

Adam Hammond

## Problematizing Literature with Digital Methods: *He Do the Police in Different Voices* and *The Brown Stocking*

Presented at HASTAC 2013, April 2013

In a recent *New York Times* feature, Matthew Jockers looked forward to a time when digital approaches like computer-assisted literary analysis would be “just part of the tool kit in the humanities, as in every other discipline.” Before that can happen—before the systematic, rigorous, unambiguous procedural methodologies of the sciences can become part of everyday life in humanities departments—some painful readjustments will be necessary. As it stands, many of what are called “traditional humanists” remain decidedly skeptical of computational approaches. In my field, English literature, the application of systematic, rigorous, unambiguous computational methodologies to the analysis of something as deeply, necessarily, and proudly ambiguous as literature, seems to many fatally wrong-headed. By seeking absolute, hard-and-fast answers in elusive, complex, polyvalent literary texts, many argue, computational analysis of literature misses the point, closing down meaning where traditional humanities approaches seek precisely the opposite, to open it up for debate.

The problem is not a new one—nor is the best response to this problem, which literary-minded digital humanists have long been making: namely, that computer-assisted analysis should be used not to disambiguate proudly ambiguous literature, but rather to find new ways of opening up literary debate. Stephen Ramsay rather bluntly argues, “A scientific literary criticism would cease to be criticism,” because “criticism”—the activity in which humanists engage—doesn’t seek *answers* (489). An algorithmic criticism worthy of the name, Ramsay argues, wouldn’t attempt to *supplant* the methods and assumptions of traditional literary scholarship, but to extend them—more quickly and to greater numbers of texts. In *Humanities Computing* (2004) Willard McCarty similarly argues that “modelling problematizes”—that computational approaches are primarily useful for upsetting some sacred assumption or critical orthodoxy, and making us begin the critical act anew. “As a tool of research,” McCarty says, “modelling succeeds

intellectually when it results in failure [...] This failure, in the sense of expectations violated, is fundamental to modelling" (26).

These are immensely attractive ideas. In practice, however, it can be difficult to find concrete examples of projects that genuinely fulfill the promise of Ramsay's algorithmic criticism or McCarty's problematizing modeling. Consider two prominent recent examples, both focusing on the 19<sup>th</sup> Century novel: Jockers's research on influence, the subject of the *New York Times* article; and Ryan Heuser's and Long Le-Khac work on "Quantitative Literary History." Both projects are premised on a "big data model," using algorithmic methods to quickly analyze vast quantities of novels. And both make bold claims: Jockers, for instance, that Austen was more "influential" than Dickens, Melville, Hardy, or Twain; Heuser and Le-Khac that the 19<sup>th</sup> Century British novel moves progressively from telling to showing.

These kinds of claims are disturbing to traditional literary scholars, for a few reasons. The first is that they present seemingly unambiguous truth claims. Grand totalizing claims are not unknown in humanities scholarship, of course—the problem is that most humanists have no idea how to argue with these particular grand claims. Without understanding how things work under the hood in the computational approach, it's unclear how to keep the conversation going, how to insert one's rejoinder. It's particularly unclear how humanists would use their primary tool of argument, analog, line-by-line, literary reading and criticism, to effect their rejoinder, since these projects aren't premised on human reading—on the way that individual readers respond subjectively to the literary texts. And so the conversation tends to stop, truth claim made and allowed to rest—or else it continues only among a subset of experts conversant in the particular digital methodologies.

My aim today is to describe two collaborative projects I have recently undertaken that aim to use computer-assisted analysis and other digital tools to carry out a true "algorithmic criticism." These projects begin from a stated aversion to truth claims; they aim rather to do what Willard McCarty calls "problematizing." Instead of trying to resolve literary dilemmas or interpretive cruxes, the projects aim to identify, highlight, and describe interpretive knots. Rather than telling readers some truth about authors, literary texts, or literary history, they aim to

provide readers with tools and resources to explore the cruxiness, the knottedness of literary language, without seeking to "disambiguate" this complexity in any definite direction.

