

Fleshing Out Models of Gender in English-Language Novels (1850 – 2000)

Jonathan Y. Cheng^a

^a*University of Nebraska, Lincoln*

ARTICLE INFO

Peer-Reviewed By: Hoyt Long, Laura Mandell

Article DOI: 10.22148/001c.11652

Dataverse DOI: 10.7910/DVN/QUGW8V

Journal ISSN: 2371-4549

ABSTRACT

Distant readers have used predictive modelling to study the strength of the relationship between characterization and binary notions of gender. This essay builds on that research, shedding light on several historical trends concerning anatomical description and its relationship to gender. Some of the evidence suggests that bodily language has long played a larger role in configuring fictional women than it did for fictional men. Other evidence implies that bodily characteristics were increasingly bifurcated along a gender binary, reflecting how characters are more and more physically sorted along a feminine-masculine axis. Taken altogether, this essay unpacks a suggestive correlation: a growing aspect of characterization was increasingly imbricated in heteronormative discourses. By weighing the discrepancies between the evidence presented in this essay, and that of its predecessors, this essay will ultimately suggest that disaggregating statistical models can unfold patterns of literary change that would otherwise remain suppressed.

Introduction

Recent computational work has analyzed the significance of gender in characterization, investigating whether character descriptions are often sorted along a feminine-masculine axis. Matthew Jockers and Gabi Kirilloff, for instance, tabulate pronoun-verb pairings, exploring the connection between characters' actions and their gendered representation in nineteenth-century novels.¹ They show evidence of a stable relationship between gendered pronouns and the verbs they perform.² Ted Underwood, David Bamman, and Sabrina Lee explore a broader range of words used in characterization, measuring the difference between the words describing fictional men and those describing fictional women. Their evidence shows that the implicit differences between gendered characters gradually becomes less clear over the past two centuries. While the former study argues that characters' actions reveal gender's steady prominence, the latter research posits that those overarching gender divisions are broadly diminishing.³ But these seemingly disparate arguments should not be taken as contradictory. Rather, these varying

conclusions should reinforce a more complicated sense of how “some forms of gender differentiation...are declining while other forms...are on the rise.”⁴ To further explore the varying degrees and modes of gender differentiation, I employ quantitative methods to investigate gender’s role in the description of characters’ bodies.

It is difficult to overstate the significance of the body in scholarly accounts of gender. Laura Mulvey famously uses the representation of women’s bodies to articulate her well-known thesis about the male gaze in cinema. The cinematic representation of women’s bodies, for Mulvey, reflects how women are often restrictively configured to be “looked at and displayed, with their appearance coded for strong visual and erotic impact, so that they can be said to connote to-be-looked-at-ness.”⁵ That thesis has extended well beyond cinema, shedding light on various heteronormative strategies, physically representing women as sexual objects for a male viewer. In a related way, Butler uses that sort of reductive embodiment to pursue her well-known critique of the sexed body as an asocial phenomenon. Anatomical representations of the body, for Butler, reflect how intensely gender norms can often be presented as a purely biological configuration. She asks readers to critically rethink how the body can come “into being in and through the mark(s) of gender? How do we reconceive the body no longer as a passive medium?”⁶ By reframing the sexed body as a socially constructed object, Butler’s framework has been foundational to a wide range of feminist and queer critiques of heteronormative strategies. And this is all to say that scholarly work on gender has often centered on representations of the body, because anatomical description has historically reflected one of the most rigid applications of gender norms.

This essay contributes to that ongoing research, analyzing anatomical characterization throughout a collection of about 13,000 English-language novels.⁷ By producing a model of gender based solely on characters’ physical features, I explore the extent that anatomical description is defined along a feminine-masculine axis. And I pursue two central claims that complicate existing models of character and gender. The first is that bodily description becomes an increasingly prominent aspect of characterization.⁸ My analysis shows that an increasing proportion of character description is devoted to the bodies of both fictional men and women. Moreover, there is evidence that those details have long played a larger role in the

representation of women. Secondly, those characteristics have increasingly been deployed along gendered lines until only very recently. As we move towards the twenty-first century, men and women are increasingly embodied using different words. Even seemingly innocuous bodily features such as “sight” or “clasping one’s hands” function as signs of gender.

