Multiple Narrative Disentanglement: Unraveling *Infinite Jest*

Byron C. Wallace
Tufts University and Tufts Medical Center
Boston, MA
byron.wallace@gmail.com

Abstract

Many works (of both fiction and non-fiction) span multiple, intersecting narratives, each of which constitutes a story in its own right. In this work I introduce the task of *multiple narrative disentanglement* (MND), in which the aim is to tease these narratives apart by assigning *passages* from a text to the sub-narratives to which they belong. The motivating example I use is David Foster Wallace's fictional text *Infinite Jest*. I selected this book because it contains multiple, interweaving narratives within its sprawling 1,000-plus pages. I propose and evaluate a novel unsupervised approach to MND that is motivated by the theory of *narratology*. This method achieves strong empirical results, successfully disentangling the threads in *Infinite Jest* and significantly outperforming baseline strategies in doing so.

1 Introduction

Both fictional and non-fictional texts often comprise multiple, intersecting and inter-related narrative arcs. This work considers the task of identifying the (sub-)narratives latent within a narrative text and the set of passages that comprise them. As a motivating example, I consider David Foster Wallace's opus *Infinite Jest* (Wallace, 1996),[1] which contains several disparate sub-narratives interleaved throughout its voluminous (meta-)story. By sub-narrative I mean, loosely, that these threads constitute their own independent stories, coherent on their own (i.e., without the broader context of the overarching narrative). I refer to the task of identifying these independent threads and untangling them from one another as *multiple narrative disentanglement* (MND).

The task is of theoretical interest because disentanglement is a necessary pre-requisite to making sense of narrative texts, an interesting direction in NLP that has received an increasing amount of attention (Elson et al., 2010; Elson and McKeown, 2010; Celikyilmaz et al., 2010; Chambers and Jurafsky, 2008; Chambers and Jurafsky, 2009). Recognizing the (main) narrative threads comprising a work provides a context for interpreting the text. Disentanglement may thus be viewed as the first step in a literary processing 'pipeline'. Identifying threads and assigning them to passages may help in automatic plot summarization, social network construction and other literary analysis tasks. Computational approaches to literature look to make narrative sense of unstructured text, i.e., construct models that relate characters and events chronologically: disentanglement is at the heart of this re-construction.

But MND is also potentially of more pragmatic import: disentanglement may be useful for identifying and extracting disparate threads in, e.g., a newsmagazine article that covers multiple (related) stories.[2] Consider an article covering a political race. It would likely contain multiple sub-narratives (the story of one candidate's rise and fall, a scandal in a political party, etc.) that may be of interest independently of the particular race at hand. Narrative dis-

[1] No relation.

[2] While *narrative* colloquially tends to refer to fictional texts, the narrative voice is also frequently used in non-fictional contexts (Bal, 1997).

2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1–10,
Montréal, Canada, June 3-8, 2012.

entanglement thus has applications outside of computational methods for fiction.

In this work, I treat MND as an unsupervised learning task. Given a block of narrative text, the aim is to identify the top k sub-narratives therein, and then to extract the passages comprising them. The proposed task is similar in spirit to the problem of *chat disentanglement* (Elsner and Charniak, 2010), in which the aim is to assign each utterance in a chat transcription to an associated conversational thread. Indeed, the main objective is the same: disentangle fragments of a monolithic text into chronologically ordered, independently coherent ‘threads’. Despite their similarities, however, narrative disentanglement is a qualitatively different task than chat disentanglement, as I highlight in Section 3.

I take inspiration from the literary community, which has studied the theoretical underpinnings of the narrative form at length (Prince, 1982; Prince, 2003; Abbott, 2008). I rely especially on the seminal work of Bal (1997), *Narratology*, which provides a comprehensive theoretical framework for treating narratives. This narratological theory motivates my strategy of *narrative modeling*, in which I first extract the entities in each passage of a text. I then uncover the latent narrative compositions of these passages by performing latent Dirichlet allocation (LDA) (Blei et al., 2003) over the extracted entities.

The main contributions of this work are as follows. First, I introduce the task of multiple narrative disentanglement (MND). Second, motivated by the theory of narratology (Section 2) I propose a novel, unsupervised method for this task (Section 5) and demonstrate its superiority over baseline strategies empirically (Section 6). Finally, I make available a corpus for this task: the text of *Infinite Jest* manually annotated with narrative tags (Section 4).

2 Narratology

I now introduce some useful definitions and concepts (Table 1) central to the theory of narratology (Bal, 1997). These constructs motivate my approach to the task of disentanglement.

