

Validating Literary Theories Using Automatic Social Network Extraction

Prashant Arun Jayannavar
Department of Computer Science
Columbia University
NY, USA
pj2271@columbia.edu

Apoorv Agarwal
Department of Computer Science
Columbia University
NY, USA
apoorv@cs.columbia.edu

Melody Ju
Columbia College
Columbia University
NY, USA
mj2558@columbia.edu

Owen Rambow
Center for Computational Learning Systems
Columbia University
NY, USA
rambow@ccls.columbia.edu

Abstract

In this paper, we investigate whether long-standing literary theories about nineteenth-century British novels can be verified using computational techniques. Elson et al. (2010) previously introduced the task of computationally validating such theories, extracting conversational networks from literary texts. Revisiting their work, we conduct a closer reading of the theories themselves, present a revised and expanded set of hypotheses based on a divergent interpretation of the theories, and widen the scope of networks for validating this expanded set of hypotheses.

1 Introduction

In his book *Graphs, Maps, Trees: Abstract Models for Literary History*, literary scholar Franco Moretti proposes a radical transformation in the study of literature (Moretti, 2005). Advocating a shift from the close reading of individual texts in a traditionally selective literary canon, to the construction of abstract models charting the aesthetic form of entire genres, Moretti imports quantitative tools to the humanities in order to inform what he calls “a more rational literary history.” While Moretti’s work has inspired both support and controversy, this reimaged mode of reading opens a fresh direction from which to approach literary analysis and historiography.

By enabling the “distant reading” of texts on significantly larger scales, advances in Natural Language Processing and applied Machine Learning can be employed to empirically evaluate existing claims

or make new observations over vast bodies of literature. In a seminal example of this undertaking, Elson et al. (2010) attempted to validate an assumption of structural difference between the social worlds of rural and urban novels using social networks extracted from nineteenth-century British novels. Extrapolating from the work of various literary theorists, Elson et al. (2010) hypothesized that nineteenth-century British novels set in urban environments feature numerous characters who share little conversation, while rural novels have fewer characters with more conversations. Using quoted speech attribution, the authors extracted *conversation networks* from 60 novels, which had been manually classified by a scholar of literature as either rural or urban. Elson et al. (2010) concluded that the results of their analysis of conversation networks, which indicated no difference between the social networks of rural and urban novels, invalidated literary hypotheses. However, we believe that Elson et al. (2010) misrepresented the original theories, and that their results actually support rather than contradict the original theories in question.

In this paper, we revisit the work of Elson et al. (2010), presenting a nuanced interpretation of their results through a closer reading of the original theories cited. We propose that Elson et al. (2010)’s results actually align with these theories. We then employ a more powerful tool for extracting social networks from texts, which allows us to examine a wider set of hypotheses and thus provide deeper insights into the original theories. Our findings confirm that the setting (rural versus urban) of a novel

in Elson et al. (2010)’s corpus has no effect on its social structure, even when one goes beyond conversations to more general and different notions of interactions. Specifically, we extend the work of Elson et al. (2010) in four significant ways: (1) we extract *interaction networks*, a conceptual generalization of conversation networks; (2) we extract *observation networks*, a new type of network with directed links; (3) we consider unweighted networks in addition to weighted networks; and (4) we investigate the number and size of communities in the extracted networks.

For extracting interaction and observation networks, we use our existing system called SINNET (Agarwal et al., 2010; Agarwal et al., 2013b; Agarwal et al., 2014). In addition to validating a richer set of hypotheses using SINNET, we present an evaluation of the system on the task of automatic social network extraction from literary texts. Our results show that SINNET is effective in extracting interaction networks from a genre quite different from the genre it was trained on, namely news articles.

The paper is organized as follows. In section 2 we revisit the theories postulated by various literary theorists and critics. In Section 3, we present an expanded set of literary hypotheses. Section 4 presents the methodology used by Elson et al. (2010) for validating their literary hypothesis. We use the same methodology for validating our expanded set of literary hypotheses. In Section 5, we give details on the difference between conversation, observation, and interaction networks. We then evaluate SINNET on the data set provided by Elson et al. (2010) (Section 6). We test our hypotheses against the data in Section 7 and conclude with future directions of research in Section 8.