By showing evidence that gender remains a prominent force in anatomical characterization, I complicate Underwood et al.’s argument that gendered differences were overall becoming less rigid. As I mentioned earlier, they model the relationship between character and gender more broadly, looking at a wider swath of words used in characterization. And, to be fair, they do concede that certain aspects of character might have been gendered to different degrees.⁹ My model leans on that anecdotal intuition, yielding evidence that the anatomical and sexual dimensions of gender have instead become increasingly conventional until more recently. But I don’t want that evidence to seem like it displaces Underwood et al.’s findings. We want to see both the overarching trend as well any significant deviations. Instead, I’ll argue that their analysis is a little too quick to aggregate the strength of gender stereotypes into a single model. Ultimately, I’ll conclude that disaggregating such models can unfold suppressed patterns of literary change.

For now, it is enough to say that my two claims form a suggestive parallel: as the body becomes an expanding aspect of characterization, that dimension is increasingly organized along gendered lines. By showing that an increasingly conventional mode of representing characters is also becoming more and more gendered, I demonstrate how quantitative approaches to gender benefit from modelling particular subsets of characterization. Moreover, on a formal register, this work also adds to scholarship on the novel, providing new evidence of the practices used in characterization.

Methods and Limitations

In order to gather the words describing each character, one first needs to tell the computer which words are describing which characters. Bamman et al.’s BookNLP

pipeline has worked well for similar problems, so I modified their process for my own purposes.¹⁰ The preprocessing narrative is as follows:¹¹

First, I implement a process called coreference resolution to locate character names in each text and identify any expressions that also refer to the same fictional person. Image 1 provides an example output of this process and reflects what coreference resolution produces:

tid	entity	main_name
6	Elizabeth	Elizabeth
16	She	Elizabeth
27	her	Elizabeth

Image 1. Output of coreference resolution. TID reflects location in text. Pronouns and proper nouns are connected to single entity.

In this example, we have instances of the pronouns “she” and “her.” Coreference resolution uses several linguistic rules to determine that those pronouns refer to the character “Elizabeth,” treating each of those pronouns as different references to the same person. As a result, we can catalog several instances when a character gets mentioned in a text, and we can do this for each character in each text (regardless of whether they are mentioned via proper noun or a pronoun). But now we want to find any words connected to each reference to a character. Doing that allows us to record which words have been used to describe which person.

So secondly, to that end, I use a process called dependency parsing to identify a variety of words connected to each character-reference. This is a process that establishes relationships between words in each sentence, allowing us to determine which words are connected to which characters.¹² I do this for each sentence containing a character coreference. Following Bamman et al.’s model of character, I extract all of the actions a character performs, the actions that they’re the object of, any adjectives modifying them, and any nouns they govern (such as body parts, like “his hands”). Moreover, for the purposes of this inquiry, I add onto their character model by extracting additional words that physically describe each character. Whenever a character’s body part is mentioned, “his hands” for instance, I also gather the verbs and adjectives modifying their bodily features, such as “his hands

grasped” or “her *blue* eyes.”¹³ As a result, I get a frequency table of the words used to describe each character in each novel. In effect, this process tabulates the same words as Bamman et al.’s model, but it additionally collects the words attributed to their physical features. Image 2 is an example output of using dependency parsing to locate words connected to each character:

Index	name	lemma	relation	count	file_id
957	elizabeth	sister	poss	31	13420
958	elizabeth	mother	poss	23	13420
959	elizabeth	father	poss	12	13420
960	elizabeth	friend	poss	10	13420
961	elizabeth	spirit	poss	10	13420
962	elizabeth	eye	poss	9	13420

Image 2. Output after coreference resolution and dependency parsing to extract words describing a character.

This approach is, however, just one way to explore physical description. In practice, this method analyzes the *anatomical* features a character possesses, the actions those features perform, the actions that those body parts receive, and any adjectives attributed to them. And this definition privileges a material definition of the body, as it sometimes suppresses the metaphorical or psychosomatic instances of bodily language. If we take two passages from Stephen King’s *Salem’s Lot*, for instance, the tradeoffs of this approach become more clear.¹⁴ In some scenes, a young boy’s fear gets rendered almost entirely through anatomical description:

Blood thudded in his temples. His hands were cold... Terror like hot iron leaped into Danny’s chest. Wires seemed to have run up his legs... Danny held his brother’s hand and they began to walk. His legs felt as if they were made up of ten thousand pencil erasers. His knees were trembling. (Emphasis is added to show the kinds of words extracted using this method.)