Actor	an agent that performs actions. Actors are not necessarily persons.
Fabula	a series of logically and chronologically related events that are caused or experienced by actors.
Story	an instantiation of a fabula, told in a particular style (a story *tells* a fabula). Stories are not necessarily told in chronological order.
Focalizer	a special actor from whose point of view the story is told.

Table 1: A small glossary of narratology.

These definitions imply that the observed narrative text has been generated with respect to some number of latent *fabulas*. A story is a particular telling of an underlying fabula, i.e., a sequence of events involving actors. Figure 1 schematizes the relationships between the above constructs. The dotted line between author and fabula implies that authors *sometimes* generate the fabula, sometimes not. In particular, an author may re-tell a widely known fabula (e.g., *Hamlet*); perhaps from a different perspective. Consider, for example, the play *Rosencrantz and Guildenstern are Dead* (Stoppard, 1967), a narrative that re-tells the fabula of *Hamlet* from the perspective of the titular characters (both of whom play a minor part in *Hamlet* itself). From a narratological view, this story is an instantiation of the *Hamlet* fabula imbued with novel *aspects* (e.g., the focalizers in this telling are Rosencrantz and Guildenstern, rather than Hamlet). In non-fictional works the fabula corresponds to the actual event sequence as it happened, and thus is not invented by the author (save for cases of outright fabrication).

Fabulas are essentially actor-driven. Further, actors tend to occupy particular places, and indeed Bal (1997) highlights locations as one of the defining elements of fabulas. Given these observations, it thus seems fruitful to attempt to identify the agents and locations (or *entities*) in each passage of a text as a first step toward disentanglement. I will return to this intuition when I present the *narrative modeling* method in Section 5. First, I place the present work in context by relating it to existing work on mining literature and chat disentanglement.

3 Relationship to Existing Work

Most similar to MND is the task of *chat disentanglement* (Shen et al., 2006; Elsner and Charniak, 2010; Elsner and Charniak, 2011), wherein utterances (perhaps overheard at a cocktail party) are to

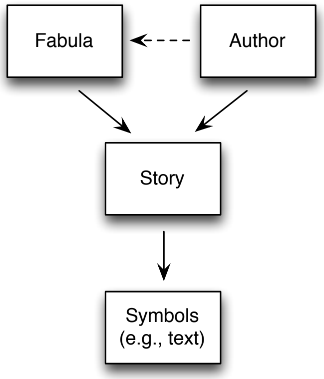

Figure 1: A schematic of the narratology theory. The dotted line between author and fabula implies that when generating a narrative text, an author may invent a fabula, or may draw upon an existing one. Together, the author and fabula jointly give rise to the story, which is communicated via the text.

be assigned to conversational threads. There are, however, important differences between these two tasks. Notably, utterances in a chat belong to a single discussion thread, motivating 'hard' assignments of utterances to threads, e.g., using graph-partitioning (Elsner and Charniak, 2010) or k-means like approaches (Shen et al., 2006). Narratives, however, often intersect: a single passage may belong to multiple narrative threads. This motivates soft, probabilistic assignments of passages to threads. Moreover, narratives are inherently hierarchical. The latter two observations suggest that probabilistic generative models are appropriate for MND.

There has also been recent interesting related work in the unsupervised induction of *narrative schemas* (Chambers and Jurafsky, 2008; Chambers and Jurafsky, 2009). In this work, the authors proposed the task of (automatically) discovering the events comprising a narrative chain. Here *narrative event chains* were defined by Chambers and Jurafsky (2008) as partially ordered sets of events involving the same protagonist. While similar in that these works attempt to make sense of narrative texts, the task at hand is quite different.

In particular, narrative schema induction presupposes a single narrative thread. Indeed, the authors explicitly make the assumption that a single protagonist participates in all of the events forming a narrative chain. Thus the discovered chains describe actions experienced by the protagonist localized within a particular narrative structure. By contrast, in this work I treat narrative texts as instantiations of fabulas, in line with Bal (1997). Fabulas can be viewed as distributions over characters, events and other entities; this conceptualization of what constitutes a narrative is broader than Chambers and Jurafsky (2008). inducing narrative schemas (Chambers and Jurafsky, 2009) may be viewed as a possible next step in a narrative induction pipeline, *subsequent* to disentangling the text comprising individual narrative threads. Indeed, the latter task might be viewed as attempting to automatically re-construct the fabula latent in a specific narrative thread.

Elsewhere, Elson et al. (2010) proposed a method for extracting social networks from literary texts. Their method relies on dialogue detection. This is used to construct a graph representing social interactions, in which an edge connecting two characters implies that they have interacted at least once; the weight of the edge encodes the frequency of their interactions. Their method is a pipelined process comprising three steps: character identification, speech attribution and, finally, graph construction. Their results from the application of this method to a large collection of novels called into question a long-held literary hypothesis: namely that there is an inverse correlation between the number of characters in a novel and the amount of dialogue it contains (Moretti, 2005) (it seems there is not). By answering a literary question empirically, their work demonstrates the power of computational methods for literature analysis.