2 Literary Theories

In section 3 of their paper, Elson et al. (2010) present a synthesis of quotations from literary theorists Mikhail Bakhtin (Bakhtin, 1937), Raymond Williams (Williams, 1975), Franco Moretti (Moretti, 1999; Moretti, 2005) and Terry Eagleton (Eagleton, 1996; Eagleton, 2013). Elson et al. (2010) simplify the quotations to derive the following hypotheses:

- EDM1: There is an inverse correlation between the number of dialogues and the number of

characters.

- EDM2: In novels set in urban environments, numerous characters share little conversational interactions. Rural novels, on the other hand, have fewer characters with more conversations.

We argue that the theories themselves are misconstrued and that the results of Elson et al. (2010)’s experiments actually support what theorists imply about the distinction between rural and urban novels as sub-genres of 19th century realist fiction. For instance, Elson et al. (2010) quote Williams (1975) as follows:

Raymond Williams used the term “knowable communities” to describe this [rural] world, in which face-to-face relations of a restricted set of characters are the primary mode of social interaction (Williams, 1975, 166). By contrast, the urban world, in this traditional account, is both larger and more complex.

On re-visiting this quotation in a larger and original context, we note that Williams (1975) actually apply the term “knowable communities” to novels in general, not to settings, and specifically not – as Elson et al. (2010) presume – to any particular setting (rural in this case). Williams (1975) states that “most novels are in some sense knowable communities”, meaning that the novelist “offers to show people and their relationships in essentially knowable and communicable ways.” However, the need or desire to portray some setting in a realistic (“knowable”) way does not automatically entail the ability to do so: evolutions in real-world social milieu may occur independently of the evolutions in novelistic technique that specifically allow such evolutions to be captured in literature.

In the same vein, Robert Alter asserts that “there may [at any point in social history] be inherent limits on the access of the novelistic imagination to objective, collective realities” (Alter, 2008, p. x). And Moretti’s central point is that a shortage of linguistic resources for reproducing the experience of an urban community persisted as literature shifted its focus toward the portrayal of urban realities in the

nineteenth century. Moretti asks, “given the over-complication of the nineteenth-century urban setting - how did novels ‘read’ cities? By what narrative mechanisms did they make them ‘legible’, and turn urban noise into information?” (Moretti, 1999, p. 79). To answer this question, Moretti points to the reductive rendering techniques of the urban genre’s first wave; these novels “don’t show ‘London’, only a small, monochrome portion of it” (Moretti, 1999, p. 79). In order to make London legible, nineteenth century British novelists, including Austen and Dickens, reduce its complexity and its randomness, thereby amputating the richer, more unpredictable interactions that could occur in a more complex city (Moretti, 1999, p. 86). Moretti compares Dickens’s London with Balzac’s Paris; unlike Dickens, Balzac allows the complications of his urban subject to flourish and inform narrative possibility. The following quote presented by Elson et al. (2010) is actually used by Moretti to describe Balzac’s Paris specifically, not urban settings in general, and specifically not Dickens’s London:

As the number of characters increases, Moretti argues (following Bakhtin in his logic), social interactions of different kinds and durations multiply, displacing the family-centered and conversational logic of village or rural fictions. “The narrative system becomes complicated, unstable: the city turns into a gigantic roulette table, where helpers and antagonists mix in unpredictable combinations” (Moretti, 1999).

In summary, the simple fact that a novel is set in an urban environment (and the evocation of the urban setting by name or choice of props) does not equate with the creation of a truly urban *space*. The latter is the key that renders possible an urban story with an urban social world; “without a certain kind of space,” Moretti declares, “a certain kind of story is simply impossible” (Moretti, 1999, p. 100).

