
Studying Literary Characters and Character Networks

Andrew Piper

andrew.piper@mcgill.ca
McGill University, Canada

Mark Algee-Hewitt

malgeehe@stanford.edu
Stanford University, United States of America

Koustuv Sinha

koustuvsinha@gmail.com
McGill University, Canada

Derek Ruths

derek.ruths@mcgill.ca
McGill University, Canada

Hardik Vala

hardik.vala@mail.mcgill.ca
McGill University, Canada

The study of character has long been one of the central concerns of literary theory. For the Russian formalists, embodied above all in work of Vladimir Propp (1968), character was primarily a “type,” one that served different narrative functions (“the hero is married and ascends the throne”). For poststructuralists that came in the wake of Propp, character was nothing more than a rhetorical “effect,” one more example of the referential phallacy of naïve readers (Barthes 1970; Culler 2002). Subsequent studies attempted to account for this dual nature of characters, the they are both rhetorical devices and also constrained by real-world references, the requirements of being human that constrain what characters can do and say (Phelan 1989; Jannidis 2004; Frow 2014). More recent research has begun to emphasize the affective or identificatory role that characters play for readers (Lynch 1998; Brewer 2011). According to this view, characters are the media through which readers come to terms with new kinds of social experience. Drawing on the field of cognitive science, other work by Zunshine (2006) and Vermeule (2011) has argued in a less historical vein that characters are useful tools through which to model “theories of mind,” means for learning

about and hypothetically experiencing human cognition.

It is within this context that our three papers situate themselves in order to understand the ways in which computation impacts the study of literature. Each is fundamentally concerned with how the increased volume of information surrounding characters impacts our understanding of the idea of character – whether it is examining several thousand plays in which characters appear, several thousand interactions between characters, or the millions of words surrounding characters’ appearance on the page. Characters are fundamentally social in literature and these computational methods are designed to better understand that sociability.

Mark Algee-Hewitt’s paper concerns itself with the study of social networks in 3,900 plays across four-centuries. It asks how the morphology of the social networks represented on stage represent (or resist) both the politics and aesthetics of a period and, more importantly, how those social networks evolve over time? In his paper, he will move beyond the network analysis of a single play by examining the network structure of a large corpus of English dramas written and performed between 1500 and 1920. By applying a series of summary statistics drawn from the field of social network analysis to the individual plays, he is able to trace the history of dramatic representations of the social sphere and shed new light on the evolution of both the protagonist and the periphery in modern drama.

Khoustiv Sinha, Andrew Piper, Derek Ruth’s paper takes a step back to ask the even more fundamental question: what *is* an interaction? Before we move to the extraction and mapping of social networks in fiction, we first need to study how readers understand the very idea of “interaction.” In this project, he examines reader annotations across a data set of over 1,000 social interactions drawn from popular contemporary fiction and non-fiction. In doing so, he addresses not only the level of agreement between readers, but also the types of interactions that produce more or less agreement among readers. What are the qualities of social relationships in literature that generate more ambiguity among readers and what does that have to tell us about the social investments of literary texts?

Finally, Andrew Piper and Hardik Vala’s paper introduces a new tool that identifies 28 different features aligned with practices of characterization. These features range across a variety of different categories, from positionality (the character’s agency), modality (behavior), descriptiveness, to social categories like

proximity to other characters or the distribution of character counts in a text. As they will show, this tool can allow us to derive novel insights about the history of character development in literary texts.

Distributed Character: Quantitative Models of the English Stage, 1500–1920

Mark Algee-Hewitt

The use of network graphs to represent social networks of characters within novels or plays has played an important role in quantitative textual analysis (see for example Agwar et al 2012, Bingenheimer et al 2011, Elson et al 2010). In this paper, I move beyond the network as a visual object, and instead, draw upon the quantitative metrics of the graph to explore the large-scale changes to the structure of English drama across four hundred years. What can the overall structure of the play can tell us both about the aesthetics of literary production of a given period, and what we can learn about the play by disaggregating the morphology from both the stagecraft and the language that, until now, have made up the two poles of dramatic criticism?

The power of networks lie in their precise, mathematical, description of a set of relationships that can be quantified, measured and aggregated in ways that are unavailable to the reader of a text. Yet, most work has been focused on the use of single networks to describe single plays. For example, in his work on character networks in *Hamlet*, Franco Moretti turns quickly from the quantitative network analysis to the qualitative approach to the plot: “I soon realized that the machine-gathering of the data, essential to large-scale quantification, was not yet a realistic possibility [...] So, from its very first section, the essay drifted from quantification to the qualitative analysis of plot” (Moretti 2011). In this paper, I introduce an automated, rule-based, parsing of the 3439 English plays in the Chadwyck Healy drama corpus in order to perform the kind of large-scale quantitative analysis that Moretti gestures towards, but is unable to realize. The algorithm uses the existing XML markup in the corpus in order to extract speeches and assign them to characters as speakers and recipients, resolving co-references to character abbreviations (this is similar to the automated method employed by Trilcke et al. 2016, although the summary statistics that I extract are quite different).