This method loses some of King’s more metaphorical expressions, but it does capture how Danny Glick’s “cold hands,” “trembling knees,” and “thudding blood”

are used to represent this character's fear. That's not at all a bad approximation of Danny's physical terror. But let us consider another example of physical description that completely gets suppressed in this approach:

“No, Danny. Really. Can't you feel it?” Danny stopped. And in the way of children, he did feel something and knew they were no longer alone. A great hush had fallen over the woods; but it was a malefic hush. Shadows, urged by the wind, twisted languorously around them.”

Even without mentioning any of Danny's physiological features, the scene's suspense relies on the reader's sense of Danny's physical and mental dread. The anatomical approach employed here will only indirectly address the psychological aspects of bodily description (e.g. her skin crawled, his spine tingled), and it will only include these instances when signaled by an anatomical feature. This approach, however, has the advantage of consistently extracting moments of anatomical description, allowing us to raise questions about how the material body is used to represent gender.

This essay also operationalizes previously employed assumptions about gender. Like previous studies on character, this essay treats gender in a Butlerian sense, as conventional roles that people –fictional or real– were expected to publicly perform in order to remain legible within heteronormative contexts. As such, authors and characters are labelled as either feminine, masculine, or unknown. Those labels, of course, do not capture the complexity of gender identity, but they can be used to help illuminate the socially pervasive and stereotypical representations of gender.¹⁵ Additionally, given that the history of sexed bodies is often associated with the enforcement of rigid gender norms, those labels are appropriate for exploring that particular history. Future research can (and should) complicate this model, and I provide a starting point for those critiques.

Like previous research on gender's significance in characterization, my method comes with its own technological challenges. In order to separate characters from each other and assign them gender identities, I use Lincoln Mullen's *Gender* package to label those names with a grammatical gender.¹⁶ Mullen's package uses U.S and North Atlantic census data to accurately predict gender of first names, accounting for shifts in time and geographic location. The problem with using proper names to

identify character, however, is that characters referred to by generic nouns are excluded, such as “the baker.” This pipeline attempts to account for this by including characters signaled by stereotypically gendered nouns, such as “the queen” or “the father,” but this does not comprehensively account for many of the generic nouns used to produce characters. Moreover, this study also does not provide any robust solutions to the first-person narrator problem. The pronoun “I” does not consistently connote a clear gender identity, so anatomical descriptions of first-person characters are excluded from this essay. As such, there are certain kinds of characters whose physical features will not be counted in this study.

Embodying Fictional Men and Women, 1850 – 2000

This essay analyzes characters’ physical descriptions throughout 13,000 English-language novels, and their publication dates range from 1850 to 2000.¹⁷ Most of this corpus was developed by the Chicago Text Lab, covering the period from 1880 to 2000. Other novels were curated by the Stanford Literary Lab, covering the period from 1750 to 1900. However, most of the novels span from 1850 to 2000, so that is where I will focus my analysis. This sample of novels will largely reflect the purchasing practices of academic libraries. This can certainly tell us a lot about the patterns of physical description underlying the literary history studied in universities. But this collection of texts will exclude representations of the body that were pervasive in other modes of publication, such as magazines and newspapers. There are a few other aspects of these corpora that can help readers raise further questions and help contextualize the findings of this article. To that end, I’ve provided an appendix with some relevant metadata about the corpora further below.¹⁸

