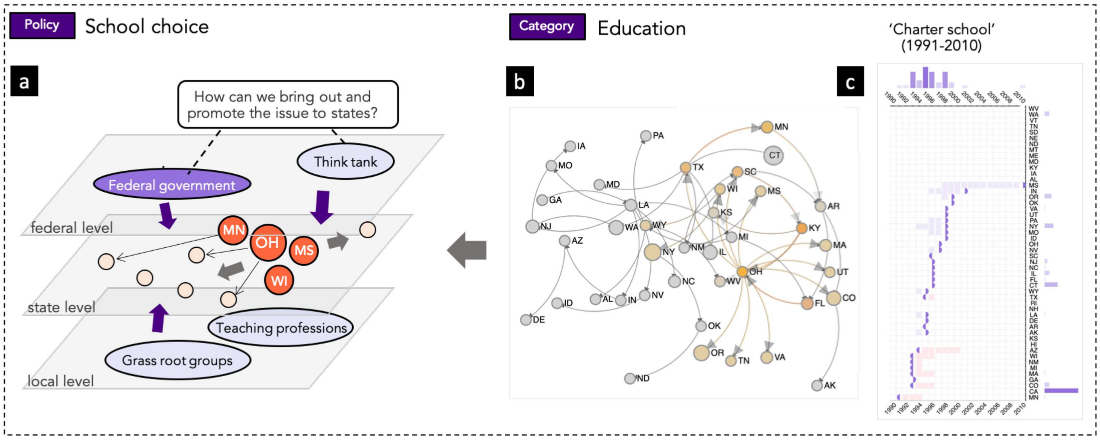

Fig. 10. The political setting and system exploration on School Choice. (a) School Choice is characterized by the active role of Midwestern states in spreading the innovation and interaction of policy entrepreneurs in the process. The overall diffusion pattern of the "Education" category demonstrates the leading role of Midwestern states (b), and `Policy-Inspection-View` summarizes the rapid spread of "Charter School" policy coming from the nationwide attention with policy entrepreneurs' involvement (c).

and other Midwestern states including MI, MN, WI, and KS also acted as great influencers. In terms of the system utility, Expert B especially mentioned that those "policy entrepreneurs" (he referred to this concept as individuals and organizations who serve as a catalyst for the diffusion of innovations around governments including think tanks, teaching professions, and grass-roots groups) can use the system to strategically plan which states should be allocated resources next, according to how a policy is likely to diffuse. During the interview, he especially acknowledged the benefit of having "the integrated tool of comprehensive policy dataset and the policy-related text information extracted from search results."

Interview with Expert C: Future Research in policy diffusion. Expert C is a political scientist who specializes in public policy diffusion research. In the brainstorming session, he mentioned his interest in exploring multi-dimensional aspects of abortion ban policies, which involves aspects such as civil rights, morality issues, and health concern. From the system design perspective, we were interested in how the system would help explore multi-dimensional aspects of abortion ban issues.

In the system, Expert C found a total of 29 policies related to abortion ban (all in the "Civil Rights" category within `Subject-Browser`). He also found that in addition to `Subject-Browser`, the system provided another semantic dimension of policies with `Content-Browser`. He found it interesting that related policies were split into three topic clusters with different keywords and diffusion patterns (Figure 11). He commented that such semantic exploration was useful in understanding the public policies in depth—how they formed as a group by a similar agenda. Expert C also noted that with the with multi-faceted filters currently provided in the system, it was helpful for him to efficiently browse various general diffusion networks with different context settings. However, "It would be more useful to have multiple filters at once (e.g., selecting multiple categories at once)" to help Expert C explore multi-dimensional aspects.

6.3 User Study

We test the usability and effectiveness of the system by conducting a user study that targets non-expert users who do not have background knowledge related to policy diffusion.

We recruited 15 participants (age: 20–28 years; gender: 6 female and 9 male). All participants were currently college or graduate students in a university, from a variety of disciplines, including