Drawing on this tagged corpus, I create a social network for each drama and extract a series of summary features based on both the eigenvector and betweenness centralities of each play. The first summary sta-

tistic that I calculate is the Gini Coefficient of the eigenvector centrality. Originally designed to measure income inequality within an economic system, the Gini Coefficient is a single number between 0 and 1 that indicates how evenly a set of resources (wealth, income or, in this case, centrality) is distributed across a population (here, of characters). In the Gini coefficient measurement in the corpus at large (Figure 1), there is a clear historical pattern being played out. Over time, between 1550 and 1900, the Gini coefficients of the plays exhibits a clear downward trend, from plays with a small core and a large, non-central periphery in the early century, to plays with a relatively large core and a small periphery in the eighteenth and nineteenth centuries and a large discontinuity between 1650 and 1700. What this metric seems to indicate, then, is the disappearance of the periphery of the English drama over time. Rather than suggesting significant structural changes to the core of the play, the largest influence in the Gini coefficient is the presence of a large periphery, whose members rarely speak (and more importantly, rarely interact with the center of power) and who therefore bring down the Gini coefficient for the entire play. Over time, then, this periphery disappears as casts get smaller and actions take place among an increasingly more tightly knit set of characters. Servants, retainers, guards, acquaintances and messengers, so important during the early modern period, disappear with increasing regularity in the later periods, echoing the reduced function such figures had in society itself, as dramas move, following Habermasian logic, from the throne room to the drawing room, becoming personal and intimate, rather than mythic, political and impersonal.

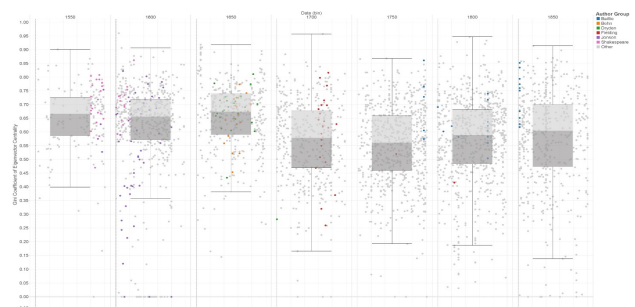


Figure 1. Gini Coefficients of Eigenvector Centrality over time. The corpus is divided into 50 year bins (with the plays in each bin arranged chronologically). Colors indicate selected canonical authors.

The second metric is the percentage of characters in the top quartile of the eigenvector centrality distribution. This measures the size of the core of the play and tells an equally striking and parallel story (Figure

2). Although the relative regularity of the measurement makes it less immediately apparent, there is a constant historical increase in the percentage of characters in the top quartile of eigenvector centrality scores. While the falling Gini Coefficients speak to the disappearance of the periphery, this metric reveals what happens to the remaining core. Rather than follow the same pattern of the early modern period (with few highly central characters), the disappearance of the periphery means that more centrality is allotted between the core characters. This speaks not just to the increasing size of the core, but, more importantly, to the tendency of having plays that feature multiple sub-networks, each with their own protagonist. In a play with a single network, it is easy for one character to dominate it, but in a play whose action is divided between competing communities, each community can have its own central figure. If we can tell the protagonist of an early modern drama by his or her high eigenvector centrality compared to the rest of the cast, then by the seventeenth century the single protagonist has been dispersed between multiple characters who all evidence a high eigenvector centrality, distributing the function of the protagonist (and/or the antagonist) among a growing number of central characters.




Figure 2. Percentage of characters in the top quartile of the eigenvector centrality distribution in the play.

As opposed to the eigenvector centrality's relationship to the protagonist, betweenness centrality speaks to the mediatedness of the drama. That is, if a high betweenness centrality indicates a character that mediates other character's interactions (such that they have to pass through her), then the scaled maximum betweenness centrality of a dramatic network overall, which measures the relatively importance of bridging characters, indicates the extent to which this mediating function is important to the drama as a whole. At the level of the corpus, the normalized maximum betweenness centrality, the relative importance of the bridging character, decreases across the century, very quickly from 1590 to 1640, and then more slowly

across the remaining two and a half centuries: the average maximum betweenness centrality in a play drops by over 750 across just the sixteenth century. Again, the largest discontinuity lies between 1650 and 1700: there is a clearly a lasting effect on the structure of dramatic networks from the puritan shuttering of the theaters during the interregnum. The English drama that returns during the restoration is evidently not the same as the English stage before Cromwell.

Understanding Reader-Identified Social Interactions in Literature

Koustuv Sinha, Andrew Piper, Derek Ruths

Social network analysis begins with the primacy of character as its object of study. In this, it fits within an already well-established area of inquiry within literary theory, one whose formal study extends back until at least the early twentieth century if not earlier (Propp 1968). Where social network analysis differs from this tradition is through the emphasis on *dynamic interactions* as a key to understanding the narrative function of character. Whether exploring the afterlife of fan fiction, theories of mind, affective identification, or the typologies of character, what all of the pre-computational work on character has in common is an emphasis on understanding character in the singular. Social network analysis argues instead that the meaning of any character is a function of his or her relationships with respect to all of the other characters introduced over the course of a story (Woloch 2009). Character networks offer a way to study not simply the types or themes or affective connections between readers and imaginary people. Rather, they afford us the ability to understand the social imaginings of writers, periods, and genres.

