Argviz: Interactive Visualization of Topic Dynamics in Multi-party Conversations

Viet-An Nguyen
Dept. of Comp. Science and UMIACS
University of Maryland
College Park, MD
vietan@cs.umd.edu

Yuening Hu
Dept. of Comp. Science and UMIACS
University of Maryland
College Park, MD
ynhu@cs.umd.edu

Jordan Boyd-Graber
iSchool and UMIACS
University of Maryland
College Park, MD
jbg@umiacs.umd.edu

Philip Resnik
Department of Linguistics and UMIACS
University of Maryland
College Park, MD
resnik@umd.edu

Abstract

We introduce an efficient, interactive framework—*Argviz*—for experts to analyze the dynamic topical structure of multi-party conversations. Users inject their needs, expertise, and insights into models via iterative topic refinement. The refined topics feed into a segmentation model, whose outputs are shown to users via multiple coordinated views.

1 Introduction

Uncovering the structure of conversations often requires close reading by a human expert to be effective. Political debates are an interesting example: political scientists carefully analyze what gets said in debates to explore how candidates shape the debate's agenda and frame issues or how answers subtly (or not so subtly) shift the conversation by dodging the question that was asked (Rogers and Norton, 2011).

Computational methods can contribute to the analysis of topical dynamics, for example through topic segmentation, dividing a conversation into smaller, topically coherent segments (Purver, 2011); or through identifying and summarizing the topics under discussion (Blei et al., 2003; Blei, 2012). However, the topics uncovered by such methods can be difficult for people to interpret (Chang et al., 2009), and previous visualization frameworks for topic models—e.g., ParallelTopics (Dou et al., 2011), TopicViz (Eisenstein et al., 2012), the Topical Guide,[1] or topic model visualization (Chaney and Blei, 2012)—are not particularly well suited for linearly structured *conversations*.

This paper describes *Argviz,* an integrated, interactive system for analyzing the topical dynamics of multi-party conversations. We bring together previous work on *Interactive Topic Modeling* (ITM) (Hu et al., 2011), which allows users to efficiently inject their needs, expertise, and insights into model building via iterative topic refinement, with *Speaker Identity for Topic Segmentation* (SITS) (Nguyen et al., 2012), a state-of-the-art model for topic segmentation and discovery of topic shifts in conversations. *Argviz*'s interface allows users to quickly grasp the topical flow of the conversation, discern when the topic changes and by whom, and interactively visualize the conversation's details on demand.

[1] http://tg.byu.edu/

2 System Overview

Our overall system consists of three steps: (1) data preprocessing, (2) interactive topic modeling, and (3) conversational topic segmentation and visualization.

Data preprocessing Preprocessing creates bags of words that can be used by models. First, stopwords and low frequency terms are removed from tokenized text. This is then used as the data for topic modeling.

Interactive topic modeling The topic modeling process then discovers—through posterior inference—the topics that best explain the conversational turns. Each of the topics is a multinomial distribution over words, which can be displayed to users along with the association of turns (documents) to these topics.

The result of topic modeling may be imperfect; we give users an opportunity to refine and curate the topics using Interactive Topic Modeling (ITM) (Hu et al., 2011). The feedback from users is encoded in the form of **correlations**: word types that should co-occur in a topic or which should not. As these correlations are incorporated into the model, the topics learned by the model change and are presented

Proceedings of the NAACL HLT 2013 Demonstration Session, pages 36–39,
Atlanta, Georgia, 10-12 June 2013.

again to the user. The process repeats over multiple iterations until the user is satisfied.

In addition, a simple but important part of the interactive user experience is the ability for users to **label** topics, i.e., to identify a "congress" topic that includes "bill", "vote", "representative", etc.

ITM is a web-based application with a HTML and jQuery[2] front end, connected via Ajax and JSON.

Topic segmentation After the user has built interpretable topics, we use SITS—a hierarchical topic model (Nguyen et al., 2012)—to jointly discover the set of topics discussed in a given set of conversations and how these topics change during each conversation. We use the output of ITM to initialize SITS[3] with a high quality user-specific set of topics. The outputs of SITS consist of (1) a set of topics, (2) a distribution over topics for each turn, and (3) a probability associated with each turn indicating how likely the topic of that turn has been shifted.

The outputs of SITS are displayed using *Argviz* (Figure 2). *Argviz* is a web-based application, built using Google Web Toolkit (GWT),[4] which allows users to visualize and manipulate SITS's outputs entirely in their browser after a single server request.

