Character-to-Character Sentiment Analysis in Shakespeare's Plays

Eric T. Nalisnick **Henry S. Baird**
Dept. of Computer Science and Engineering
Lehigh University
Bethlehem, PA 18015, USA
{etn212,hsb2}@lehigh.edu

Abstract

We present an automatic method for analyzing sentiment dynamics between characters in plays. This literary format's structured dialogue allows us to make assumptions about who is participating in a conversation. Once we have an idea of who a character is speaking to, the sentiment in his or her speech can be attributed accordingly, allowing us to generate lists of a character's enemies and allies as well as pinpoint scenes critical to a character's emotional development. Results of experiments on Shakespeare's plays are presented along with discussion of how this work can be extended to unstructured texts (i.e. novels).

1 Introduction

Insightful analysis of literary fiction often challenges trained human readers let alone machines. In fact, some humanists believe literary analysis is so closely tied to the human condition that it is impossible for computers to perform. In his book *Reading Machines: Toward an Algorithmic Criticism*, Stephen Ramsay (2011) states:

> Tools that can adjudicate the hermeneutical parameters of human reading experiences...stretch considerably beyond the most ambitious fantasies of artificial intelligence.

Antonio Roque (2012) has challenged Ramsay's claim, and certainly there has been successful work done in the computational analysis and modeling of narratives, as we will review in the next section. However, we believe that most previous work (except possibly (Elsner, 2012)) has failed to directly address the root cause of Ramsay's skepticism: can computers extract the emotions encoded in a narrative? For example, can the love that Shakespeare's Juliet feels for Romeo be computationally tracked? Empathizing with characters along their journeys to emotional highs and lows is often what makes a narrative compelling for a reader, and therefore we believe mapping these journeys is the first step in capturing the human reading experience.

Unfortunately but unsurprisingly, computational modeling of the emotional relationships described in natural language text remains a daunting technical challenge. The reason this task is so difficult is that emotions are indistinct and often subtly conveyed, especially in text with literary merit. Humans typically achieve no greater than 80% accuracy in sentiment classification experiments involving product reviews (Pang et al., 2002) (Gamon, 2004). Similar experiments on fiction texts would presumably yield even higher error rates.

In order to attack this open problem and make further progress towards refuting Ramsay's claim, we turn to shallow statistical approaches. Sentiment analysis (Pang and Lee, 2008) has been successfully applied to mine social media data for emotional responses to events, public figures, and consumer products just by using emotion lexicons–lists that map words to polarity values (+1 for positive sentiment, -1 for negative) or valence values that try to capture degrees of polarity. In the following paper, we describe our attempts to use modern sentiment lexicons and dialogue structure to algorithmically track and model–with no domain-specific customization–the emotion dynamics between characters in Shakespeare's plays.[1]

2 Sentiment Analysis and Related Work

Sentiment analysis (SA) is now widely used commercially to infer user opinions from product reviews and social-media messages (Pang and Lee,

[1] XML versions provided by Jon Bosak: http://www.ibiblio.org/xml/examples/shakespeare/

Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, pages 479–483,
Sofia, Bulgaria, August 4-9 2013.

2008). Traditional machine learning techniques on n-grams, parts of speech, and other bag of words features can be used when the data is labeled (*e.g.* IMDB's user reviews are labeled with one to ten stars, which are assumed to correlate with the text's polarity) (Pang et al., 2002). But text annotated with its true sentiments is hard to come by so often labels must be obtained via crowdsourcing.

Knowledge-based methods (which also typically rely on crowdsourcing) provide an alternative to using labeled data (Andreevskaia and Bergler, 2007). These methods are driven by sentiment lexicons, fixed lists associating words with "valences" (signed integers representing positive and negative feelings) (Kim and Hovy, 2004). Some lexicons allow for analysis of specific emotions by associating words with degrees of fear, joy, surprise, anger, anticipation, etc. (Strapparava and Valitutti, 2004) (Mohammad and Turney, 2008). Unsurprisingly, methods which, like these, lack deep understanding often work more reliably as the length of the input text increases.

Turning our attention now to automatic semantic analysis of fiction, it seems that narrative modeling and summarization has been the most intensively studied application. Chambers and Jurafsky (2009) described a system that can learn (without supervision) the sequence of events described in a narrative, and Elson and McKeown (2009) created a platform that can symbolically represent and reason over narratives.