Moretti exposes another reductive rendering technique used by Dickens: the narrative crux of the family romance. This technique, he asserts, “is a further instance of the tentative, contradictory path followed by urban novels: as London’s random and

unrelated enclaves increase the ‘noise’, the ‘dissonance’, the complexity of the plot – the family romance tries to reduce it, turning London into a coherent whole” (Moretti, 1999, p. 130). Alter agrees, arguing that in Dickens’ *London*, “representation of human solidarity characteristically sequesters it in protected little enclaves within the larger urban scene” (Alter, 2008, p. 55) and that “in these elaborately plotted books of Dickens’s, almost no character is allowed to go to waste; each somehow is linked with the others as the writer deftly brings all the strands together in the complication and resolution of the story” (Alter, 2008, p. 67). In terms of the “perception of the fundamental categories of time and space, the boundaries of the self, and the autonomy of the individual” (Alter, 2008, p. xi), Dickens essentially writes a rural fiction, but in an urban setting.

To summarize these arguments: when novelists – like Dickens – employ narrative techniques not originally evolved for the portrayal of urban areas in novels with an urban *setting*, they fail to create in the novel the type of urban *space* in which an urban story with an urban social world is possible. Setting is sociological (it exists outside of novels), but space is literary (it exists only in novels): it is only the development of new practices in writing that are able to create truly urban spaces, those which reflect the fundamental transformations in the nature of human experience by the city: “Urban crowds and urban dwellings may reinforce a sense of isolation in individuals which, pushed to the extreme, becomes an incipient solipsism or paranoia. This feeling of being cut off from meaningful human connections finds a congenial medium in modes of narration – pioneered by Flaubert – that are rigorously centered in the consciousness of the character” Alter (2008, p. 107).

We now turn to Elson et al. (2010)’s presentation of the literary theories. They muddy the difference between setting and space, a serious flaw in interpreting Bakhtin. Urban setting does not equal urban space, and space – not setting – is what concerns Bakhtin’s chronotope. The nature of the space of a novel, not its explicit setting, defines what can happen in it, including its social relationships. Each text in Elson et al. (2010)’s corpus of 60 novels is classified as either rural or urban by the following defini-

tions. They define urban to mean set in a metropolitan zone, characterized by multiple forms of labor (not just agricultural), where social relations are largely financial or commercial in character. Conversely, rural is defined to mean set in a country or village zone, where agriculture is the primary activity, and where land-owning, non-productive, rent-collecting gentry are socially predominant. Thus, the distinction between rural and urban for Elson et al. (2010) is clearly one of setting, not one of space. Hypothesis EDM2 of Elson et al. (2010) is therefore *not* a correct representation of the theories they cite.

Interestingly, Elson et al. (2010) cannot validate their own hypothesis EDM2: their results suggest that the “urban” novels within the corpus do *not* belong to a fundamentally separate class of novels, insofar as basic frameworks of time and space inform the essential experience of the characters. They conclude:

We would propose that this suggests that the form of a given novel – the standpoint of the narrative voice, whether the voice is “omniscient” or not – is far more determinative of the kind of social network described in the novel than where it is set or even the number of characters involved.

Put differently, differences in novels’ social networks are more related to *literary* differences (space) than to non-literary differences (setting).

3 Expanded Set of Literary Hypotheses

In light of the analysis in the previous section, we propose that Elson et al. (2010)’s results, though they invalidate hypotheses EDM1 and EDM2, actually align with the parent theories from which they are derived. In direct opposition to EDM1 and EDM2, we expect our analysis to confirm the absence of correlation between setting and social network within our corpus of novels. While Elson et al. (2010)’s approach is constricted to examining social networks from the perspective of conversation, we obtain deeper insight into the novels by exploring an expanded set of hypotheses which takes general interaction and observation into account. If a comprehensive look at the social networks in our corpus confirms a lack of structural difference between

the social worlds of rural- and urban-set novels, we confirm the need to look beyond setting in order to pinpoint facets of novelistic form that do determine social networks.