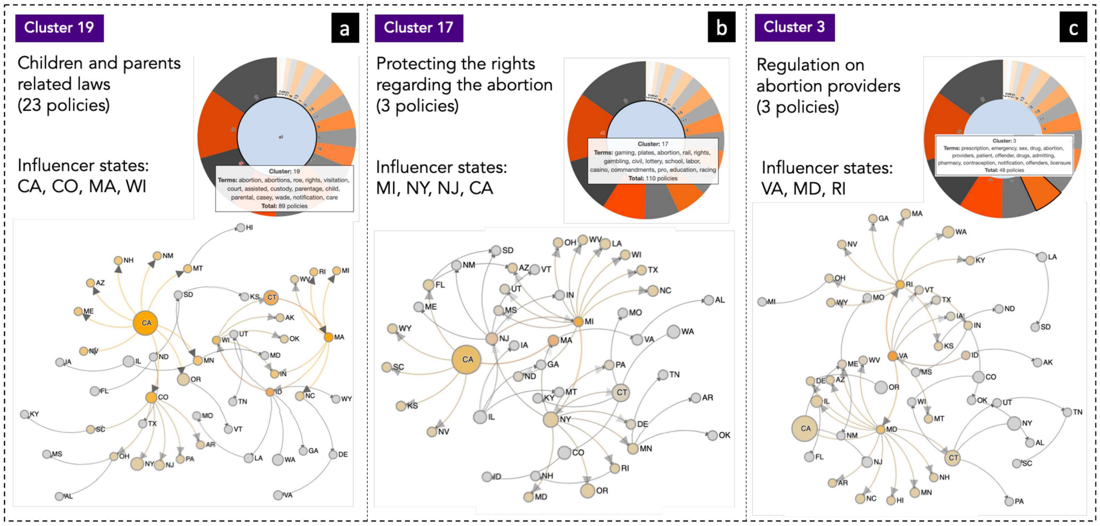

Fig. 11. Three topical clusters where 29 abortion ban policies belong. We capture that those policies are divided and grouped in three clusters by semantic aspects. This figure includes the interpretation of three sets of abortion policies and diffusion patterns with influencer states listed.

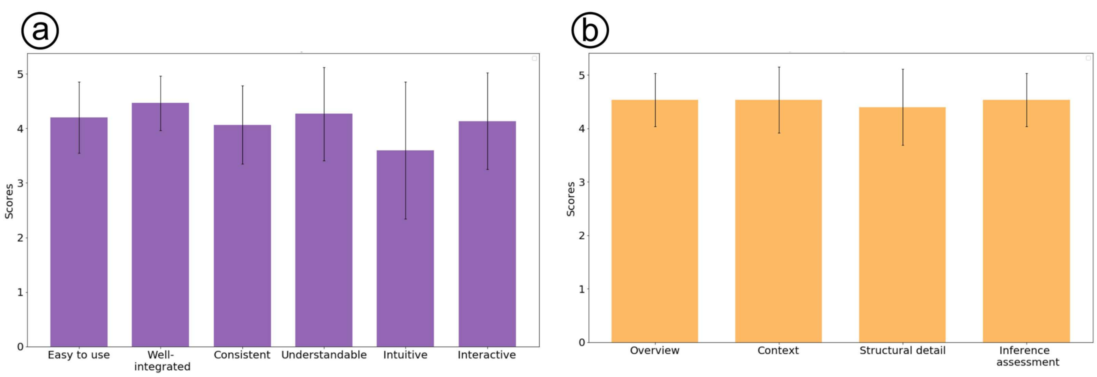

Fig. 12. User study results: usability (a) and effectiveness (b).

information and computer science, political science, public policy, economics, business, and statistics.

Study procedure. The study was conducted using three steps. First, we briefly introduced the background knowledge of policy diffusion, the goal of PolicyFlow, as well as the system functionality for 20 minutes. Second, we provided a tutorial of the system and let users play with it to become familiar with its interface. Third, to help them better understand the system's functionality, we prepared 10 tasks to let users explore the system. For each task, we specified a policy, state, or time period of interest to seek an answer (e.g., Could you identify policies that might have similar cascade sequence to the policy "Film Tax Credits"?). Finally, we asked participants to complete a set of survey questions. The survey includes two types of subjective aspects: (1) *usability*: whether the system is easy to use and intuitive, and (2) *effectiveness*: how the system meets its design requirements and helps users understand policy diffusion. A total of 11 questions, each on a 5-point Likert scale, were asked. We also collected open-ended feedback after the questionnaires.

Results. As shown in Figure 12, the overall result ($M = 4.15$, $SD = 0.81$) indicated that the participants agreed that PolicyFlow has reasonably good usability and has overall achieved its system requirement. For the usability, participants agreed that the system is "easy to use" (M = 4.2, SD = 0.65) and "understandable" (M = 4.27, SD = 0.85). In the open-ended feedback section, most of the participants mentioned that they appreciated the multi-dimensional analysis of policies and adoption cases. Three participants mentioned that the inferred network looked appealing; it intuitively shows the relationship among states and clearly represents the diffusion pattern. The "intuitive" aspect had rated the lowest (M = 3.6, SD = 1.25). Some participants felt that the system's learning curve is high because they need to have background knowledge of policy adoption to understand the concept of network inference.

