

# Visualizing the Flow of Discourse with a Concept Ontology

Baoxu Shi  
University of Notre Dame  
bshi@nd.edu

Tim Weninger  
University of Notre Dame  
tweninge@nd.edu

## ABSTRACT

Understanding and visualizing human discourse has long been a challenging task. Although recent work on argument mining have shown success in classifying the role of various sentences, the task of recognizing concepts and understanding the ways in which they are discussed remains challenging. Given an email thread or a transcript of a group discussion, our task is to extract the relevant concepts and understand how they are referenced and re-referenced throughout the discussion. In the present work, we present a preliminary approach for extracting and visualizing group discourse by adapting Wikipedia's category hierarchy to be an external concept ontology. From a user study, we found that our method achieved better results than 4 strong alternative approaches, and we illustrate our visualization method based on the extracted discourse flows.

### ACM Reference Format:

Baoxu Shi and Tim Weninger. 2018. Visualizing the Flow of Discourse with a Concept Ontology. In *The 2018 Web Conference Companion, April 23-27, 2018, Lyon, France*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3184558.3186943>

## 1 INTRODUCTION

Language has long been one of the most efficient forms of communication between people. Technology that can parse and extract information from these conversations currently exists and operates with reasonable accuracy; however, there is a gap in our ability to understand and visualize these conversational statements. Current and previous work in the analysis of news articles and social posts have demonstrated the ability to extract and quantify written ideas [1, 2, 4]; however, these tools operate over large text corpora, so they are not able to discover concept flows of individual conversations (or documents). Related work in natural language processing aims to extract named entities or important concepts and entities from sentences. Although Named Entity Recognition may be able to extract high-quality entities, which could be viewed as concepts, they are usually limited to a few entity types.

The goal of the present work is different. Here we transform a group conversation into a network over concepts in order to visualize the concept flows so that we might better understand the latent communication patterns and group dynamics. Our key insight is to treat human group conversations as trails

over a graph of concepts. With this perspective, an individual's ideas as expressed through language can be mapped to explicit entities or concepts, and, therefore, a single argument can be treated as a path over the graph of concepts.

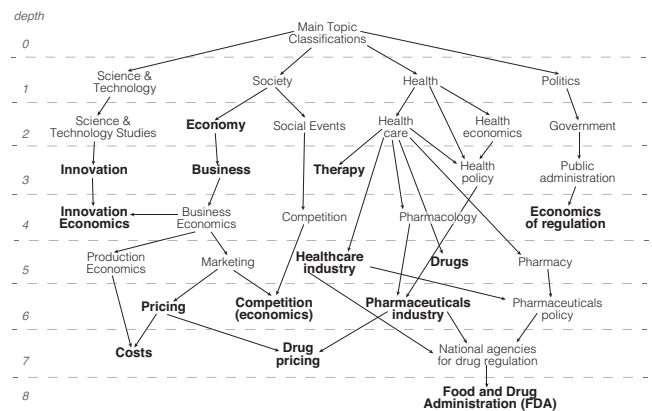
We overcome the limitations mentioned above by distilling a high-quality concept ontology from Wikipedia and using its entity surface forms to detect concepts in human discourse. We then find concept flows by computing sentence similarities using a joint text and concept similarity. The code and data are available at <https://github.com/bxshi/DiscourseVisualization>.

## 2 DISCOURSE GRAPHIFICATION

Here we assume each Wikipedia article represents a unique concept and further treat the categories that an article belongs to as more general concepts. This solution assumes that the Wikipedia category hierarchy is a clean ontology. This is not the case. So the first step is to perform some pre-processing to transform the Wikipedia category hierarchy into an ontology.

We use the October 2017 English dump of all Wikipedia articles and categories. We begin by removing all maintenance, tracking, chronological and list-like pages. This results in a graph rooted at category Main Topic Classifications with 976,163 category nodes, 1,901,706 fine-to-coarse concept edges, and 11,967,618 unique leaf-concepts corresponding to Wikipedia articles. Each article belongs to 4.75 categories on average. A snippet of this ontology is illustrated in Fig. 1.

The next step is to extract concepts from group conversation transcripts and link the sentences to form concept flows over the ontology. To extract concepts from the discourse, we match the surface forms of concepts  $E_i$  from the  $i^{\text{th}}$  sentence  $S_i$  within in transcript  $D$  against the Wikipedia article titles (leaves in



**Figure 1: Illustration of a snippet of the concept graph. Bold texts are extracted concepts recognized as important in the Intelligence<sup>2</sup> debate referenced in Fig. 2.**

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW'18 Companion, April 23-27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3186943>