Narrative structure has also been studied by representing character interactions as networks. Mutton (2004) adapted methods for extracting social networks from Internet Relay Chat (IRC) to mine Shakespeare's plays for their networks. Extending this line of work to novels, Elson and McKeown (2010) developed a reliable method for speech attribution in unstructured texts, and then used this method to successfully extract social networks from Victorian novels (Elson et al., 2010)(Agarwal et al., 2012).

While structure is undeniably important, we believe analyzing a narrative's emotions is essential to capturing the 'reading experience,' which is a view others have held. Alm and Sproat (2005) analyzed *Brothers Grimm* fairy tales for their 'emotional trajectories,' finding emotion typically increases as a story progresses. Mohammad (2011) scaled-up their work by using a crowdsourced emotion lexicon to track emotion dynamics over the course of many novels and plays, including Shakespeare's. In the most recent work we are aware of, Elsner (2012) analyzed emotional trajectories at the character level, showing how Miss Elizabeth Bennet's emotions change over the course of *Pride and Prejudice*.

3 Character-to-Character Sentiment Analysis

Character	Hamlet's Sentiment Valence Sum
Guildenstern	31
Polonius	25
Gertrude	24
Horatio	12
Ghost	8
Marcellus	7
Osric	7
Bernardo	2
Laertes	-1
Ophelia	-5
Rosencrantz	-12
Claudius	-27

Figure 1: The characters in *Hamlet* are ranked by Hamet's sentiment towards them. Expectedly, Claudius draws the most negative emotion.

We attempt to further Elsner's line of work by leveraging text structure (as Mutton and Elson did) and knowlege-based SA to track the emotional trajectories of interpersonal relationships rather than of a whole text or an isolated character. To extract these relationships, we mined for character-to-character sentiment by summing the valence values (provided by the *AFINN* sentiment lexicon (Nielsen, 2011)) over each instance of continuous speech and then assumed that sentiment was directed towards the character that spoke immediately before the current speaker. This assumption doesn't always hold; it is not uncommon to find a scene in which two characters are expressing feelings about someone offstage. Yet our initial results on Shakespeare's plays show that the instances of face-to-face dialogue produce a strong enough signal to generate sentiment rankings that match our expectations.

For example, Hamlet's sentiment rankings upon the conclusion of his play are shown in Figure 1. Not surprisingly, Claudius draws the most negative sentiment from Hamlet, receiving a score of -27. On the other hand, Gertrude is very well liked by Hamlet (+24), which is unexpected (at least to us) since Hamlet suspects that his mother was involved in murdering King Hamlet.

Figure 2: The above chart tracks how Gertrude's and Hamlet's sentiment towards one another changes over the course of the play. Hamlet's sentiment for Gertrude is denoted by the black line, and Gertrude's for Hamlet is marked by the opposite boundary of the dark/light gray area. The drastic change in *Act III Scene IV: The Queen's Closet* is consistent with the scene's plot events.

3.1 Peering into the Queen's Closet

To gain more insight into this mother-son relationship, we examined how their feelings towards one another change over the course of the play. Figure 2 shows the results of dynamic character-to-character sentiment analysis on Gertrude and Hamlet. The running total of Hamlet's sentiment valence toward Gertrude is tracked by the black line, and Gertrude's feelings toward her son are tracked by the opposite boundary of the light/dark gray area. The line graph shows a dramatic swing in sentiment around line 2,250, which corresponds to Act iii, Scene iv.

In this scene, entitled *The Queen's Closet*, Hamlet confronts his mother about her involvement in King Hamlet's death. Gertrude is shocked at the accusation, revealing she never suspected Hamlet's father was murdered. King Hamlet's ghost even points this out to his son: "But, look, amazement on thy mother sits" (3.4.109). Hamlet then comes to the realization that his mother had no involvement in the murder and probably married Claudius more so to preserve stability in the state. As a result, Hamlet's affection towards his mother grows, as exhibited in the sentiment jump from *-1* to *22*. But this scene has the opposite affect on Gertrude: she sees her son murder an innocent man (Polonius) and talk to an invisible presence (she cannot see King Hamlet's ghost). Gertrude is coming to the understanding that Hamlet is not just depressed but possibly mad and on a revenge mission. Because of Gertrude's realization, it is only natural that her sentiment undergoes a sharply negative change (*1* to *-19*).