3 *Argviz:* Coordinated Conversational Views

Given the limited screen of a web browser, *Argviz* follows the multiple coordinated views approach (Wang Baldonado et al., 2000; North and Shneiderman, 2000) successfully used in Spotfire (Ahlberg, 1996), Improvise (Weaver, 2004), and SocialAction (Perer and Shneiderman, 2006). *Argviz* supports three main coordinated views: *transcript*, *overview* and *topic*.

Transcript occupies the prime real estate for a close reading. It has a *transcript panel* and a *speaker panel*. The *transcript panel* displays the original transcript. Each conversational turn is numbered and color-coded by speaker. The color associated with each speaker can be customized using the *speaker panel*, which lists all the speakers.

[2] http://jquery.com/

[3] Through per-word topic assignments

[4] https://developers.google.com/web-toolkit/

Overview shows how topics gain and lose prominence during the conversation. SITS's outputs include a topic distribution and a topic shift probability for each turn in the conversation. In *Argviz,* these are represented using a *heatmap* and *topic shift column*.

In the *heatmap*, each turn-specific topic distribution is displayed by a heatmap row (Sopan et al., 2013). There is a cell for each topic, and the color intensity of each cell is proportional to the probability of the corresponding topic of a particular turn. Thus, users can see the topical flow of the conversation through the vertical change in cells' color intensities as the conversation progresses. In addition, the *topic shift column* shows the topic shift probability (inferred by SITS) using color-coded bar charts, helping users discern large topic changes in the conversation. Each row is associated with a turn in the conversation; clicking on one shifts the *transcript view*.

Topic displays the set of topics learned by SITS (primed by ITM), with font-size proportional to the words' topic probabilities. The *selected topic panel* goes into more detail, with bar charts showing the topic-word distribution. For example, in Figure 2, the Foreign Affairs topic in panel E has high probability words "iraq", "afghanistan", "war", etc. in panel F.

4 Demo: Detecting 2008 Debate Dodges

Visitors will have the opportunity to experiment with the process of analyzing the topical dynamics of different multi-party conversations. Multiple datasets will be preprocessed and set up for users to choose and analyze. Examples of datasets that will be available include conversation transcripts from CNN's *Crossfire* program and debates from the 2008 and 2012 U.S. presidential campaigns. For this section, we focus on examples from the 2008 campaign.

Interactive topic refinement After selecting a dataset and a number of topics, the first thing a user can do is to label topics. This will be used later in *Argviz* and helps users build a mental model of what the topics are. For instance, the user may rename the second topic "Foreign Policy".

After inspecting the "Foreign Policy" topic, the user may notice the omission of Iran from the most probable words in the topic. A user can remedy that by adding the words "Iran" and "Iranians" into the