Similar to the approach of Elson et al. (2010), our hypotheses concern (a) the implications of an increase in the number of characters, and (b) the implications of the dichotomy between rural and urban settings. However, unlike Elson et al. (2010), we do not claim any hypothesized relation between the increase in number of characters and the social structure. We formulate our own hypotheses (H1.1, H1.2, H3.1, H3.2, H5) concerning the increase in number of characters and study them out of curiosity and as an exploratory exercise. Furthermore, unlike Elson et al. (2010), we claim that literary theorists did not posit a relation between setting and social structure. Following is the set of hypotheses we validate in this paper:

- H0: As setting changes from rural to urban, there is no change in the number of characters. The number of characters is given by the formula 1 in table 1.
- H1.1: There is a positive correlation between the number of interactions and the number of characters. The number of interactions is given by the formula 3 in table 1.
- H1.2: There is a negative correlation between the number of characters and the number of other characters a character interacts with (unweighted version of H1.1, formula 4).
- H2.1: As setting changes from rural to urban, there is no change in the total number of interactions that occur. The number of interactions is given by the formula 3 in table 1.
- H2.2: As setting changes from rural to urban, there is no change in the average number of characters each character interacts with.
- H3.1: There is a positive correlation between the number of observations and the number of characters. The number of observations is given by the formula 3 in table 1.

- H3.2: There is a negative correlation between the number of characters a character observes (formula 4), and the number of characters.
- H4.1: As setting changes from rural to urban, there is no change in the total number of observations that occur. The number of observations is given by the formula 3 in table 1.
- H4.2: As setting changes from rural to urban, there is no change in the average number of observations performed by each character. (This hypothesis is the unweighted version of H4.1, formula 4, and the OBS counterpart of H2.2).
- H5: As the number of characters increases, the number of communities increases, but the average size of communities decreases.
- H6: As setting changes from rural to urban, there is no change in the number nor the average size of communities.

4 Methodology for Validating Hypotheses

Elson et al. (2010) provide evidence to invalidate EDM1. They report a positive Pearson’s correlation coefficient (**PCC**) between the number of characters and the number of dialogues to show that the two quantities are not inversely correlated. We use the same methodology to examine our hypotheses related to the number of characters.

Elson et al. (2010) provide evidence to invalidate EDM2. They extract various features from the social networks of rural and urban novels and show that these features are not statistically significantly different. They use the **homoscedastic t-test** to measure statistical significance (with $p < .05 \implies$ statistical significance). We employ the same methodology to examine our hypotheses related to the rural/urban dichotomy.

The features that Elson et al. (2010) use to invalidate EDM2 are as follows: (a) average degree, (b) rate of cliques, (c) density, and (d) rate of characters’ mentions of other characters. EDM2 posits that the number of characters in urban settings share lesser conversation as compared to the rural settings. The average degree (count of the number of conversations normalized by the number of characters, see formula 4 in Table 1) seems to be the metric that

is relevant for (in)validating EDM2. It is unclear why Elson et al. (2010) report the correlation between other features to invalidate EDM2. We therefore, validate our formulation of the theory (similar to EDM2) using only the average degree metric.

5 Types of Networks

This section provides definitions for the three different types of networks we consider in our study.

5.1 Conversation Network

Elson et al. (2010) defined a conversation network as a network in which nodes are characters and links are conversations. The authors defined a conversation as a continuous span of narrative time in which a set of characters exchange dialogues. Since dialogues are denoted in text by quotation marks, Elson et al. (2010) used simple regular expressions for dialogue detection. However, associating dialogues with their speakers (a task known as *quoted speech attribution*) turned out to be non-trivial (Elson and McKeown, 2010; He et al., 2013). In a separate work, Elson and McKeown (2010) presented a feature-based, supervised machine learning system for performing quoted speech attribution. Using this system, Elson et al. (2010) successfully extracted conversation networks from the novels in their corpus. We refer to the system as **EDM2010** throughout this paper.

5.2 Observation and Interaction Networks

In our past work (Agarwal et al., 2010), we defined a social network as a network in which nodes are characters and links are *social events*. We defined two broad categories of social events: observations (OBS) and interactions (INR). Observations are defined as unidirectional social events in which *only one* entity is cognitively aware of the other. Interactions are defined as bidirectional social events in which *both* entities are cognitively aware of each other *and* of their mutual awareness.