3.2 Analyzing Shakespeare's Most Famous Couples

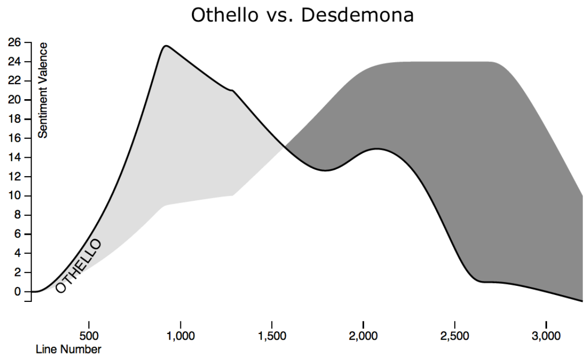

Figure 3: Othello's sentiment for Desdemona is denoted by the black line, and Desdemona's for Othello is marked by the opposite boundary of the dark/light gray area. As expected, the line graph shows Othello has very strong positive emotion towards his new wife at the beginning of the play, but this positivity quickly degrades as Othello falls deeper and deeper into Iago's deceit.

After running this automatic analysis on all of Shakespeare's plays, not all the results examined were as enlightening as the Hamlet vs. Gertrude example. Instead, the majority supported our already held interpretations. We will now present what the technique revealed about three of Shakespeare's best known relationships. Figure 3 shows Othello vs. Desdemona sentiment dynamics. We clearly see Othello's love for his new bride climaxes in the first third of the play and then rapidly degrades due to Iago's deceit while Desdemona's feelings for Othello stay positive until the very end of the play when it is clear Othello's love for her has become poisoned. For an example of a contrasting relationship, Figure 4 shows Romeo vs. Juliet. As expected, the two exhibit rapidly increasing positive sentiment for each other that only tapers when the play takes a tragic course in the latter half. Lastly, Figure 5 shows Petruchio vs. Katharina (from *The Taming of the Shrew*). The phases of Petruchio's courtship can be seen: first he is neutral to her, then 'tames' her with a period of negative sentiment, and finally she embraces him, as shown by the increasingly positive sentiment exhibited in both directions.

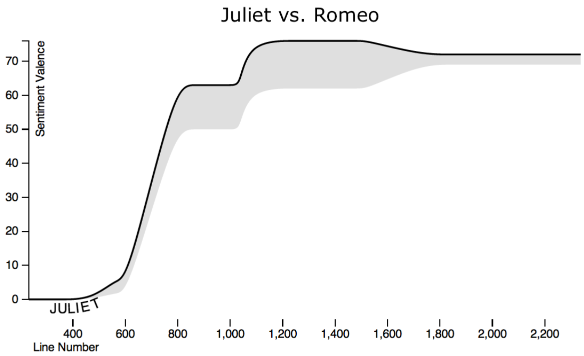

Figure 4: Juliet's sentiment for Romeo is denoted by the black line, and Romeo's for Juliet is marked by the opposite boundary of the gray area. Aligning with our expectations, both characters exhibit strong positive sentiment towards the other throughout the play.

Unfortunately, we do not have room in this paper to discuss further examples, but a visualization of sentiment dynamics between any pair of characters in any of Shakespeare's plays can be seen at *www.lehigh.edu/∼etn212/ShakespeareExplorer.html*.

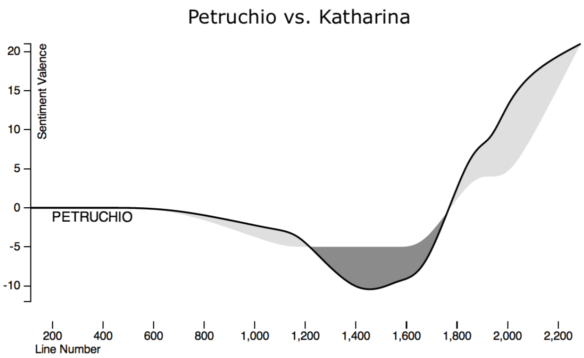

Figure 5: Petruchio's sentiment for Katharina is denoted by the black line, and Katharina's for Petruchio is marked by the opposite boundary of the dark/light gray area. The period from line 1200 to line 1700, during which Petruchio exhibits negative sentiment, marks where he is 'taming' the 'shrew.'