The first project, *He Do the Police in Different Voices*, focuses on T. S. Eliot's *The Waste Land*. The overall aim of this project is to make a difficult modernist poem more accessible. Readers approach *The Waste Land* with a certain amount of dread—a dread that comes, I believe, primarily from the perception that there is, behind the surface confusions of the poem, some singular, coherent latent meaning that they can recover if they mine enough sources. This is a perception for which Eliot himself is largely responsible. Though the famous notes to *The Waste Land* were not his idea—they in fact began as nothing more than a sneaky way to fill some extra pages in the book version of the poem—they established the image of *The Waste Land* as a difficult poem comprehensible only to the immensely learned, or the incredibly studious. Eliot himself later came to regret these footnotes, and the vision of the poem they promoted. In *The Use of Poetry and the Use of Criticism*, he complained of the false perception among many readers, of

the author's having left something out which the reader is used to finding; so that the reader, bewildered, gropes about for what is absent, and puzzles his head for a kind of meaning which is not there, and is not meant to be there. (144)

What Eliot is saying here is what many humanists say when complaining about computer-assisted analysis: *The Waste Land* is not a poem with a determinate meaning, and should not be analyzed as such. To seek to revolve its dilemmas is to fundamentally to misconstrue the act of literary reading. It has always seemed to me that the best way of understanding this point in relation to the *Waste Land* to read the poem aloud. When you do that, you realize that *The Waste Land* is quite literally a poem made up of many irreconcilable voices—voices young and old, rich and poor, earth-bound and eternal, speaking all manner of languages and class dialects. One of the great breakthroughs in Faber's fabulously successful *Waste Land* iPad app was to introduce so many excellent readings, all of which bring this multi-voicedness to life.

Our website takes its name from Eliot's working title for the poem: He Do the Police in Different Voices. The aim of the website is to go a step further than a good reading like Guinness's—to help students understand *The Waste Land* by giving them a number of resources in which to explore the voices in the poem.

- Step 1: working with undergraduate students in "The Digital Text," a second-year English class at U of T, produce a "class version."
- Step 2: Working with computational linguistics Julian Brooke and Graeme Hirst of the University of Toronto, do an algorithmic reading of *The Waste Land*. Allow it to try to produce a definitive reading, using the usual computational means. But in the context of the website, make it just one of several "voices."
- Step 3: Crucially, build an interface for users to produce their own readings.

The website's "argument" unfolds in the interaction of its three parts. The "class reading" is not intended as definitive; indeed, one of its main purposes is to provoke disagreement—particular places where a reader feels our interpretation doesn't hold. The algorithmic reading serves the same function: it's meant both to problematize the "class reading"—to identify particular points where perhaps we should have indicated a voice switch, to spur new investigation in that manner—and also to provoke the rejoinders of readers who see limitations in the algorithmic reading of the poem. The key to the website is the third part, where visitors make their own readings.

Our second, ongoing project focuses on another modernist text, Virginia Woolf's *To the Lighthouse*. This time the project takes its name from "The Brown Stocking," the final chapter of Erich Auerbach's great philological work *Mimesis*. Auerbach, a German Jew who was living in exile in Turkey during WW2, when he wrote *Mimesis*, focuses in this chapter on a particular feature of modernist style he calls the "multipersonal representation of consciousness," which he finds exemplified in *To the Lighthouse*. At a time when totalitarian states were violently enforcing their singular readings of the world, Auerbach found political import in a modernist narrative strategy in which there is

not one order and one interpretation, but many, which may either be those of different persons or of the same person at different times; so that overlapping, complementing, and contradiction yield something that we might call a synthesized cosmic view. (549)

The principal technical device by which Woolf achieves the "multipersonal representation of consciousness" is what we now call free indirect discourse, a narrative strategy for blending voices. The technique is ubiquitous in *To the Lighthouse*, for instance in the last sentences in the novel:

It was done; it was finished. Yes, she thought, laying down her brush in extreme fatigue, I have had my vision. (III:13)

Although not enclosed in quotation marks, "Yes" and "I have had my vision" are direct discourse, giving the exact words that the character Lily thinks. Everything else is probably FID, mixing together the words of the narrator with those of Lily—in the phrase "It was done; it was finished," for example, the past tense of "was" clearly belongs to the objective narrator, while the words "done" and "finished" and the semi-coloned sentence structure seem to originate in the subjective experience of Lily. There is necessarily a great deal of uncertainty in FID, however. In this context, it's still not entirely certain whether "It was done; it was finished" is the narrator's rewriting of Lily's thoughts or whether it includes Lily's actual words; likewise, it's unclear whether the phrase "extreme fatigue" comes from Lily's subjective perspective or belongs to the narrator.