4 Corpus (*Infinite Jest)*

I introduce a new corpus for the task of multiple narrative disentanglement (MND): David Foster Wallace's novel *Infinite Jest* (Wallace, 1996) that I have manually annotated with narrative tags.[3] *Infinite Jest* is an instructive example for experimenting with MND, as the story moves frequently between a few mostly independent – though ultimately connected and occasionally intersecting – narrative threads.

[3] Available at http://github.com/bwallace/computationaljest. I also note that the text comprises $\sim$100 pages of footnotes, but I did not annotate these.

Annotation, i.e., manually assigning text to one or more narratives, is tricky due primarily to having to make decisions about new thread designation and label granularity.[4] Start with the first. There is an inherent subjectivity in deciding what constitutes a narrative thread. In this work, I was liberal in making this designation, in total assigning 49 unique narrative labels. Most of these tell the story of particular (minor) characters, who are themselves actors in a 'higher-level' narrative – as previously mentioned, narrative structures are inherently hierarchical. This motivates my liberal introduction of narratives: lesser threads are subsumed by their parent narratives, and can thus simply be ignored during analysis if one is uninterested in them. Indeed, this work focuses only on the three main narratives in the text (see below).

Granularity poses another challenge. At what level ought the text be annotated? Should each *sentence* be tagged with associated threads? Each *paragraph*? I let context guide this decision: in some cases tags span a single sentence; more often they span paragraphs. As an example, consider the following example of annotated text, wherein the AFR briefly narrative intersects the story of the ETA (see Table 2).

> <**AFR**>Marathe was charged with this operation's details ... <**ETA**>A direct assault upon the Academy of Tennis itself was impossible. A.F.R.s fear nothing in this hemisphere except tall and steep hillsides. ... </**ETA**></**AFR**>

Here the ellipses spans several paragraphs. Precision probably matters less than context in MND: identifying only sentences that involve a particular sub-narrative, sans context, would probably not be useful. Because the appropriate level of granularity depends on the corpus at hand, the task of segmenting the text into useful chunks is a sub-task of MND. I refer to the segmented pieces of text as *passages* and say that a passage belongs to all of the narrative threads that appear anywhere within it. Hence in the above example, the passage containing this excerpt would be designated as belonging to both the **ETA** and **AFR** threads.

[4] These complexities seem to be inherent to disentanglement tasks in general: Elsner and Charniak (2010) describe analogues issues in the case of chat.

AFR	This is the tale of the *wheelchair assassins*, a Quèbècois terrorist group, and their attempts to seize an original copy of a dangerous film. *Focalizer*: Marathe.
EHDRH	The Ennet House Drug Recovery House (sic). This narrative concerns the going-ons at a drug recovery house. *Focalizer*: Don Gately.
ETA	This narrative follows the students and faculty at the Enfield Tennis Academy. *Focalizer*: Hal.

Table 2: Brief summaries of the main narratives comprising *Infinite Jest*.

narrative	# of passages	prevalence
AFR	30	16%
EHDRH	42	23%
ETA	69	38%

Table 3: Summary statistics for the three main narratives.

Infinite Jest is naturally segmented by breaks, i.e., blank lines in the text which typically indicate some sort of context-shift (functionally, these are like mini-chapters). There are 182 such breaks in the book, demarcating 183 passages. Each of these comprises about 16,000 words and contains an average of 4.6 (out of 49) narratives, according to my annotations.

There are three main narrative threads in *Infinite Jest*, summarized briefly in Table 2.[5] I am not alone in designating these as the central plot-lines in the book.[6] Nearly all of the other threads in the text are subsumed by these (together the three cover 72% of the passages in the book). These three main threads are ideal for evaluating an MND system, for a few reasons. First, they are largely independent of one another, i.e., overlap only occasionally (though they do overlap). Second, they are relatively unambiguous: it is mostly clear when a passage tells a piece of one of these story-lines, and when it does not. These narratives are thus well-defined, providing a minimal-noise dataset for the task of MND. That I am the single annotator of the corpus (and hence inter-annotator agreement cannot be assessed) is unfortunate; the difficulty of finding someone both qualified and willing to annotate the 1000+ page book precluded this possibility. I hope to address

[5] I include these only for interested readers: the descriptions are not technically important for the work here, and one may equivalently substitute 'narrative 1', 'narrative 2', etc.