4 Future Work

While this paper presents experiments on just Shakespeare's plays, note that the described technique can be extended to *any* work of fiction written since the Elizabethan Period. The sentiment lexicon we used, *AFINN*, is designed for modern English; thus, it should only provide better analysis on works written after Shakespeare's. Furthermore, character-to-character analysis should be able to be applied to novels (and other unstructured fiction) if Elson and McKeown's (2010) speaker attribution technique is first run on the work.

Not only can these techniques be extended to novels but also be made more precise. For instance, the assumption that the current speaker's sentiment is directed toward the previous speaker is rather naive. A speech could be analyzed for context clues that signal that the character speaking is not talking about someone present but about someone out of the scene. The sentiment could then be redirected to the not-present character. Furthermore, detecting subtle rhetorical features such as irony and deceit would markedly improve the accuracy of the analysis on some plays. For example, our character-to-character analysis fails to detect that Iago hates Othello because Iago gives his commander constant lip service in order to manipulate him–only revealing his true feelings at the play's conclusion.

5 Conclusions

As demonstrated, shallow, un-customized sentiment analysis can be used in conjunction with text structure to analyze interpersonal relationships described within a play and output an interpretation that matches reader expectations. This character-to-character sentiment analysis can be done statically as well as dynamically to possibly pinpoint influential moments in the narrative (which is how we noticed the importance of Hamlet's *Act 3, Scene 4* to the Hamlet-Gertrude relationship). Yet, we believe the most noteworthy aspect of this work lies not in the details of our technique but rather in the demonstration that detailed emotion dynamics can be extracted with simplistic approaches–which in turn gives promise to the future work of robust analysis of interpersonal relationships in short stories and novels.

References

A. Agarwal, A. Corvalan, J. Jensen, and O. Rambow. 2012. Social network analysis of alice in wonderland. *NAACL-HLT 2012*, page 88.

Cecilia Ovesdotter Alm and Richard Sproat. 2005. Emotional sequencing and development in fairy tales. In *Affective Computing and Intelligent Interaction*, pages 668–674. Springer.

Alina Andreevskaia and Sabine Bergler. 2007. Clac and clac-nb: knowledge-based and corpus-based approaches to sentiment tagging. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, SemEval '07, pages 117–120, Stroudsburg, PA, USA. Association for Computational Linguistics.

Nathanael Chambers and Dan 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: Volume 2-Volume 2*, pages 602–610. Association for Computational Linguistics.

Micha Elsner. 2012. Character-based kernels for novelistic plot structure. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '12, pages 634–644, Stroudsburg, PA, USA. Association for Computational Linguistics.

David K Elson and Kathleen R McKeown. 2009. Extending and evaluating a platform for story understanding. In *Proceedings of the AAAI 2009 Spring Symposium on Intelligent Narrative Technologies II*.

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.

Michael Gamon. 2004. Sentiment classification on customer feedback data: noisy data, large feature vectors, and the role of linguistic analysis. In *Proceedings of the 20th international conference on Computational Linguistics*, COLING '04, Stroudsburg, PA, USA. Association for Computational Linguistics.

Soo-Min Kim and Eduard Hovy. 2004. Determining the sentiment of opinions. In *Proceedings of the 20th international conference on Computational Linguistics*, COLING '04, Stroudsburg, PA, USA. Association for Computational Linguistics.

Saif M Mohammad and Peter D Turney. 2008. Crowdsourcing the creation of a word–emotion association lexicon.

S. Mohammad. 2011. From once upon a time to happily ever after: Tracking emotions in novels and fairy tales. In *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 105–114. Association for Computational Linguistics.

P. Mutton. 2004. Inferring and visualizing social networks on internet relay chat. In *Information Visualisation, 2004. IV 2004. Proceedings. Eighth International Conference on*, pages 35–43. IEEE.

F. Å. Nielsen. 2011. Afinn, March.

Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing - Volume 10*, EMNLP '02, pages 79–86, Stroudsburg, PA, USA. Association for Computational Linguistics.

Stephen Ramsay. 2011. *Reading Machines: Toward an Algorithmic Criticism.* University of Illinois Press.

Antonio Roque. 2012. Towards a computational approach to literary text analysis. *NAACL-HLT 2012*, page 97.

C. Strapparava and A. Valitutti. 2004. Wordnet-affect: an affective extension of wordnet. In *Proceedings of LREC*, volume 4, pages 1083–1086.