In terms of effectiveness (how well the system meets its requirements), participants agreed that overall, the system is well integrated (M = 4.6, SD = 0.49). All four system requirements received scores above 4.4, indicating that the system met most of its design requirements. Three users mentioned that the exploration of policies in context was useful, especially the ability to perform categorical and topical selection (filtering) of policies. Four participants additionally suggested that the system could provide more detailed information about state politics (the ideological orientation of a state and who governs the state, etc.).

7 DISCUSSION

We found that overall, `PolicyFlow` was highly appreciated by the expert and non-expert users who participated in our interviews and user study. The overall strengths, as mentioned in both studies, were the following:

- *Capacity of multi-faceted filtering/selection to meet users' heterogeneous needs*: This allows for narrowing down the complicated policy data according to the users' interest in specific political settings. In the case studies and expert interviews, we demonstrated the system's capacity to help connect the diffusion network patterns with particular political contexts through `Policy-Inspection-View` with the concept of expected and deviant patterns. Even though the experts and users we interviewed showed their heterogeneous interest and needs, in most scenarios the multi-faceted filtering was able to meet the users' political interests coming from their heterogeneous expertise and focus of study.
- *Diverse application scenarios throughout the different political settings*: In the expert interviews, we also showed that `PolicyFlow` with our system can benefit political actors in different settings: our system is capable of not only serving state governments but also other political actors across the federal and local levels of the political system such as Congress and grass root groups as potential users. The inferred relationship based on historical policy adoptions enabled them to identify the leading state as policy advocates within specific political context, then any political actors who want to communicate with them proactively take action according to the analysis result in the system. The federal-level policy officers are also expected to have our system monitor state-level policy diffusion.
- *Comparative analysis between the inferred relationship and actual adoptions*: `Policy-Inspection-View` that allows for a comparison between the expected and deviant patterns was also highly appreciated by experts as a useful feature that allows users to see how expected or predictable the adoption sequence of a given policy was from the model and the historical data.

The evaluation process also served to identify the main limitations or suggestions gathered in our study:

- *More flexible policy filtering*: The system currently allows users to perform multi-faceted selection of policies (on spatial, temporal, and topical facets), whereas only a single filter can be applied on each facet. Experts B and C suggested that it would be more useful for users to apply multiple filters simultaneously (e.g., to apply keyword filtering together with topical and subject filters, or other criteria). This will extend `PolicyFlow`'s capability to deal with inter-related aspects in the policy data.
- *Ideological aspect of policies*: During the expert interviews, it was suggested that the ideological aspect of polices is crucial in interpreting the policy diffusion patterns. As Expert B commented, policymakers make their arguments with stances through policies and bills—the ideological aspect provides a possible explanation as to why sets of policies were adopted in similar or different ways.
- *Relationship beyond the state level*: Although our system aims at capturing the policy influence between states, the political system as a whole involves more complex interactions beyond the state level. Expert A stressed that policy is made and adopted more like a ping-pong game where state legislatures and other political actors such as the Supreme Court, federal congresses, and other local and nation-wide organizations negotiate and make modifications to the legislation. Capturing the multi-level interactions among various types of political actors will help further understand the dynamics of real-world political interactions and how they manifest in policy diffusion.

8 CONCLUSION AND FUTURE WORK

In this article, we presented `PolicyFlow`, an interactive framework for exploring diffusion patterns of policies in context. The policy-making process in the public policy sector requires multi-faceted exploration and careful inspection of the observable policy data within a complex context. The visual components in our system not only provide an overview of diffusion patterns over states but also enable contextual filtering on spatial, temporal, and topical aspects. We went a step further by providing analytic components that allows users to interactively inspect the similarity of policies, improving `Policy-Inspection-View` to make it interact with users and focus on a brief range of states and year, and assess the inference networks generated by the network inference algorithm. In the future, we plan to incorporate the experts' suggestions and address the aforementioned limitations, such as providing a more flexible policy filtering mechanism, providing a ideological view of policies, and integrating multiple levels of interactions among different types of political actors, to help users better understand policy diffusion in the highly complex context.

ACKNOWLEDGMENTS

We thank the anonymous referees for their useful suggestions.

REFERENCES

[1] Fernanda Viégas, Martin Wattenberg, Jack Hebert, Geoffrey Borggaard, Alison Cichowlas, Jonathan Feinberg, Jon Orwant, and Christopher Wren. 2013. Google+ ripples: A native visualization of information flow. In *Proceedings of the 22nd International World Wide Web Conference*. 1389–1398.