[6] e.g., http://www.sampottsinc.com/ij/file/IJ_Diagram.pdf.

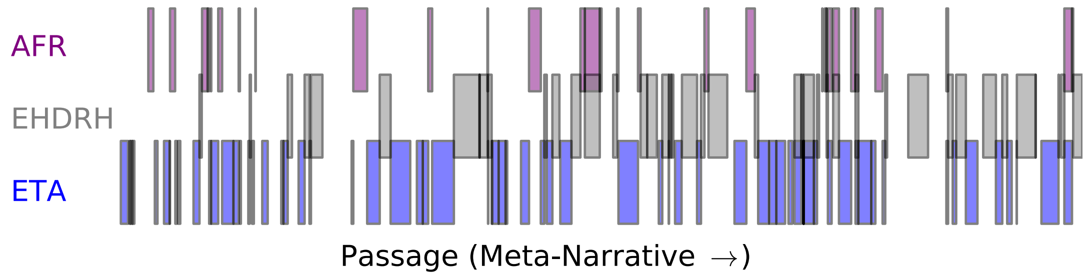

Figure 2: The three main narratives in *Infinite Jest*. A colored box implies that the corresponding narrative is present in the passage at that location in the text; these are scaled relative to the passage length.

this shortcoming in future work.

Figure 2 depicts the location and duration of these sub-narratives within the text. Passages run along the bottom axis. A colored box indicates that the corresponding narrative is present in the passage found at that location in the book. Passages are normalized by their length: a wide box implies a long passage. The aim of MND, then, is to automatically infer this structure from the narrative text.

5 Narrative Modeling for Multiple Narrative Disentanglement

The proposed method is motivated by the theory of narratology (Bal, 1997), reviewed in Section 2. Specifically I assume that passages are mixtures of different narratives with associated underlying fabulas. Fabulas, in turn, are viewed as distributions over *entities*. Entities are typically actors, but may also be locations, etc.; they are what fabulas are *about*. The idea is to infer from the observed passages the probable latent fabulas.

This is a generative view of narrative texts, which lends itself naturally to a topic-modeling approach (Steyvers and Griffiths, 2007). Further, this generative vantage allows one to exploit the machinery of latent Dirichelet allocation (LDA) (Blei et al., 2003). LDA is a generative model for texts (and discrete data, in general) in which it is assumed that each document in a corpus reflects a mixture of (latent) topics. The words in the text are thus assumed to be generated by these topics: topics are multinomials over words. Graphically, this model is depicted by Figure 3. All of the parameters in this model must be estimated; only the words in documents are observed. To uncover the topic mixtures latent in documents, standard inference procedures can be used for parameter estimation (Jordan et al., 1999).

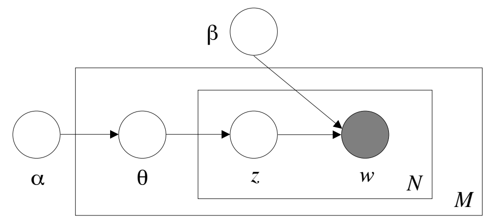

Figure 3: The graphical model of latent Dirichlet allocation (LDA; Figure from Blei et al. (2003)). Θ parameterizes the multinomial governing topics, i.e., zs. The observed words w are then assumed to be drawn from a multinomial conditioned on z. Here the plates denote that there are N (observed) words and M topics.

I propose the following approach for MND, which I will refer to as *narrative modeling*. (This pipeline is also described by Figure 4).

1. Segment the raw text into *passages*. It is at the level of this unit that narratives will be assigned: if a given narrative tag is anywhere in a passage, that passage is deemed as being a part of said narrative.[7] In many cases (including the present one) this step will be relatively trivial; e.g., segmenting the text into chapters or paragraphs.

2. (Automatically) extract from each of these segments *named entities*. The idea is that these include the primary players in the respective narratives, i.e., important actors and locations.

3. Perform latent Dirichelet analysis (LDA) over the entities extracted in (2). When this topic mod-

[7] This is analogous to a multi-label scenario.

eling is performed over the entities, rather than the text, I shall refer to it as *narrative modeling*.

As mentioned above, Step (1) will be task-specific: what constitutes a passage is inherently subjective. In many cases, however, the text will lend itself to a 'natural' segmenting, e.g., at the chapter-level. Standard statistical techniques for named entity recognition (NER) can be used for Step (2) (McCallum and Li, 2003).

Algorithm 1 The story of LDA over extracted entities for multiple narrative disentanglement.

