

# Multiple Narrative Disentanglement

## ft. identification of narrative threads in *Infinite Jest*



Maggie Engler<sup>1</sup>, Brett Harvey<sup>2,3</sup>

<sup>1</sup> Department of Electrical Engineering, Stanford University, 350 Serra Mall, Stanford, CA 94305

<sup>2</sup> Department of Physics, Stanford University, 382 Via Pueblo, Stanford, CA 94305

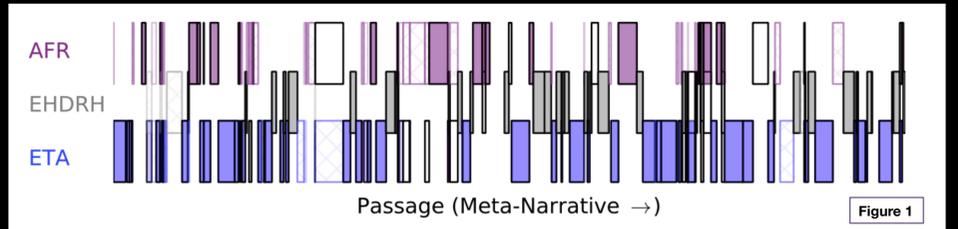
<sup>3</sup> Institute for Computational and Mathematical Engineering, Stanford University, 475 Via Ortega, Stanford, CA 94305



### Background

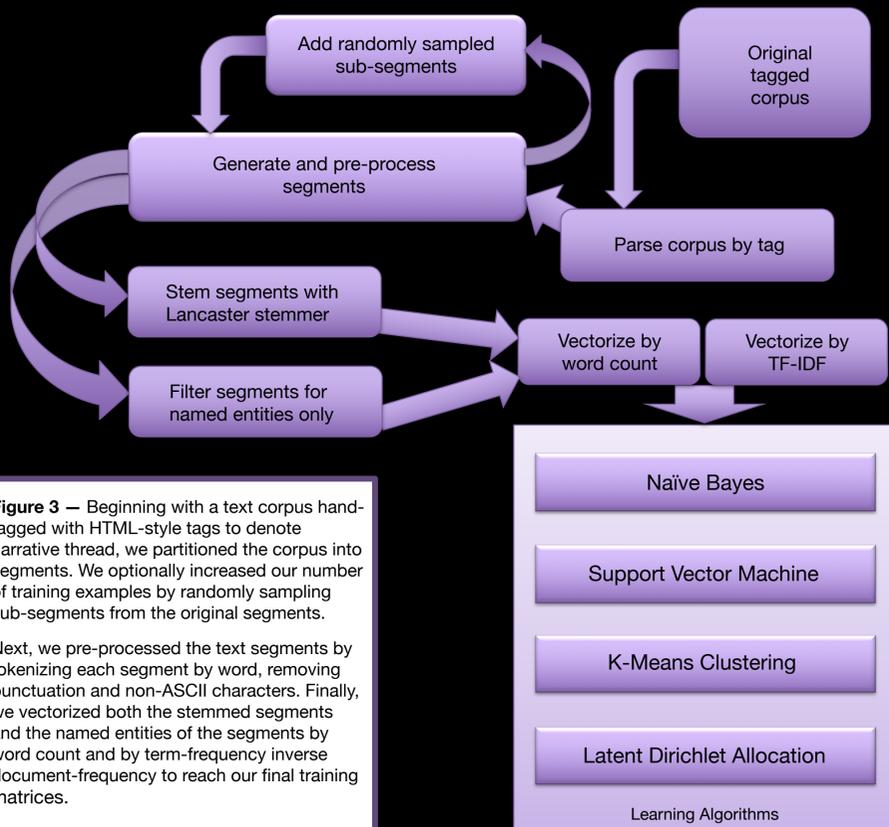
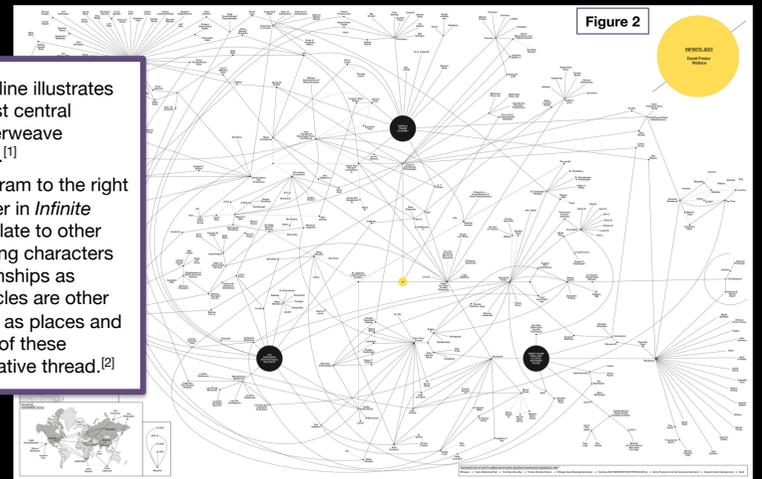
Complex texts usually have multiple narrative threads interwoven rather than a standard linear narrative structure. The task of multiple narrative disentanglement is performed by segmenting texts and reconstructing the separate narrative threads. It can be used for automatic plot summarization, thematic reorganization of a chronological text, and other literary analysis.