[2] NYTLabs. 2011. NYTLabs Cascade: How Information Propagates through Social Media. Retrieved May 9, 2018 from http://nytlabs.com/projects/cascade.html.

[3] Guardian Interactive Team. 2016. How Riot Rumours Spread on Twitter. Retrieved May 9, 2018 from https://www.theguardian.com/uk/interactive/2011/dec/07/london-riots-twitter.

[4] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. 2017. A taxonomy and survey of dynamic graph visualization. *Computer Graphics Forum* 36 (2017), 133–159.

[5] Matthew Berger, Katherine McDonough, and Lee M. Seversky. 2017. cite2vec: Citation-driven document exploration via word embeddings. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 691–700.

[6] William D. Berry and Brady Baybeck. 2005. Using geographic information systems to study interstate competition. *American Political Science Review* 99, 4 (2005), 505–519.

[7] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research* 3 (Jan. 2003), 993–1022.

[8] Nan Cao, Yu-Ru Lin, David Gotz, Jimeng Sun, and Huamin Qu. 2011. ChronAtlas: A visualization for dynamic topic exploration. In *Proceedings of the 2011 IEEE Information Visualization Conference (InfoVis'11)*.

[9] Nan Cao, Yu-Ru Lin, Xiaohua Sun, David Lazer, Shixia Liu, and Huamin Qu. 2012. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2649–2658.

[10] Nan Cao, Jimeng Sun, Yu-Ru Lin, David Gotz, Shixia Liu, and Huamin Qu. 2010. FacetAtlas: Multifaceted visualization for rich text corpora. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1172–1181.

[11] Siming Chen, Shuai Chen, Zhenhuang Wang, Jie Liang, Xiaoru Yuan, Nan Cao, and Yadong Wu. 2016. D-Map: Visual analysis of ego-centric information diffusion patterns in social media. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST'16)*. IEEE, Los Alamitos, CA, 41–50.

[12] Isaac Cho, Wewnen Dou, Derek Xiaoyu Wang, Eric Sauda, and William Ribarsky. 2015. VAiRoma: A visual analytics system for making sense of places, times, and events in Roman history. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 210–219.

[13] Weiwei Cui, Shixia Liu, Li Tan, Conglei Shi, Yangqiu Song, Zekai Gao, Huamin Qu, and Xin Tong. 2011. TextFlow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2412–2421.

[14] Weiwei Cui, Shixia Liu, Zhuofeng Wu, and Hao Wei. 2014. How hierarchical topics evolve in large text corpora. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2281–2290.

[15] Bruce A. Desmarais, Jeffrey J. Harden, and Frederick J. Boehmke. 2015. Persistent policy pathways: Inferring diffusion networks in the American states. *American Political Science Review* 109, 2 (2015), 392–406.

[16] Wenwen Dou, Xiaoyu Wang, Drew Skau, William Ribarsky, and Michelle X. Zhou. 2012. LeadLine: Interactive visual analysis of text data through event identification and exploration. In *Proceedings of the 2012 IEEE Conference on Visual Analytics Science and Technology (VAST'12)*. IEEE, Los Alamitos, CA, 93–102.

[17] Wenwen Dou, Li Yu, Xiaoyu Wang, Zhiqiang Ma, and William Ribarsky. 2013. HierarchicalTopics: Visually exploring large text collections using topic hierarchies. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2002–2011.

[18] Glenn Ellison and Drew Fudenberg. 1995. Word-of-mouth communication and social learning. *Quarterly Journal of Economics* 110, 1 (1995), 93–125.

[19] Siwei Fu, Yong Wang, Yi Yang, Qingqing Bi, Fangzhou Guo, and Huamin Qu. 2018. VisForum: A visual analysis system for exploring user groups in online forums. *ACM Transactions on Interactive Intelligent Systems* 8, 1 (Feb. 2018), 1–21. DOI: https://doi.org/10.1145/3162075

[20] Siwei Fu, Jian Zhao, Hao Fei Cheng, Haiyi Zhu, and Jennifer Marlow. 2018. T-Cal: Understanding team conversational data with calendar-based visualization. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 500.

[21] Samah Gad, Waqas Javed, Sohaib Ghani, Niklas Elmqvist, Tom Ewing, Keith N. Hampton, and Naren Ramakrishnan. 2015. ThemeDelta: Dynamic segmentations over temporal topic models. *IEEE Transactions on Visualization and Computer Graphics* 21, 5 (2015), 672–685.