Seq	Speaker	Sentence
A <sub>1</sub>	J. Donovan	if you think of how many times you've taken <b>antibiotics</b> , casually saving yourself from <b>death by infection</b> .
A <sub>2</sub>	J. Donovan	if you think about <b>vaccines</b> and the way they have shielded you against deadly and disabling <b>disease</b> , if you think about the fact that <b>hiv infection</b> is no longer a <b>death</b> sentence,
B <sub>1</sub>	P. Howard	some <b>patients</b> , particularly those with serious <b>chronic illnesses</b> , are paying too much out of pocket for their <b>medicines</b> , and we need to find a <b>solution</b> for that
C <sub>1</sub>	E. Emanuel	but we don't want those <b>drug prices</b> of <b>\$150,000</b> per year to take those <b>drugs</b>
B <sub>2</sub>	E. Emanuel	it took a <b>disease</b> that was a <b>chronic disease</b> – but blasted off and killed people in six months and basically made people live a very long time with the <b>disease</b> .
C <sub>2</sub>	E. Emanuel	there are multiple <b>drugs</b> out there on the market that are about <b>\$150,000</b> per year, don't <b>cure</b> anyone, ameliorate the <b>disease</b> , but are hugely <b>expensive</b> .
B <sub>3</sub>	L. Reilly	today, it's <b>chronic disease</b> and <b>treating patients</b> with <b>chronic disease</b> that are responsible for 90 percent of all <b>healthcare costs</b> .
C <sub>3</sub>	E. Emanuel	the problem is we have these super high <b>drug prices</b> , <b>150,000</b> , \$300,000 drugs that don't <b>cure</b> anyone, and they're still exorbitantly <b>expensive</b> .
C <sub>4</sub>	P. Howard	micromanage who gets what <b>price</b> and who can do what, and that is one of the biggest things standing in front of us, especially the <b>fda's drug regulations</b> that make it so difficult to <b>innovate</b> and so <b>expensive</b> to <b>innovate</b> .
C <sub>5</sub>	P. Howard	and if we change <b>pricing</b> without changing how we <b>innovate</b> , all we're going to wind up is with fewer <b>drugs</b> and a <b>drug</b> that you don't have for a serious <b>disease</b> is infinitely <b>expensive</b> because you can't buy it.
C <sub>6</sub>	E. Emanuel	we're here to debate <b>drug pricing</b> and <b>drugs</b> outrageously high <b>prices</b> .




Figure 2: A partial example of the Intelligence<sup>2</sup> debate *Blame Big Pharma for Out-of-Control Health Care Costs*. Sentences are color-coded by speaker. Bold texts are extracted concepts shown in Fig. 1. At bottom is a flow diagram over concepts mentioned during the debate.

the Wikipedia concept ontology). Because each concept-leaf is associated with one or more parent and ancestor concepts, we say that each sentence is associated with a concept tree as an induced subgraph  $C_i$  from the concept ontology.

Next we need to link concepts across sentences in the discourse. This requires some notion of concept similarity. There are many ways to do this. We initially tried to adapt the Jaccard coefficient, but this did not work well because it fails to consider the concept granularity and instead treats all concepts, regardless their position in the ontology, equally. To properly weight the concepts, we apply TF-IDF weighting to the extracted concepts by treating them as “words” and the sentences as “documents”. We further define the concept feature vector  $V_i$  of the  $i^{\text{th}}$  sentence as

$$V_i = \left\{ \mathbb{I}(c_k \in C_i) \times \left( 1 + \log \frac{N}{\sum_{j=1}^N \mathbb{I}(c_k \in C_j)} \right) \mid k \in \{1 \dots m\} \right\}, \quad (1)$$

in which  $m$  and  $N$  are the number of concepts and sentences respectively. We can get the word feature vector  $U_i$  using the same method. We now define the sentence similarity as the combination of the word and concept cosine similarities:

$$\text{sim}(S_i, S_j) = \theta(V_i, V_j) + \theta(U_i, U_j). \quad (2)$$

Using Eq. 2 we can now construct concept flows by linking similar sentences and highlighting important concepts in the sentences. For each sentence  $S_i \in \mathcal{D}$ , we find the most similar sentence  $S_j$  in which  $i < j$  and illustrate the concept relationships using a concept network as shown in Fig. 2.

### 3 EXPERIMENTS

We performed a user study to evaluate how well this model captures sentence-level semantic similarities. We compared our model with the results of TopicFlow (LDA) [3], word overlap, averaged sentence embeddings from GloVe, and a text-only version of our model using only  $\theta(U_i, U_j)$  from Eq. 2. Our dataset consisted of four debates from intelligence<sup>2</sup> covering Politics, Health, Science and Economics. We randomly selected 20 sentences from each debate and used the methods mentioned above to find the most similar sentence for each selected sentence. Then for each sentence pair, we asked 10 human annotators to rank the similarity on a 0 to 4 Likert Scale. The results are in Fig. 3 with 95% confidence intervals.




Figure 3: Annotated sentence semantic similarity scores.

The proposed method (Concept in Fig. 3) finds more coherent sentence pairs compared to others. We believe this is because it can better distinguish concept-level similarity. For example, TF-IDF returns *So what is wrong with that argument?* as the most similar sentence to *So, what is wrong with the FDA...*, which ignores FDA, whereas our model returns *Look, the FDA is the biggest barrier here.* instead. LDA performs poorly because the corpus size is limited. Another finding is that averaged word embeddings perform the same as word overlap.

### 4 CONCLUSIONS

We describe a Wikipedia-based concept ontology and a method to find semantically similar sentences. We further present a preliminary visualization using the proposed method to discover concept flows in debates. As for future work, we will employ entity disambiguation into this model to improve the entity detection accuracy, create an interactive visualization tool, and investigate how to model concept shifts in discourse.

### 5 ACKNOWLEDGEMENT

This work is sponsored by the Army Research Office under contract W911NF-17-1-0448.

### REFERENCES

- [1] Mauro Dragoni, Célia Da Costa Pereira, Andrea GB Tettamanzi, and Serena Villata. 2016. Smack: An argumentation framework for opinion mining. In *IJCAI*. IJCAI/AAAI Press, 4242–4243.
- [2] Vinodh Krishnan and Jacob Eisenstein. 2016. Nonparametric Bayesian Storyline Detection from Microtexts. In *EMNLP Workshop*.
- [3] Sana Malik, Alison Smith, Timothy Hawes, Panagis Papadatos, Jianyu Li, Cody Dunne, and Ben Shneiderman. 2013. TopicFlow: visualizing topic alignment of twitter data over time. In *ASONAM*. ACM, 720–726.
- [4] Dafna Shahaf, Carlos Guestrin, and Eric Horvitz. 2012. Trains of thought: Generating information maps. In *WWW*. ACM, 899–908.