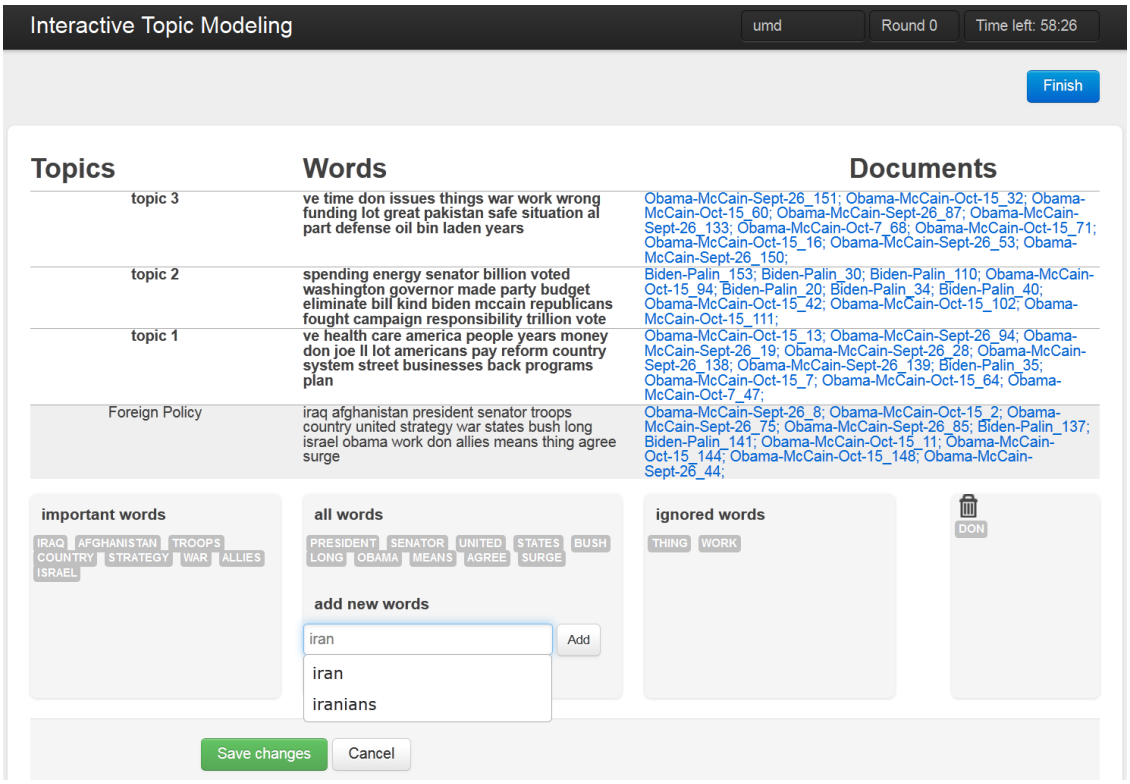

Figure 1: ITM user interface for refining a topic. Users can iteratively put words into different "bins", label topics, and add new words to the topic. Users can also click on the provided links to show related turns for each topic in context.

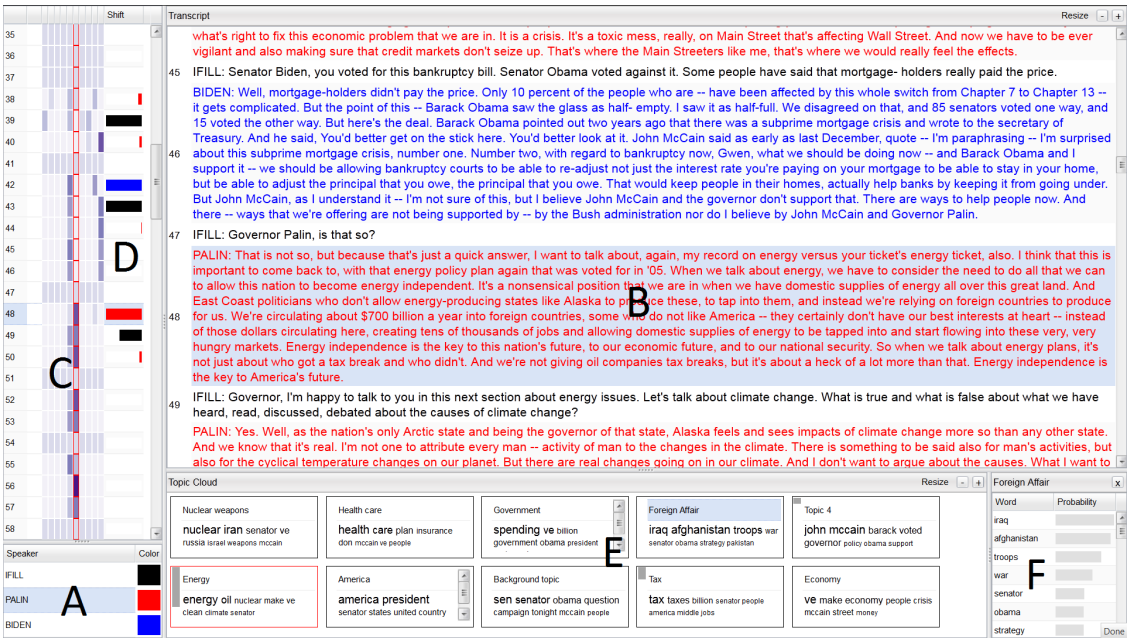

Figure 2: The *Argviz* user interface consists of *speaker panel* (A), *transcript panel* (B), *heatmap* (C), *topic shift column* (D), *topic cloud panel* (E), *selected topic panel* (F).

important words bin (Figure 1). Other bins include *ignored words* for words that should be removed (e.g., "thing" and "work" from this topic) from the topic and *trash* (e.g., "don", which is a stop word).