[22] Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. 2012. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data* 5, 4 (2012), 21.

[23] Erin R. Graham, Charles R. Shipan, and Craig Volden. 2013. The diffusion of policy diffusion research in political science. *British Journal of Political Science* 43, 3 (2013), 673–701.

[24] Florian Heimerl, Qi Han, Steffen Koch, and Thomas Ertl. 2016. CiteRivers: Visual analytics of citation patterns. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 190–199. DOI: https://doi.org/10.1109/TVCG.2015.2467621

[25] Enamul Hoque and Giuseppe Carenini. 2018. Interactive topic hierarchy revision for exploring a collection of online conversations. *Information Visualization* 18, 3 (2018), 318–338.

[26] Minjeong Kim, Kyeongpil Kang, Deokgun Park, Jaegul Choo, and Niklas Elmqvist. 2016. TopicLens: Efficient multilevel visual topic exploration of large-scale document collections. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 151–160.

[27] Jure Leskovec and Andreas Krause. 2010. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'10)*.

[28] Jie Li, Siming Chen, Gennady Andrienko, and Natalia Andrienko. 2018. Visual exploration of spatial and temporal variations of tweet topic popularity. In *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA'18)*. DOI: https://doi.org/10.2312/eurova.20181105

[29] Quan Li, Ziming Wu, Lingling Yi, Kristanto Sean Njotoprawiro, Huamin Qu, and Xiaojuan Ma. 2020. WeSeer: Visual analysis for better information cascade prediction of WeChat articles. *IEEE Transactions on Visualization and Computer Graphics* 26, 2 (Feb. 2020), 1399–1412.

[30] Yu-Ru Lin, Jimeng Sun, Nan Cao, and Shixia Liu. 2010. Contextour: Contextual contour visual analysis on dynamic multi-relational clustering. In *Proceedings of the 2010 SIAM International Conference on Data Mining*. 418–429.

[31] Shixia Liu, Xiting Wang, Jianfei Chen, Jim Zhu, and Baining Guo. 2014. TopicPanorama: A full picture of relevant topics. In *Proceedings of the 2014 IEEE Conference on Visual Analytics Science and Technology (VAST'14)*. IEEE, Los Alamitos, CA, 183–192.

[32] Mark Newman. 2010. *Networks: An Introduction*. Oxford University Press.

[33] Deokgun Park, Seungyeon Kim, Jurim Lee, Jaegul Choo, Nicholas Diakopoulos, and Niklas Elmqvist. 2017. ConceptVector: Text visual analytics via interactive lexicon building using word embedding. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 361–370.

[34] W. Russell Neuman, Lauren Guggenheim, S. Mo Jang, and Soo Young Bae. 2014. The dynamics of public attention: Agenda-setting theory meets big data. *Journal of Communication* 64, 2 (2014), 193–214.

[35] Yang Shi, Chris Bryan, Sridatt Bhamidipati, Ying Zhao, Yaoxue Zhang, and Kwan-Liu Ma. 2018. MeetingVis: Visual narratives to assist in recalling meeting context and content. *IEEE Transactions on Visualization and Computer Graphics* 24, 6 (2018), 1918–1929.

[36] Guodao Sun, Yingcai Wu, Shixia Liu, Tai-Quan Peng, Jonathan J. H. Zhu, and Ronghua Liang. 2014. EvoRiver: Visual analysis of topic coopetition on social media. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1753–1762.

[37] Tatiana Von Landesberger, Felix Brodkorb, Philipp Roskosch, Natalia Andrienko, Gennady Andrienko, and Andreas Kerren. 2015. MobilityGraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 11–20.

[38] Jack L. Walker. 1969. The diffusion of innovations among the American states. *American Political Science Review* 63, 3 (1969), 880–899.

[39] Xiting Wang, Shixia Liu, Yang Chen, Tai-Quan Peng, Jing Su, Jing Yang, and Baining Guo. 2016. How ideas flow across multiple social groups. In *Proceedings of the 2016 IEEE Conference on Visual Analytics Science and Technology (VAST'16)*. IEEE, Los Alamitos, CA, 51–60. DOI: https://doi.org/10.1109/VAST.2016.7883511

[40] Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. 2014. OpinionFlow: Visual analysis of opinion diffusion on social media. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1763–1772.

[41] Jin Xu, Yubo Tao, Yuyu Yan, and Hai Lin. 2018. VAUT: A visual analytics system of spatiotemporal urban topics in reviews. *Journal of Visualization* 21, 3 (2018), 471–484.

Received April 2019; revised February 2020; accepted February 2020