Draw a mixture of narrative threads $\theta \sim Dir(\alpha)$
for each entity in the passage e_i **do**
  Draw a narrative thread $t_i \sim Multinomial(\theta)$
  Draw e_i from $p(e_i|t_i)$
end for

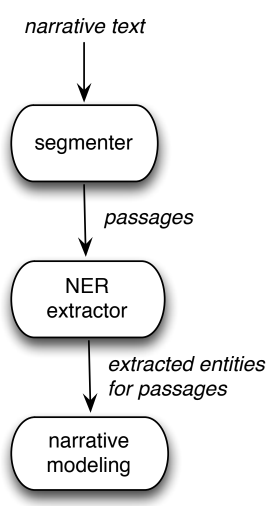

Figure 4: The MND pipeline.

For the narrative modeling Step (3), I use LDA (Blei et al., 2003); the generative story for narrative modeling is told by Algorithm 1.[8] This squares with the narratological view: entities are observed in the text with probability proportional to their likelihood of being drawn from the corresponding latent fabulas (which we are attempting to recover). Focusing on these entities, rather than the raw text, is crucial if one is to be compatible with the narratological view. The text is merely a particular telling of the underlying fabula, made noisy by story specific aspects; extracting entities from the passages effectively removes this noise, allowing the model to operate over a space more closely tied to the fabulas. In the following section, I demonstrate that this shift to the entity-space substantially boosts MND performance.

[8] Liu and Liu (2008) have also proposed topic models over NEs, though in a very different context.

The aim is to uncover the top k most salient narrative threads in a text, where k is a user-provided parameter. Indeed one *must* specify the number of threads he or she is interested in identifying (and disentangling), because because, due to the hierarchical nature of narratives, there is no single 'right number' of them. Consider that the input block of text constitutes a perfectly legitimate (meta-)narrative on its own, for example. A related issue that must be addressed is that of deciding when to assign a passage to multiple threads. That is, given the (estimated) narrative mixtures for each passage as an input, to which (if any) narrative threads ought this passage be assigned?

My approach to this is two-fold. First, I set a threshold probability α such that a passage p_i cannot be assigned to a narrative thread t if the estimated mixture component is $\leq \alpha$. I use $\alpha = 1/k$, as this value implies that the passage is dominated by other threads (in the case that all k threads contribute equally to a passage, the corresponding mixture elements would all be $1/k$). Second, I enforce a constraint that in order to be assigned to the narrative t, a passage must contain at least one of the top l entities involved in t (according to the narrative model). This constraint encodes the intuition that the main actors (and locations) that constitute a given fabula are (extremely) likely to be present in any given passage in which it is latent. I set $l = 100$, reflecting intuition. These were the first values I used for both of these parameters; I did not tune them to the corpus at hand. I did, however, experiment with other values after the primary analysis to assess sensitivity. The proposed algorithm is not terribly sensitive to either parameter, though both exert influence in the expected directions: increasing α decreases recall, as passages are less likely to be assigned to narratives. Decreasing l has a similar effect, but does not substantially impact performance unless extreme values are used.[9]

5.1 Focalizer Detection

Recall that the focalizer of a narrative is the agent responsible for perception: it is from their point of view that the story is told (Bal, 1997). One can easily exploit the narrative modeling method above to

[9] Fewer than 10 or more than 500, for example.

automatically identify the (main) focalizer of the uncovered narratives.[10] To this end, I simply identify the highest ranking entity from each narrative that has also been labeled as a 'person' (as opposed, e.g., to an 'organization').

6 Empirical Results

I now present experimental results over the *Infinite Jest* corpus, described in Section 4. The task here is to uncover the three main narratives in the text, depicted in Figure 2. To implement the proposed narrative modeling method (Section 5), I first chunked the text into passages, delineated in *Jest* by breaks in the text. I performed entity extraction over these passages using the NLTK toolkit (Bird et al., 2009). I then performed LDA via Mallet (McCallum, 2002) to estimate the narrative mixture components of each passage.

$$recall = TP/(TP + FN) \quad (1)$$

$$precision = TP/(TP + FP) \quad (2)$$

$$F = 2 \cdot \frac{precision \cdot recall}{precision + recall} \quad (3)$$

I compare the narrative modeling approach presented in the preceding section to three baselines. The simplest of these, **round-robin** and **all-same** are similar to the baselines used for chat disentanglement (Elsner and Charniak, 2010). Respectively, these strategies designate each passage as: belonging to the next narrative in a given sequence ('narrative 1', 'narrative 2', 'narrative 3'), and, belonging to the majority narrative. In both cases I show the best result attainable using the method: thus in the case of the former, I report the best scoring results from all 3! possible thread sequences (with respect to macro-averaged F-score) and in the latter case I use the true majority narrative.