As we continue to develop more robust samples of literary history, this sample of 13,000 novels can, at the very least, make it possible to explore some of the more prominent trends in characterization. For example, we have several different theories about the extent that characterizing men and women relied on anatomical description, and there is an unclear sense of how best to navigate those disparate expectations. On the one hand, we have some intuitions about the overall space afforded to fictional women. Underwood et al. have shed light on a discouraging story about the steadily diminishing space afforded to female characters over the

past two centuries.¹⁹ That diminishment might lead one to expect that the physical description of women would similarly decline. On the other hand, there is also a sense that bodily language was steadily becoming more important to fiction. Ryan Heuser and Long Le-Khac have shown that body parts were becoming more and more conventional in nineteenth-century British novels.²⁰ It would also be reasonable then to expect more physical description for both female and male characters. But then, if we follow feminist research about the intense objectification of women's bodies, it's also reasonable to expect that an increasing amount of physical description gets attributed to women. On top of that, there are also hypotheses about increasing masculine anxieties regarding physical appearance, which could lead one to expect that the body was becoming more central to describing men.²¹

We can start to untangle these various hypotheses by asking a simple question: is the percentage of characterization, that is bodily description, different between fictional men and women? Exploring this question helps us see whether the body becomes a larger proportion of characterizing men or women. Let's start by taking the number of words physically describing female characters and divide that by the number of *all* words describing those characters. Then we perform that calculation for each year. In effect, we're just plotting the proportion of words describing women that are about their bodies.²² The same calculation is also performed for the fictional men. Once we have those proportions, for each year, and for both men and women, we can compare those proportions to see if there are any clear gendered differences. That comparison is plotted below in Figure 1.




Figure 1. The percentage of words used in characterization connected to characters' physical features. Separated out by character gender.

Two clear long-term patterns emerge. First, *body language becomes a growing aspect of all characters as we get closer to the twenty-first century*. Over time, more and more of the words used in characterization concern the body. If we want to describe the direction and strength of that trend, we can measure their correlation and measure the rate of that change. Those descriptive statistics are provided below in Table 1.

Correlated with Publication Year	r	p-value	slope
Percentage of Physical Description of Men	0.88239	2.2e-16	0.060010
Percentage of Physical Description of Women	0.80049	2.2e-16	0.029894

Table 1. Correlations between publication year and proportion of physical description. P-value shown for statistical significance. Slope shown to provide sense of rate of increase. Separated by character gender.

Given that the correlations (r) are both above 0.8, we have reason to suspect that there is a very strong relationship between the two variables, suggesting that the proportion of bodily characterization does grow as we move forward to the present.²³ And, given that we're looking at small slopes, we're also looking at a gradual (yet steady) change in characterization. Even though that rate of change seems marginally incremental, this adds up over 150 years. For fictional men in this corpus, the percentage of physical description grows from 7% to about 15%. So, when it comes to characterizing men, the amount of physical description has nearly doubled. For fictional women, that percentage grows from around 12% to about 17%, suggesting that nearly one-fifth of the words about women have become about their bodies.

Readers will also note that the magnitude of this change seems to vary across the two centuries. More specifically, there appears to be a rapid growth in the physical description of men from the 1880s to the 1900s. But I'd caution that this jump might be an artifact of the data, rather than a real cultural effect. If we look at the rate of growth before that cliff, the incline looks much more gradual (For men: 0.040943, For women: 0.000438). Given the strength of the correlation, we're still looking at a steady growth in bodily description, but we'll want some clarification about that spike at the turn of the century. Instead of a literary or social explanation, one potential reason for that uptick is that this is where the data shifts from the Stanford corpus to the Chicago one. The Chicago corpus begins right at 1880, which is around when we see a sudden increase in an otherwise gradual pattern. If this is indeed the case, then there's reason to doubt a dramatic shift at the close of the nineteenth century, leaving us with a story of gradual change.

So, on the one hand, Figure 1 reflects a well-known story: the body becomes a growing aspect of producing characters. These two slopes provide further evidence of Heuser's and Le-Khac's claim that the body becomes steadily more important in fiction. This isn't to claim that the body was becoming *the most* important aspect of character. Andrew Piper, for example, has provided evidence that signs of interiority (such as acts of thinking, reflection, and sense perception) are a distinctive feature of female characters in nineteenth-century novels written by women.²⁴ And this isn't to say that the body was only growing in importance in the past two centuries. There's a lot of evidence to the contrary.²⁵ This is just to say that, since the 1850s,

more and more space for describing characters has been allocated to their physical features, appearances, and gestures.