The user can commit these changes by pressing the *Save changes* button. The back end relearns given the user's feedback. Once users are satisfied with the topic quality, they can click on the *Finish* button to stop updating topics and start running the SITS model, initialized using the final set of refined topics.

Visual analytic of conversations After SITS finishes (which takes just a few moments), users see the dataset's conversations in the *Argviz* interface. Figure 2 shows *Argviz* displaying the 2008 vice presidential debate between Senator Joe Biden and Governor Sarah Palin, moderated by Gwen Ifill.

Users can start exploring the interface from any of the views described in Section 3 to gain insight about the conversation. For example, a user may be interested in seeing how the "Economy" is discussed in the debates. Clicking on a topic in the *topic cloud panel* highlights that column in the *heatmap*. The user can now see where the "Economy" topic is discussed in the debate. Next to the heatmap, the *topic shift column* when debate participants changed the topic. The red bar in turn 48 shows an interaction where Governor Palin dodged a question on the "bankruptcy bill" to discuss her "record on energy". Clicking on this turn shows the interaction in the *transcript view*, allowing a closer reading.

Users might also want to contrast the topics that were discussed before and after the shift. This can be easily done with the coordination between the *heatmap* and the *topic cloud panel*. Clicking on a cell in the *heatmap* will select the corresponding topic to display in the *selected topic panel*. In our example, the topic of the conversation was shifted from "Economy" to "Energy" at turn 48.

5 Conclusion

Argviz is an efficient, interactive framework that allows experts to analyze the dynamic topical structure of multi-party conversations. We are engaged in collaborations with domain experts in political science exploring the application of this framework to political debates, and collaborators in social psychology exploring the analysis of intra- and inter-cultural negotiation dialogues.

References

[Ahlberg, 1996] Ahlberg, C. (1996). Spotfire: an information exploration environment. *SIGMOD*, 25(4):25–29.

[Blei, 2012] Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, 55(4):77–84.

[Blei et al., 2003] Blei, D. M., Ng, A., and Jordan, M. (2003). Latent Dirichlet allocation. *JMLR*, 3.

[Chaney and Blei, 2012] Chaney, A. J.-B. and Blei, D. M. (2012). Visualizing topic models. In *ICWSM*.

[Chang et al., 2009] Chang, J., Boyd-Graber, J., Wang, C., Gerrish, S., and Blei, D. M. (2009). Reading tea leaves: How humans interpret topic models. In *NIPS*.

[Dou et al., 2011] Dou, W., Wang, X., Chang, R., and Ribarsky, W. (2011). ParallelTopics: A probabilistic approach to exploring document collections. In *VAST*.

[Eisenstein et al., 2012] Eisenstein, J., Chau, D. H., Kittur, A., and Xing, E. (2012). TopicViz: interactive topic exploration in document collections. In *CHI*.

[Hu et al., 2011] Hu, Y., Boyd-Graber, J., and Satinoff, B. (2011). Interactive topic modeling. In *ACL*.

[Nguyen et al., 2012] Nguyen, V.-A., Boyd-Graber, J., and Resnik, P. (2012). SITS: A hierarchical nonparametric model using speaker identity for topic segmentation in multiparty conversations. In *ACL*.

[North and Shneiderman, 2000] North, C. and Shneiderman, B. (2000). Snap-together visualization: a user interface for coordinating visualizations via relational schemata. In *AVI*, pages 128–135.

[Perer and Shneiderman, 2006] Perer, A. and Shneiderman, B. (2006). Balancing systematic and flexible exploration of social networks. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):693–700.

[Purver, 2011] Purver, M. (2011). Topic segmentation. In *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*.

[Rogers and Norton, 2011] Rogers, T. and Norton, M. I. (2011). The artful dodger: Answering the wrong question the right way. *Journal of Experimental Psychology: Applied*, 17(2):139–147.

[Sopan et al., 2013] Sopan, A., Freier, M., Taieb-Maimon, M., Plaisant, C., Golbeck, J., and Shneiderman, B. (2013). Exploring data distributions: Visual design and evaluation. *JHCI*, 29(2):77–95.

[Wang Baldonado et al., 2000] Wang Baldonado, M. Q., Woodruff, A., and Kuchinsky, A. (2000). Guidelines for using multiple views in information visualization. In *AVI*, pages 110–119.

[Weaver, 2004] Weaver, C. (2004). Building highly-coordinated visualizations in Improvise. In *INFOVIS*.