The *Brown Stocking* project began with a TEI tagging assignment in the Fall 2012 session of "The Digital Text." Each of the 160 students in the class was assigned a 100-150 word section from the first four chapters of the novel. Students were asked to mark up every instance of character speech for the following features:

- Type of discourse, whether direct, indirect, or FID
- Identity of speaker, from a list compiled of all characters in the novel
- Whether speech is aloud or silent

Since we acknowledged from the start that there are often multiple possible correct interpretations of a given passage, we assigned each passage to three or four students, giving

them freedom to mark it up as they saw fit. This tagging was graded, and each student was asked to prepare a 2-page assignment justifying their interpretation.

In the Spring 2013 session of the class, the assignment was repeated with another 160 students. Working on the last seven chapters of the novel this time, we modified our tagset slightly, and increased the length of passages to around 250 words.

This tagging produced the following, effectively a "reader's map" of *To the Lighthouse*, showing how students actually respond to the text. The diversity of their readings is perhaps the distinguishing feature—sometimes the students produced invalid readings, for instance attributing speech to a character who is clearly not involved in the scene, but more often the disagreements are revealing.

Our aim is again to develop this project into a website—one that would feature the above "reading maps," and also include an interface whereby visitors to the site can produce their own readings of the novel.

But it's another aspect of the project that is most relevant to this talk: using machine learning, we—myself and my colleagues in computational linguistics—want to see if we can develop an algorithm that can learn from the students' readings of the annotated chapters and produce a meaningful interpretation of FID in the rest of the novel.

Reading FID in *To the Lighthouse* presented a significant challenge for the student readers: the "reading map" certainly bears this out. The task of automatically detecting FID presents perhaps an even greater challenge to computational analysis. Literary scholarship has long said that ambiguity is a positive, indeed a defining trait of literary language—not a fault of style. Computational analysis, however, almost invariably treats ambiguity as a problem to be resolved—for instance treating disagreement between annotators as troublesome "noise." We see this project as an ideal test case to see whether it's possible to develop an algorithm that can—like a good reader of modernist literature—"exist in uncertainty." The algorithm's task will be the same as the students': to identify spans of character speech, classify them as direct, indirect, or free indirect discourse, then produce an interpretation of that passage (who is speaking, aloud or silent), and produce a judgment of the likely validity of the reading. Since we're asking the

algorithm to make and reflect on interpretations, the aim is not strict "disambiguation"—but even an algorithm whose only aim was to identify spans of FID would be treading deeply in the sea of ambiguity, since such spans are, by definition, passages where it is impossible to say definitely who is speaking.

We don't know yet if this will work. But if it does, we feel we'll have achieved something quite new. Unlike Jockers, Heuser and Le-Khac, we're not working on the "big data" model of many lightly annotated texts; instead, we're working from very few, highly annotated texts. Beginning from human responses—conflicting, "messy" responses to a complex literary work—and expanding outward, we're building a computational literary analysis that we believe will yield meaningful results—and also one that traditional literary scholars will be able to understand and respond to. We certainly feel that Auerbach would understand it. In the "multipersonal" modernist novel, he said, there is

not one order and one interpretation, but many, which may either be those of different persons or of the same person at different times; so that overlapping, complementing, and contradiction yield something that we might call a synthesized cosmic view. (549)

Our project, *The Brown Stocking*—built from the overlapping, complementing, contradicting responses of human readers—aims at precisely such a synthesized cosmic view.

## Works Cited

Auerbach, Erich. *Mimesis: The Representation of Reality in Western Literature*. 1946. Trans.

Willard R. Trask. Princeton, NJ: Princeton UP, 1953.

Eliot, T. S. *The Use of Poetry and the Use of Criticism*. 1933. London: Faber, 1964.

Jockers, Matthew. Qtd. in Steve Lohr, "Dickens, Austen and Twain, Through a Digital Lens." *New York Times*. 26 January 2013. Web.

McCarty, Willard. *Humanities Computing*. New York: Palgrave Macmillan, 2005.

Ramsay, Stephen. "Algorithmic Criticism." *A Companion to Digital Literary Studies*. Eds. Susan Schreibman and Ray Siemens. Oxford: Blackwell, 2008.