“So why did the body become increasingly prominent in novels?” While I will not pose any causal hypotheses for this trend, literary scholars have already produced a wide range of literary and social explanations. On a literary register, scholarly accounts of the novel tend to associate participation in realism with a commitment to minute physical details, such as bodily characteristics. Fredric Jameson, at the very least, sees the construction of a “secular or bourgeois body” as one of the central projects of realism in the nineteenth century.²⁶ Configuring the body, for Jameson, was necessary, because physical sensations and sensory perceptions produced an affective immediacy, allowing realism to feel urgently present. Which is to say that Figure 1 *could* reflect theories about changing literary ideals.

On a more cultural register, however, other scholars correlate the body’s growing prominence to a variety of cultural shifts. The breadth of research on this subject is difficult to cover. Catherine Gallagher, for example, approaches the body in literature as an economic phenomenon. Her work focuses on the tendency of market discourses to emphasize physical expressions as signs of economic happiness, seeing its resonance in the characters of Victorian novels.²⁷ Focusing more on scientific discourses, Lucy Hartley correlates an increasing attention to facial expressions to professional and social investments in physiognomic debates. Hartley’s research describes how interpretive debates about facial expressions figured their way into literary practices.²⁸ There are, of course, several other ways to tell the story of the body and characterization. While Figure 1 cannot verify each one of those different hypotheses, it does align with a general sense that the body has become more important over the past two centuries. And, perhaps excitingly, this picture also reflects how that growing prominence has yet to conclude.

But there’s another significant story to come out of Figure 1: *physical description consistently tends to be a larger proportion of characterizing women than men*. In fact, as we move from the 1850s to the 2000s, at no point is the proportion of bodily description greater for men. The pattern remains intact even when physical characteristics were becoming more prominent for all characters. This evidence suggests that the body has historically played a larger role in representing women. That’s an important detail about literary history, because it underscores the extent

that women's physical descriptions are imbricated in gender discourses. Feminist scholars have, of course, already captured important parts of this story. Butler's argument about gendered bodies, for instance, hinges upon her claim that there are "cultural associations of mind with masculinity and body with femininity."²⁹ The evidence in Figure 1 doesn't completely verify a rigid alignment of men with the mind and women with the body. Rather, it is congruent with the latter part of Butler's claim, showing that the representation of women tends to rely more heavily on bodily description.

Are those two patterns, however, broadly consistent across all novels? I have shown that fictive women are, proportionally, configured with more body language than men. But it is entirely possible that certain kinds of novels, or those written by a particular demographic, embody men and women to different degrees. Along Mulveyian lines, for instance, one could reasonably suspect that books by men have a stronger tendency to focus on women's bodies. Moreover, given that my corpus does contain more male authors, we'll want to see whether the patterns in Figure 1 persist among different demographics. We can explore this possibility by separating the proportions in Figure 1 by author-gender. In other words, we're seeing if the proportion of physical characterization dramatically changes when we separate books written by men and those by women. When we perform that analysis, we are confronted with a broader historical pattern. That analysis is plotted below in Figures 2 and 3.




Figure 2. The percentage of body language used in characterization, in books by women.




Figure 3. The percentage of body language used in characterization, in books by men.

There are some very subtle differences between those two images, but they broadly tend to mirror the patterns we saw earlier in Figure 1. We can again calculate correlations to describe the direction and strength of these trends, and we can also calculate slopes to see any major difference in rates of change.

Correlation with Publication Year	<i>r</i>	p-value	slope
Percentage of Physical Description of Men in Books by Women	0.7349	2.2e-16	0.061635
Percentage of Physical Description of Women in Books by Women	0.3475	2.2e-16	0.029894
Percentage of Physical Description of Men in Books by Men	0.734	2.2e-16	0.058309
Percentage of Physical Description of Women in Books by Men	0.5737	2.2e-16	0.034156

Table 2. Correlations between publication year and proportion of physical description. P-value shown for statistical significance. Slope shown to provide sense of rate of increase. Separated out by author and character-gender.

What immediately stands out in Table 2 is that the direction of these trends is the same, but the strength of the correlation between publication year and percentage of physical description wanes. Some of that instability is to be expected, as we have split a larger sample of novels into two smaller samples. However, the correlation has dropped a fair amount for female characters written by women (a drop from 0.80049 to 0.3475). This could be due to the decrease in sample size. But that sharp decline could also suggest that, in female authored novels, there is a less poignant

relationship between historical progression and the amount of physical description attributed to women. Moreover, Table 2 could also suggest that the growing amount of prose on women's bodies we saw earlier might have been more prominent in books by men. The slightly higher rate of change in books by men (0.0341 to 0.029894) could also be evidence of this trend.