We explore several different approaches to the task, drawing from both supervised and unsupervised learning, and demonstrate them on *Infinite Jest*, the nearly-1100 page magnum opus of the late David Foster Wallace, famous for its size and complexity.



**Figure 1** — The timeline illustrates how three of the most central narrative threads interweave throughout the novel.<sup>[1]</sup>

**Figure 2** — The diagram to the right shows each character in *Infinite Jest* and how they relate to other characters by denoting characters as nodes and relationships as edges. The black circles are other named entities, such as places and organizations. Many of these correspond to a narrative thread.<sup>[2]</sup>



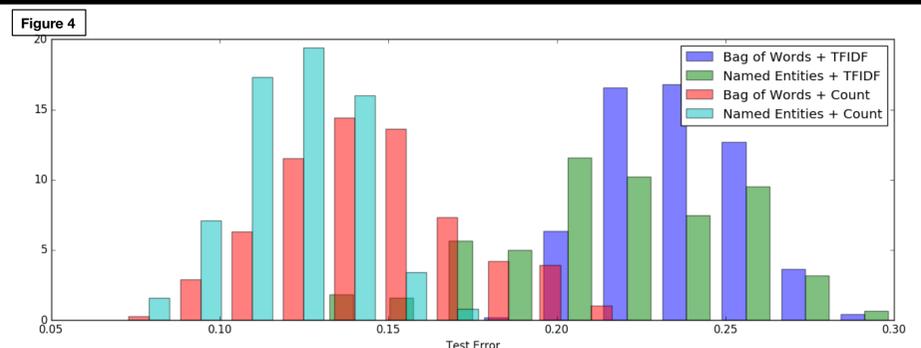
**Figure 3** — Beginning with a text corpus hand-tagged with HTML-style tags to denote narrative thread, we partitioned the corpus into segments. We optionally increased our number of training examples by randomly sampling sub-segments from the original segments.

Next, we pre-processed the text segments by tokenizing each segment by word, removing punctuation and non-ASCII characters. Finally, we vectorized both the stemmed segments and the named entities of the segments by word count and by term-frequency inverse document-frequency to reach our final training matrices.

### Text Feature Extraction

Our various methods for feature extraction are shown at left in Figure 3. Below, in Figure 4, we present several overlaid histograms of test errors for the four different methods of feature extraction indicated in the legend.

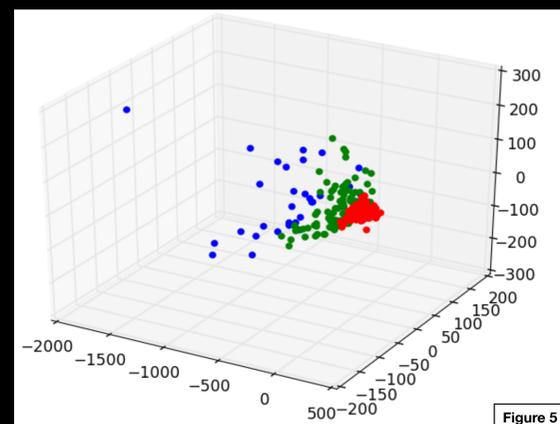
The bimodal distribution shows the relative improvement of a word-frequency method over a term-frequency/inverse document-frequency (TFIDF) method. We had assumed that TFIDF would produce the better test error, but since our results show otherwise, it is possible that segment length is not independent of the particular narrative thread. Moreover, the small difference in peaks in the two modes of the distribution show how using named entity extraction improved our test error.



**Figure 4** — Each histogram shows the test error over 250 trials using different types of feature extraction according to the legend and a Naïve Bayes classifier.

**Figure 5** — To the right is a graph of the three clusters that the k-means algorithm determined, using a bag-of-words model for feature extraction. We used Primary Component Analysis (PCA), with the axes representing the first through third primary components.

**References**  
<sup>[1]</sup> Wallace, Byron C. "Multiple Narrative Disentanglement: Unraveling *Infinite Jest*." *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2012): 1-10.  
<sup>[2]</sup> Potts, Sam. "Infinite Jest Diagram." *Sam Potts, Inc.*, 30 Mar. 2015. [www.sampottsinc.com/ij/](http://www.sampottsinc.com/ij/). Accessed 12 Dec. 2016.



**Figure 5**

### Narrative Models

In addition to classifying the segments into their narrative thread with supervised learning algorithms, we also explored the data with two unsupervised methods.

We used Latent Dirichlet Allocation (LDA) to associate topics with groups of segments. Performing LDA on the unfiltered segments yields groups that combine narratives around female characters and those around male characters, due to the prevalence of gendered pronouns. When we performed LDA on only the named entities, the top words of each group of entities corresponded most often to characters and places who are part of the same larger narrative thread. K-means clustering showed varying degrees of similarity between clusters (Fig. 5).