In Example 1, Mr. Woodhouse is *talking about* Emma. He is therefore cognitively aware of Emma. However, there is no evidence that Emma is also aware of Mr. Woodhouse. Since only one character is aware of the other, this is an observation event directed from Mr. Woodhouse to Emma.

No	Name	Formula	Weighted?	Remark
1	# of characters	$ V $		
2	# of interaction pairs	$ E $	unw.	
3	# of interactions	$W := \sum_{i=1}^{ E } w_i$	weighted	
4	average degree	$\frac{\sum_{v \in V} E_v }{ V } = \frac{2 E }{ V }$	unw.	number of characters a character interacts with on average
5	average weighted degree	$\frac{\sum_{u \in V} \sum_{v \in E_u} w_{u,v}}{ V } = \frac{2W}{ V }$	weighted	number of interactions a character has on average

Table 1: Table connecting the social network terminology to the natural language interpretation, along with the formulae. Interactions may be replaced with observations to obtain the corresponding formula.

- (1) “[Emma] never thinks of herself, if she can do good to others,” {rejoined} [Mr. Woodhouse] **OBS**

In Example 2, Mr. Micawber is talking about going home with Uriah. Since Mr. Micawber is talking about Uriah, there is a directed OBS link from Mr. Micawber to Uriah. Since he went home with Uriah, they must both have been aware of each other and of their mutual awareness. Thus, there is a bidirectional INR link between the two characters.

- (2) [Mr. Micawber] {said}_{OBS}, that [he] had {gone home with}_{INR} [Uriah] **OBS and INR**

In Example 3, the author (Jane Austen) states a fact about three characters (Elton, Mr. Knightley, and Mr. Weston). However, the author does not tell us about the characters’ cognitive states, and thus there is no social event between the characters.

- (3) [Elton]’s manners are superior to [Mr. Knightley]’s or [Mr. Weston]’s. **NoEvent**

As these examples demonstrate, the definition of a social event is quite broad. While quoted speech (detected by Elson et al. (2010)) represents only a strict sub-set of interactions, social events may be linguistically expressed using other types of speech as well, such as reported speech.

In our subsequent work (Agarwal and Rambow, 2010; Agarwal et al., 2013b; Agarwal et al., 2014), we leveraged and extended ideas from the relation extraction literature (Zelenko et al., 2003; Kambhatla, 2004; Zhao and Grishman, 2005; GuoDong et al., 2005; Harabagiu et al., 2005; Nguyen et al.,

2009) to build a tree kernel-based supervised system for automatically detecting and classifying social events. We used this system for extracting observation and interaction networks from novels. We will refer to it as **SINNET** throughout this paper.

5.3 Terminology Regarding Networks

A network (or graph), $G = (V, E)$, is a set of vertices (V) and edges (E). The set of edges incident on vertex v is written E_v . In weighted networks, each edge between nodes u and v is associated with a weight, denoted by $w_{u,v}$. In the networks we consider, weight represents the frequency with which two people interact or observe one another. An edge may be directed or undirected. Interactions (INR) are undirected edges and observations (OBS) are directed edges. Table 1 presents the name and the mathematical formula for social network analysis metrics we use to validate the theories.

Edges in a network are typed. We consider four types of networks in this work: networks with undirected interaction edges (INR), with directed observation edges (OBS), with a combination of interaction and observation edges (INR + OBS), and with a combination of interaction, observation, and undirected conversational edges (CON). We denote these networks by $G_{INR} = (V, E_{INR})$, G_{OBS} , $G_{INR+OBS}$, and $G_{INR+OBS+CON}$ respectively. Each of these networks may be weighted or unweighted.

6 Evaluation of SINNET

In our previous work, we showed that SINNET adeptly extracts the social network from one work of fiction, *Alice in Wonderland* (Agarwal et al., 2013b).