I also evaluate a simple topic-modeling baseline, which is the same as narrative modeling, except that: 1) LDA is performed over the full-text (rather than the extracted entities) and, 2) there is no constraint enforcing that passages reference an entity associated with the assigned narrative. I evaluate results with respect to per-narrative recall, precision and F-score (Equations 1-3) (where TP=true positive, FN=false negative, etc.). I also consider micro- and macro-averages of these.

To calculate the micro-average, one considers each passage at a time by counting up the TPs, FPs, TNs and FNs therein for each narrative under consideration (w.r.t. the model being evaluated). The micro-average is then calculated using these tallied counts. Note that in this case certain narratives may contribute more to the overall result than others, e.g. those that are common. By contrast, to calculate the macro-average, one considers each narrative in turn and calculates the average of the metrics of interest (recall, precision) w.r.t. this narrative over all passages. An average is then taken over these mean performances. This captures the average performance of a model over all of the narratives, irrespective of their prevalence; in this case, each thread contributes equally to the overall result. Finally, note that none of the methods explicitly labels the narratives they uncover: this assignment can be made by simply matching the returned narratives to the thread labels (e.g., **ETA**) that maximize performance. This labeling is strictly aesthetic; the aim is to recover the latent narrative threads in text, not to label them.

Table 4 presents the main empirical results. Neither of the simple baseline methods (**round-robin** and **all-same**) performed very well. Both cases, for example, completely failed to identify the EHDRH thread (though this is hardly surprisingly in the **all-same** case, which identifies only one thread by definition). The macro-averaged precisions and F-measures are thus undefined in these cases (these give rise to a denominator of 0). With respect to micro-averaged performance, **all-same** achieves a substantially higher F-score than **round-robin** here, though in general this will be contingent on how dominated the text is by the majority thread.

Next consider the two more sophisticated strategies, including the proposed narrative modeling method. Start with the performance of **full-text TM**, i.e., performing standard topic-modeling over the full-text. This method improves considerably on the baselines, achieving a macro-averaged F-score of .545.[11] But the narrative modeling method (Section 5) performs substantially better, boosting the

[10] Technically, there may be multiple focalizers in a narrative, but more often there is only one.

[11] In the full-text case, I evaluated the performance of every possible assignment of topics to threads, and report the best scoring result.

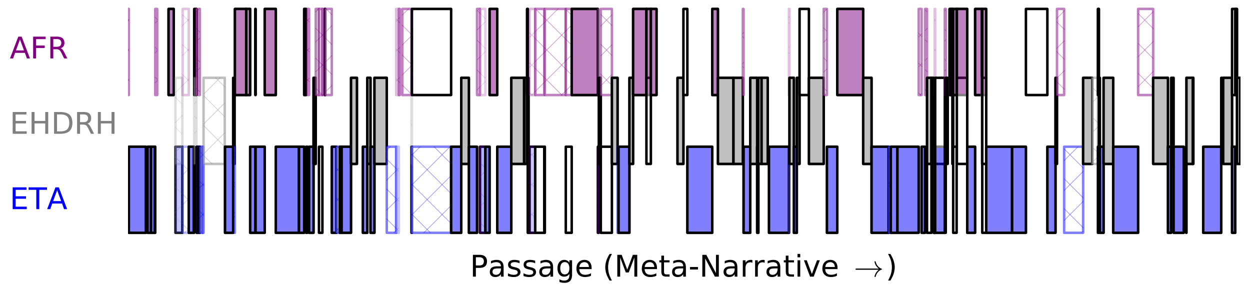

Figure 5: The unsupervised re-construction of the three main narratives using the narrative modeling approach. Hatched boxes denote false-positives (designating a passage as belonging to a narrative when it does not); empty boxes false negatives (failing to assign a passage to narrative to which it belongs).

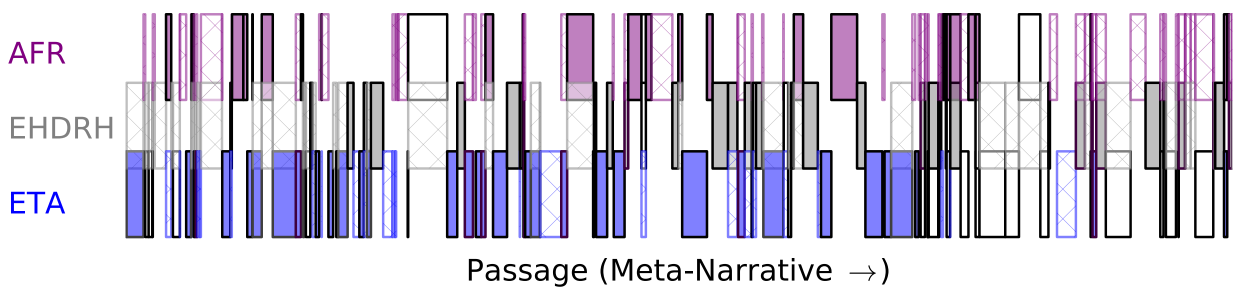

Figure 6: Results using full-text topic modeling (see above caption).

macro-averaged F-score by over 15 points (a percent gain of nearly 30%).