Several initial attempts to introduce social network analysis into the study of literature have already been made. Character networks have been studied within three major European epics to understand their relation to contemporary models of social networks (MacCaron/Kenna 2012); an abridged version of a single well-known literary work (*Alice in Wonderland*) to test differences between interactions and observations on character centrality (Agarwal 2012); nineteenth-century novels to understand the correlation between dialogue and setting (Elson 2010); as a form of narrative generation (Sack 2013); and the genre of classical drama to better understand the notion of tragic conflict (Moretti 2013; Karsdorp et al 2015).

Each of these works has added to our understanding of the relationship between character and literary form in important ways. And yet at the core of each of

these studies lies a fundamental assumption about the self-evident nature of an “interaction.” Initial attempts to use machine learning to derive interactions on prose texts have shown very poor performance (Agarwal 2012 reports a maximum F1 score of 0.61). What this indicates at least in part is that interactions are highly complex verbal constructions which we cannot easily assume pre-exist our attempts at extracting them.

To counter this problem, we have designed a study to explore reader agreement across a variety of text passages (1,000) drawn from popular contemporary fiction and non-fiction. Rather than begin with a stable set of interaction types, however, our goal is to infer possible classes of interactions and then understand which of these classes generate more ambiguity among readers. We perform this in three phases. In the first phase, we ask coders to identify minimally defined interactions using a standardized web interface (where an interaction consists of two entities and an action linking them). Our goal here is not to pre-define types of interactions as in other studies (Agarwal), but to better understand how readers intuitively understand social interactions between characters. As we have shown in another study, readers indicate very high agreement in identifying character aliases (i.e. determining what is an entity (Vala et al.)). In the second phase, we use unsupervised clustering techniques to identify different interaction “types” based on syntactic and lexical features of the labeled interactions. Third, we then measure reader agreement across these different types. While we want to know overall how well readers agree on defining interactions, we also want to understand if different types of interactions across different types of writing (fiction/non-fiction) illustrate significantly higher levels of disagreement. This is a first step in understanding the unique ways literary texts generate social complexity, not simply through the quantity of interactions but also importantly through their qualities.

Emma: A Feature Space for Studying Character

Andrew Piper and Hardik Vala

This paper will argue that computation has an important role to play in understanding the nature of characters and the process of what we might generally term characterization – the writerly act of generating animate entities through language. With an estimated 86 characters per novel in the nineteenth century and a conservative estimate of 20,000 novels published during this period in the English language, there are

over 1.7 million unique characters that appear in that one century and one language alone. Even if we condition on main characters, we are still looking at several thousand distinct entities. At the same time, there are not only a great number of characters in literature, but there is also a tremendous amount of information surrounding even one primary character. Like other highly frequent textual features such as conjunctions or punctuation, characters are abundant across the pages of individual novels. Personal pronouns alone account for roughly 12% of all tokens, and if one adds in proper names the number of character occurrences is closer to 16% – or one in every six words! Like the abundance of characters, such semiotic abundance surrounding characters poses problems for inherited critical methods. How can we be sure that our claims about “character” are capturing the broad and potentially diverse ways that characters are depicted in novels, this larger mass of fictional beings and what it means to be fictional?

In order to address this question, we have developed a computational tool designed for the study of character. Its aim is to identify 28 different features that relate to qualities that characters may possess. These range across categories like distinctiveness (how distinctive is the main character from other characters within the novel); positionality (how often is the character the agent or object of a sentence or a possessor of some object); centrality (how important is the protagonist relative to other characters in the novel); and modality (what kinds of behaviors and descriptions inform this character’s identity, such as cogitation, perception, motion, embodiment and even clothing or dress).

Rather than start with known “types” of character, this tool allows us to implement a more multi-dimensional understanding of character and use that representation to think about the relationships between novels. Prior work on stylistic analysis has not differentiated between various aspects of texts when comparing them to each other. The novel is taken as a unified whole. Our character feature tool allows readers to begin to explore these different sub-domains of a novel, which in our case refers to the language used to construct character. In our presentation, we will discuss the mechanics that underlie the tool, which implements a modified version of BookNLP (Bamman 2014) and the Stanford dependency parser in order to identify words related to character. We will also discuss a case study in which we explore the identity of “introversion” in novels from the nineteenth century to the present. As we will show, the character feature

tool allows us to construct not only familiar narratives about the history of the novel – wherein the representation of interiority is strongly gendered around female protagonists – but also novel and nuanced insights about that tradition when we follow these features across a broader swath of time. As we will show, interiority no longer remains the distinctive quality of feminine heroines but is transposed onto a very different generic and gender scene – the male hero of science fiction.