Novel Excerpt	CONV-GOLD		INT-GOLD					
	EDM2010	SINNET	EDM2010			SINNET		
	R	R	P	R	F1	P	R	F1
Emma	0.40	0.70	1.0	0.13	0.22	0.86	0.48	0.61
Study in Scarlet	0.50	0.50	1.0	0.18	0.31	0.69	0.41	0.51
David Copperfield	0.70	0.80	1.0	0.22	0.36	0.80	0.63	0.70
Portrait of a Lady	0.66	0.66	1.0	0.22	0.36	0.73	0.44	0.55
Micro-Average	0.56	0.68	1.0	0.18	0.30	0.79	0.50	0.61

Table 3: Performance of the two systems on the two gold standards.

Novel Excerpt	# of char.	# of links	
	pairs	CG	IG
Emma	91	10	40
Study in Scarlet	55	8	22
David Copperfield	120	10	32
Portrait of a Lady	55	6	18

Table 2: A comparison of the number of links in the two gold standards; CG is CONV-GOLD and IG is INT-GOLD

In this paper, we determine the effectiveness of SINNET on an expanded collection of literary texts. Elson et al. (2010) presented a gold standard for measuring the performance of EDM2010, which we call CONV-GOLD. This gold standard is not suitable for measuring the performance of SINNET because SINNET extracts a larger set of interactions beyond conversations. We therefore created another gold standard more suitable for evaluating SINNET, and refer to it as INT-GOLD.

6.1 Gold standards: CONV-GOLD and INT-GOLD

Elson et al. (2010) created their gold standard for evaluating the performance of EDM2010 using excerpts from four novels: Austen’s *Emma*, Conan Doyle’s *A Study in Scarlet*, Dickens’ *David Copperfield*, and James’ *The Portrait of a Lady*. The authors enumerated all pairs of characters for each novel excerpt. If a novel features n characters, its corresponding list contains $\frac{n*(n-1)}{2}$ elements. For each pair of characters, annotators were asked to mark “1” if the characters *converse* (defined in Section 5) and “0” otherwise. Annotators were asked to identify conversations framed with both direct (quoted) and indirect (unquoted) speech.

As explained in previous sections, *conversations* are a strict subset of general *interactions*. Since SINNET aims to extract the entire set of observations and interactions, the gold standard we created records all forms of observation and interaction between characters. For each pair of characters, annotators were asked to mark “1” if the characters *observe* or *interact* and “0” otherwise.

Table 2 presents the number of character pairs in each novel excerpt, the number of character pairs that converse according to CONV-GOLD and the number of character pairs that observe or interact according to INT-GOLD. The difference in the number of links between CONV-GOLD and INT-GOLD suggests that the observation and interaction of many more pairs of characters is expressed through reported speech in comparison to conversational speech. For example, the number of conversational links identified in the excerpt from *Emma* by Jane Austen was 10, while the number of interaction links identified was 40.

6.2 Evaluation and Results

Table 3 presents the results for the performance of EDM2010 and SINNET on the two gold standards (CONV-GOLD and INT-GOLD). The recall of SINNET is significantly better than that of EDM2010 on CONV-GOLD (columns 2 and 3), suggesting that most of the links expressed as quoted conversations are also expressed as interactions via reported speech. Note that, because SINNET extracts a larger set of interactions, we do not report the precision and F1-measure of SINNET on CONV-GOLD. By definition, SINNET will predict links between characters that may not be linked in CONV-GOLD; therefore the precision (and thus

Hypothesis	As # of characters ↑ ...		As settings go from rural to urban ...	
	PCC	Valid?	t-test	Valid?
[H0] ... # of characters ~			$p > 0.05$	✓
[H1.1] ... # of interactions ↑	0.83	✓		
[H1.2] ... # of characters interacted with ↓	-0.36	✓		
[H2.1] ... # of interactions ~			$p > 0.05$	✓
[H2.2] ... # of characters interacted with ~			$p > 0.05$	✓
[H3.1] ... # of observations ↑	0.77	✓		
[H3.2] ... # of characters observed ↓	-0.36	✓		
[H4.1] ... # of observations ~			$p > 0.05$	✓
[H4.2] ... # of characters observed ~			$p > 0.05$	✓
[H5] ... # of communities ↑	0.98	✓		
[H5] ... average size of communities ↓	-0.26	✓		
[H6] ... # of communities ~			$p > 0.05$	✓
[H6] ... average size of communities ~			$p > 0.05$	✓

Table 4: Hypotheses and results. All correlations are statistically significant. ~ denotes no significant change. As an example, hypothesis H0 may be read as: As settings go from rural to urban ... the number of characters does not change significantly.