But the margins are subtle enough that we ought to resist immediate interpretation. At the very least, it seems safe to say that both graphs appear to reflect the two overarching patterns exhibited in Figure 1. This isn't to suggest that female and male authors embody their characters in the same way. I will soon present evidence that argues otherwise. Rather, it reflects how the growing prominence of physical description, and its relatively greater importance to characterizing women, seems to flow more broadly across the novels contained in this corpus.

If those patterns prove durable, they lead to several interesting questions about the embodiment of men and women. If we take the Butlerian line, one question immediately stands out, "to what extent were women culturally associated with the body?" I have shown that fictive women are consistently composed of more physical description than fictive men. That's a start, as a consistent 5-10% margin isn't negligible evidence. Moreover, this evidence also aligns with previous research, so multiple approaches lend weight to that association. Underwood et al. has shown how the words "*lips, eyes, face, and voice* are all associated with women...*body* itself, in fact, is gendered feminine" (Underwood et. al). Those anecdotal examples present individual instances where body parts appear to align with female characters. But, on the other hand, we will want other forms of social and quantitative evidence to see if the pattern resurfaces in different genres. If, for instance, the pattern persists in poetry and other forms of fiction, then it would be much more difficult to argue against an association between femininity and the body. Realistically, more research is needed. For now, we have evidence of the gendered distribution of bodily description. That's exciting, because we can now start to ask more direct questions about the gendering of the physical characteristics themselves.

Gendering the Body

We have seen evidence that bodily description is allocated an increasing amount of space in characterization. We've also seen that the body is consistently a greater proportion of characterizing women than it is of men. But neither of those sheds light on the extent that the physical features themselves are allocated along gendered lines. In short, it's unclear whether fictive men and women are embodied using different body parts, gestures, and attributes. Moreover, it's unclear whether those differences are allocated to a rigidly consistent degree. And, because we have seen that the body becomes an increasingly prominent aspect of character, we'll want a sense of the extent that bodily features were gendered across the timeline. To explore that history, we can start by asking two simple questions: Were fictive men embodied in different ways than fictive women? And does the strength of that difference change as we move towards the twenty-first century?

In order to explore those questions, binary classification methods have proven effective at modelling the strength of social categories.³⁰ At its core, this method tests to see how well individual samples can be sorted into two related categories. In our case, we want to test whether bodily nouns, verbs, adjectives, etc. are consistently attributed to either fictive men or women. This means, first, taking multiple random samples of characters from each decade, tabulating the words that physical describe them. Then, each of those words is labelled with the character's grammatical gender. By showing a model many of these labelled words, we train it to develop a stereotypical sense of what attributes constitute a stereotypically "feminine" or "masculine" body in that decade. Finally, we instruct the model to use its sense of gendered bodies to make gender predictions about characters it hasn't seen yet. If, for instance, physical characteristics are predominantly distributed along gendered lines, this will allow the model to consistently make accurate gender predictions. On the other hand, if body language and gender are generally unrelated, then the model will be less capable of accurately inferring a character's gender from that language.

What can we expect to see? Once again, it's hard to say. On the one hand, previous research has shown that several differences between fictive men and women have diminished over time. Underwood et al. uses a similar classification method to show

that “the implicit differences between masculine and feminine characters get steadily harder to discern from the middle of the nineteenth century to the beginning of the twenty-first” (Underwood et al.) Moving from that evidence, one could reasonably assume that the same is true for bodily characteristics. But, anecdotally, that feels like a hard sell. If we restrict characters to bodily imagery, intuitively, it’s hard to imagine gender as a diminishing force. Though, with Underwood et al.’s evidence, one can hope that those intuitions give way to subtle pushbacks against embodied notions of gender.