Figures 5 and 6 depict the unsupervised re-construction of the narrative threads using narrative modeling and the full-text topic modeling approach, respectively. Recall that the aim is to re-construct the narratives depicted in Figure 2. In these plots, an empty box represents a false negative (i.e., implies that this passage contained the corresponding narrative but this was not inferred by the model), and a hatched box denotes a false positive (the model assigned the passage to the corresponding narrative, but the passage did not belong to it). One can see that the narrative modeling method (Figure 5) re-constructs the hidden threads much better than does the full-text topic modeling approach (Figure 6). Once can see that the latter method has particular trouble with the EHDRH thread.

I also experimented with the focalizer detection method proposed in Section 5.1. This simple strategy achieved 100% accuracy on the three main narratives, correctly identifying by name each of the corresponding focalizers (see Table 2).

6.1 A More Entangled Thread

The preceding results are positive, insofar as the proposed method substantially improves on baselines and is able to disentangle threads with relatively high fidelity. These results considered the three main narratives that comprise the novel (Figure 2). This is the sort of structure I believe will be most common in narrative disentanglement, as it is likely that one will mostly be interested in extracting coherent threads that are largely independent of one another.

That said, I will next consider a more entangled thread to see if the method handles these well. More specifically, I introduce the narrative **INC**, which relates the story of the Incandenza family. This family is (arguably) the focus of the novel. The story of the Incandenza's overlaps extremely frequently with the three main, mostly independent narratives considered thus far (see Figure 6). This thread is thus difficult from an MND perspective.

I apply the same methods as above to this task, requesting four (rather than three) sub-narratives, i.e., $k = 4$. Results are summarized in Table 5.[12] We ob-

[12] I omit the two baseline strategies due to space constraints;

narrative	round-robin recall	prec.	F	all-same recall	prec.	F	full-text TM recall	prec.	F	narrative modeling recall	prec.	F
AFR	0.433	0.210	0.283	0.000	*undef.*	*undef.*	0.900	0.300	0.450	0.933	0.359	0.519
EHDRH	0.000	*undef.*	*undef.*	0.000	*undef.*	*undef.*	0.786	0.402	0.532	0.929	0.736	0.821
ETA	0.369	0.348	0.393	1.000	0.375	0.545	0.667	0.639	0.653	0.855	0.694	0.766
macro-avg.	0.260	*undef.*	*undef.*	0.333	*undef.*	*undef.*	0.752	0.447	0.545	0.906	0.596	0.702
micro-avg.	0.262	0.300	0.280	0.489	0.375	0.425	0.752	0.434	0.551	0.894	0.583	0.706

Table 4: Empirical results using different strategies for MND. The top three rows correspond to performance for individual narratives; the bottom two provide micro- and macro-averages, which are taken over the individual passages and the narrative-level results, respectively.

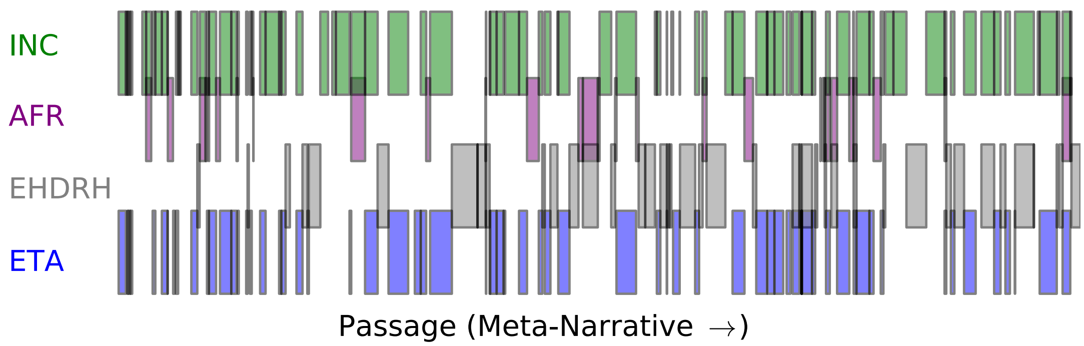

Figure 7: The INC narrative thread (green, top). This narrative is substantially more entangled than the others, i.e., more frequently intersects with the other narratives.