F1-measure) of SINNET will be low (and uninterpretable) on CONV-GOLD.

Table 3 additionally presents the performance of the two systems on INT-GOLD (the last six columns). These results show that EDM2010 achieves perfect precision, but significantly lower recall than SINNET (0.18 versus 0.50). This is expected, as EDM2010 was not trained (or designed) to extract any interactions besides conversation.

6.3 Discussion of Results

If there are any conversational links that EDM2010 detects but SINNET misses, then the two systems should be treated as complementary. To determine whether or not this is the case, we counted the number of links in all four excerpts that are detected by EDM2010 and missed by SINNET. For Austen’s *Emma*, SINNET missed two links that EDM2010 detected (with respect to INT-GOLD). For the other three novels, the counts were SINNET two, zero, and one, respectively. In total, the number of links that SINNET missed and EDM2010 detected is five out of 112. Since the precision of EDM2010 is perfect, it seems advantageous to combine the output of the two systems.

7 Results for Testing Literary Hypotheses

Table 4 presents the results for all hypotheses (H0-H6) formulated in this paper. There are two broad categories of hypotheses: (1) ones that comment on social network analysis metrics (the rows) based on the increase in the number of characters (columns 2 and 3), and (2) ones that comment on the social network analysis metrics based on the type of setting (rural versus urban, columns 4 and 5).

The results show that as settings change from rural to urban, there is no significant change in the number of characters (row H0, column t-test). Furthermore, as the number of characters increases, the number of interactions also increases with a high Pearson correlation coefficient of 0.83 (row H1.1, column PCC). Similarly, for all other hypotheses, the relation between the number of characters and the setting of novels behaves as expected in terms of various types of networks and social network analysis metrics. Our results thus provide support for the cogency of the original theories.

These results highlight one of the critical findings of this paper: while network metrics are significantly correlated with the number of characters, there is **no correlation at all between setting and number of characters** within our corpus (hypothesis H0 is valid). If H0 were invalid, then all hy-

potheses concerning the effects of setting would be false. However, since H_0 is true, we may conclude that setting (as defined by our rural/urban classification) has **no predictive effect** on any of the aspects of social networks that we investigate.

We also consider whether examining different network types (interaction, observation, and combination) in conjunction produces the same results as examining each individually. The results indeed align with those in Table 4, but with slightly different correlation numbers. We give one example: we find that the correlation between number of characters and number of interactions (hypothesis H1.1) increases from 0.83 for G_{INR} alone (as shown in Table 4) to 0.85 for $G_{INR+OBS}$ and also 0.85 for $G_{INR+OBS+CONV}$. This pattern is observed for all hypotheses.

8 Conclusion and Future Work

In this paper, we investigated whether social network extraction confirms long-standing assumptions about the social worlds of nineteenth-century British novels. Namely, we set out to verify whether the social networks of novels explicitly located in urban settings exhibit structural differences from those of rural novels. Elson et al. (2010) had previously proposed a hypothesis of difference as an interpretation of several literary theories, and provided evidence to invalidate this hypothesis on the basis of conversational networks. Following a closer reading of the theories cited by Elson et al. (2010), we suggested that their results, far from invalidating the theories themselves, actually support their cogency. To extend Elson et al. (2010)’s findings with a more comprehensive look at social interactions, we explored the application of another methodology for extracting social networks from text (called SINNET) which had previously not been applied to fiction. Using this methodology, we were able to extract a rich set of observation and interaction relations from novels, enabling us to build meaningfully on previous work. We found that the rural/urban distinction proposed by Elson et al. (2010) indeed has no effect on the structure of the social networks, while the number of characters does.