We see two overarching patterns below in Figure 4. Up until the 1960s, the model gradually has an easier time inferring the gender of characters from their physical characteristics. During that period, the model’s percentage of correct gender predictions rises from about 76% to about 83%. What this suggests is that words describing the body are becoming increasingly bifurcated along normative gender lines. After the 1960s, however, the percentage of correct predictions drops back down to roughly 77% as we reach the 2000s. This would suggest that the model was less and less able to use body language to correctly infer a character’s gender. This would also suggest that the bodily differences between fictive men and women has only more recently begun to diminish.




Figure 4. Model’s accuracy at inferring gender from bodily descriptors. 7200 characters drawn from two corpora of novels, 225 men and 225 women sampled in each decade.

More quantitative evidence will, of course, be needed to feel confident about these figures, but this approach takes a few steps in order to test the strength of this pattern. This figure was produced by running fifteen different models within each decade. Each of those models is produced by randomly selecting 450 characters at a time (225 men and 225 women), classifying them using the top ten percent of the most commonly occurring words in each sample.³¹ These steps help ensure that we're making fair comparisons between uniform samples. This winnowing strategy comes at the cost of ignoring less frequent physical descriptions. For example, when it comes to describing eye-color, the adjective "hazel" is less common than "black," "brown," or "blue," so it is often excluded from each model. These sparse features can certainly be significant signs of gender, but the benefit of this approach is that it analyzes gender's prominence within pervasively deployed physical features, such as having a "nose" or being attributed with "brown" hair.

The patterns in Figure 4 complicate Underwood et al.'s evidence of a steady decline in the gendering of characters, and this discrepancy merits further unpacking.³² On the one hand, this pattern does align with their research in a clear way. From the 1960s forward, we see that the models in Figure 4 become decreasingly accurate at inferring a character's gender from body language, suggesting that the representation of men's and women's bodies has steadily become less different over time. In this regard, both studies do provide evidence of diminishing gender divisions. Before 1950, however, this picture tells a different story. The classifier's improving accuracy suggests that embodied gender differences were instead becoming steadily sharper. Taken altogether, the evidence I present suggests that a growing aspect of character was also becoming increasingly gendered.

This raises an interesting question for distant reading: how do we want to interpret the discrepancy between these two pieces of quantitative evidence? I would argue that Underwood et al. are too quick to aggregate "the strength of gender stereotypes" into a single unified trend. This sort of "aggregation bias" has become a concern for data scientists, because a generalized model can sometimes draw misleading conclusions about its various subgroups. In the field of medicine, for instance, how a condition behaves in the overall population can be very different from the way it behaves in specific demographics, and we don't want doctors to overlook those subpopulations when diagnosing patients.³³ In the field of literary criticism,

however, a generalized model can be unfolded to explore complicated stories of historical change. Disaggregating Underwood et al.'s model, for example, sheds light on an aspect of representation that has a particular historical relationship to gender. By demonstrating how those particularities impact the way we draw historical trendlines, we can produce new evidence of literary change through a critique of existing models. Taken altogether, we're not dealing with contradicting pieces of evidence where one model subsumes the other. Instead, we're looking at two different scales of analysis, uncovering related patterns that have, at times, moved in opposite directions. By unfolding generalized models, we can concretize different cultural perspectives by comparing their effect on the historical record.

If we set aside larger methodological reflections for now, there are several scholarly accounts of literary history that help to explain gender's growing prominence in bodily description. Kate Flint, for instance, explores the expanding popularity of natural philosophy in the late-eighteenth and early-nineteenth centuries. As more attention gets paid to the eye's anatomy, new uncertainties emerged about the functional symbolic capacity of the human eye. That uncertainty lead to explicitly gendered theories about how the eye operated differently for men and women. Popularizers of natural philosophy, such as Joseph Turnley, would claim that "the eye of man is the most *firm*; woman's the most *flexible*...Man's *surveys* and *observes*; woman's *glances*" (emphasis added).³⁴ These sort of physiological arguments were central to configurations of biological sex and could be one explanation for Figure 4.