narrative	full-text TM recall	prec.	F	narrative modeling recall	prec.	F
AFR	0.60	0.30	0.40	0.83	0.50	0.63
EHDRH	0.83	0.57	0.67	0.79	0.75	0.77
ETA	0.67	0.69	0.68	0.67	0.89	0.76
INC	0.57	0.46	0.51	0.43	0.75	0.54
macro-avg.	0.67	0.50	0.56	0.68	0.72	0.67
micro-avg.	0.65	0.50	0.57	0.62	0.72	0.67

Table 5: Results when the fourth narrative, more entangled narrative (INC) is added.

serve that the narrative modeling strategy again bests the baseline strategies, achieving a macro-averaged F-score of about 10 points greater than that achieved using the full-text TM method (a ∼20% gain).

Focalizer identification is tricky in this case because there are multiple focalizers. However I note that using the proposed strategy, four members of the Incandenza clan rank in the top five entities associated with this narrative, an encouraging result.[13]

7 Conclusions

I have introduced the task of multiple narrative disentanglement (MND), and provided a new annotated corpus for this task. I proposed a novel method (narrative modeling) for MND that is motivated by the theory of narratology. I demonstrated that this method is able to disentangle the narrative threads comprising *Infinite Jest* and that it substantially outperforms baselines in terms of doing so. I also extended the method to automatically identify narrative focalizers, and showed that it is possible to do so with near-perfect accuracy.

Interesting future directions include exploring *supervised* narrative disentanglement, combining MND with narrative induction (Chambers and Jurafsky, 2009) and applying MND to non-fictional texts.

Acknowledgments

Thanks to Kevin Small and Carla Brodley for suggesting improvements to this work, and to all of the members of the Inman Square Existentialist Book Club for insightful discussions about *Jest*.

both performed worse than the displayed methods.

[13] The fifth top-ranking entity is Joelle, a girl who plays an important part in the family saga.

References

H.P. Abbott. 2008. *The Cambridge introduction to narrative*. Cambridge Univ Pr.

M Bal. 1997. *Narratology: Introduction to the theory of narrative, 3rd ed.* University of Toronto Press.

S. Bird, E. Klein, and E. Loper. 2009. *Natural language processing with Python*. O'Reilly Media.

D.M. Blei, A.Y. Ng, and M.I. Jordan. 2003. Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022.

A. Celikyilmaz, D. Hakkani-Tur, H. He, G. Kondrak, and D. Barbosa. 2010. The actortopic model for extracting social networks in literary narrative. In *NIPS Workshop: Machine Learning for Social Computing*.

N. Chambers and D. Jurafsky. 2008. Unsupervised learning of narrative event chains. *Proceedings of ACL-08: HLT*, pages 789–797.

N. Chambers and D. Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610. Association for Computational Linguistics.

M. Elsner and E. Charniak. 2010. Disentangling chat. *Computational Linguistics*, 36(3):389–409.

M. Elsner and E. Charniak. 2011. Disentangling chat with local coherence models. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 1179–1189. Association for Computational Linguistics.

D.K. Elson and K.R. McKeown. 2010. Automatic attribution of quoted speech in literary narrative. In *Proceedings of AAAI*.

D.K. Elson, N. Dames, and K.R. McKeown. 2010. Extracting social networks from literary fiction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147. Association for Computational Linguistics.

M.I. Jordan, Z. Ghahramani, T.S. Jaakkola, and L.K. Saul. 1999. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233.

Y. Liu and F. Liu. 2008. Unsupervised language model adaptation via topic modeling based on named entity hypotheses. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 4921–4924. IEEE.

A. McCallum and W. Li. 2003. Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003-Volume 4*, pages 188–191. Association for Computational Linguistics.

Andrew Kachites McCallum. 2002. Mallet: A machine learning for language toolkit. http://mallet.cs.umass.edu.

F. Moretti. 2005. *Graphs, Maps, Trees: Abstract models for a literary history*. Verso Books.

G. Prince. 1982. *Narratology: The form and functioning of narrative*. Mouton Berlin.

G. Prince. 2003. *A dictionary of narratology*. University of Nebraska Press.

D. Shen, Q. Yang, J.T. Sun, and Z. Chen. 2006. Thread detection in dynamic text message streams. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 35–42. ACM.

M. Steyvers and T. Griffiths. 2007. Probabilistic topic models. *Handbook of latent semantic analysis*, 427(7):424–440.

T. Stoppard. 1967. *Rosencrantz & Guildenstern are dead: a play in three acts*. Samuel French Trade.

D.F. Wallace. 1996. *Infinite Jest*. Little Brown & Co.