As our findings support our literary hypothesis that the urban novels within Elson et al. (2010)’s

original corpus do not belong to a fundamentally separate class of novels, insofar as the essential experience of the characters is concerned possible directions for future research include expanding our corpus in order to identify novelistic features that do determine social worlds. We are particularly interested in studying novels which exhibit innovations in narrative technique, or which occur historically in and around periods of technological innovation. Lastly, we would like to add a temporal dimension to our social network extraction, in order to capture information about how networks transform throughout different novels.

Acknowledgments

We would like to thank anonymous reviewers for their insightful comments. We would also like to thank David Elson for providing the gold standard the data set used in their previous work. This paper is based upon work supported in part by the DARPA DEFT Program. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- Apoorv Agarwal and Owen Rambow. 2010. Automatic detection and classification of social events. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1024–1034, Cambridge, MA, October. Association for Computational Linguistics.
- Apoorv Agarwal, Owen C. Rambow, and Rebecca J. Passonneau. 2010. Annotation scheme for social network extraction from text. In *Proceedings of the Fourth Linguistic Annotation Workshop*.
- Apoorv Agarwal, Anup Kotalwar, and Owen Rambow. 2013a. Automatic extraction of social networks from literary text: A case study on alice in wonderland. In *the Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP 2013)*.
- Apoorv Agarwal, Anup Kotalwar, Jiehan Zheng, and Owen Rambow. 2013b. Sinnet: Social interaction network extractor from text. In *Sixth International Joint Conference on Natural Language Processing*, page 33.
- Apoorv Agarwal, Sriramkumar Balasubramanian, Anup Kotalwar, Jiehan Zheng, and Owen Rambow. 2014. Frame semantic tree kernels for social network extraction from text. *14th Conference of the European Chapter of the Association for Computational Linguistics*.
- Robert Alter. 2008. *Imagined cities: urban experience and the language of the novel*. Yale University Press.
- Mikhail M Bakhtin. 1937. Forms of time and of the chronotope in the novel: Notes toward a historical poetics. *Narrative dynamics: Essays on time, plot, closure, and frames*, pages 15–24.
- Terry Eagleton. 1996. *Literary theory: An introduction*. U of Minnesota Press.
- Terry Eagleton. 2013. *The English novel: an introduction*. John Wiley & Sons.
- David K Elson and Kathleen McKeown. 2010. Automatic attribution of quoted speech in literary narrative. In *AAAI*.
- David K. Elson, Nicholas Dames, and Kathleen R. McKeown. 2010. Extracting social networks from literary fiction. *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147.
- Zhou GuoDong, Su Jian, Zhang Jie, and Zhang Min. 2005. Exploring various knowledge in relation extraction. In *Proceedings of 43th Annual Meeting of the Association for Computational Linguistics*.
- Sanda Harabagiu, Cosmin Adrian Bejan, and Paul Morarescu. 2005. Shallow semantics for relation extraction. In *International Joint Conference On Artificial Intelligence*.
- Hua He, Denilson Barbosa, and Grzegorz Kondrak. 2013. Identification of speakers in novels. In *ACL (1)*, pages 1312–1320.
- Nanda Kambhatla. 2004. Combining lexical, syntactic, and semantic features with maximum entropy models for extracting relations. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, page 22. Association for Computational Linguistics.
- Franco Moretti. 1999. *Atlas of the European novel, 1800-1900*. Verso.
- Franco Moretti. 2005. *Graphs, Maps, Trees: Abstract Models for a Literary History*. Verso.
- Truc-Vien T. Nguyen, Alessandro Moschitti, and Giuseppe Riccardi. 2009. Convolution kernels on constituent, dependency and sequential structures for relation extraction. *Conference on Empirical Methods in Natural Language Processing*.
- Raymond Williams. 1975. *The country and the city*. Oxford University Press.
- Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. 2003. Kernel methods for relation extraction. *The Journal of Machine Learning Research*, 3:1083–1106.
- Shubin Zhao and Ralph Grishman. 2005. Extracting relations with integrated information using kernel methods. In *Proceedings of the 43rd Meeting of the ACL*.