But I should also impress that scientific discourses were not the only aspect of culture that embodied gender divisions. Peter J. Capuano reminds us that etiquette manuals are constantly preoccupied with the bodily behaviors of women. For Capuano, that preoccupation about women's hand gestures aligns with the class conflicts emerging out of Britain's industrialization.³⁵ In those conflicts, surveilling women's physical manners was one strategy for landowners to differentiate themselves from the rising professional class. According to Capuano, that strategy (and its subversion) find its way into the social drama of Victorian novels, distinguishing wealthy women from other people.³⁶ Both of these scholarly accounts describe how physical characteristics, whether they be eyes or hands, were becoming

a significant site of gender differentiation. And the evidence displayed in Figure 4 aligns with those claims while extending their effects to the mid-twentieth century.

But do books by men and women gender bodily characteristics to the same degree? We saw in Figure 4 that inferring gender from characters' bodily features gets easier until around the 1960s. Exploring whether that pattern varies throughout different groups of novels can tell us a lot about that historical arc. One can imagine, for instance, a historical arc where books by men tend to deploy bodily characteristics along gendered lines more often than books by women. There are hints of that kind of story in Underwood et al.'s research. When they compare the degree of gender differentiation in books by men and in books by women, they find that the divide between fictive men and women is *slightly* more apparent in stories written by men. Intuitively, along Mulveyian lines, it seems reasonable to expect a similar story when we narrow our analysis to just bodily description. The following picture, however, underscores how dramatically apparent embodied divisions are in male authored novels. When we compare the difference between fictive men's and women's bodies in only books by men, and the difference only in books by women, the models trained on books by men often tend to be more accurate (by 7% on average).³⁷ This is especially true, during the twentieth century, when several of those models report more than 90% accuracy, suggesting that there is a sharp divide between fictive men's and women's bodies in novels by men.




Figure 5. Model's accuracy at inferring gender from bodily characteristics. Those accuracies are separated out by author gender. Classification parameters same as in Figure 4.

There's another curious detail about this picture that will require future research to fully map out its implications. The models trained on books only written by men, and those trained on books only by women, are more accurate at inferring gender than the initial models trained on both groups of authors. On average, the initial models were 80% accurate at inferring gender (Figure 4). Whereas models trained only on books by men were 87% accurate while those trained only on books by women were 83% accurate (Figure 5). Those aren't large disparities. We should also remember that it is unclear whether a 3-7% difference in accuracy is culturally significant. But those differences could suggest that discerning gender divisions becomes slightly easier when we isolate the writing tendencies of varying groups of authors. One could imagine, for instance, that embodied gender differentiation becomes discernable throughout books by men in some ways (ex. those books frequently configuring women with lips and rarely so with men), but then that differentiation becomes apparent in books by women for a different set of reasons (ex. these books often mention men as having beards and not so with women). It's perhaps possible then training models on both groups of authors diluted some of those gender divisions, resulting in slightly lower accuracies. So, if the patterns in

Figure 5 prove consistent, it would raise several questions about the kinds of novels that men and write and their different relationships to gendered body language.³⁸

Because we have seen that embodied gender differentiation is more apparent in books by men than in those by women, there's a very tempting claim to advance. Male authors have a stronger tendency to institute a heteronormative representation of characters' bodies. If we want to fully map out that observation, we will need a better understanding of the variabilities reflected in Figure 5. For example, when it comes to books by men, we will need to explain why the gendered differences between fictive men's and women's bodies become increasingly apparent after 1900. While the gendered differences between fictive men and women become increasingly sharp during the twentieth century, this is less the case during the mid-to-late nineteenth century. The gendering of the body just wasn't that much more apparent in books written by men than in those by women. Advancing this claim, will require a better understanding of that gradual transition. On the other hand, when it comes to books by women, the task of interpretation is a bit more difficult. When the task of gender inference was becoming easier in books by men, they were erratically declining for books by women (occasionally falling below 75% accuracy). It isn't immediately clear why the division between fictive men's and women's bodies became more diffuse. One might point to the way that pattern aligns with the rise in second-wave feminism during the mid-twentieth century, and it's focus on women's bodies and sexuality. But more research is needed to understand this important part of the story.

Transformations of the Gendered Body

Setting aside broad cultural hypotheses for now, how can we analyze the historical transformation of individual bodily characteristics? The previous analyses reflect the fluctuating degree to which the body was gendered over time. But we know that notions of embodied gender shift in qualitative ways as well. In order to explore those transformations, we need a way of visualizing the gendering of particular features across the timeline.

There are many ways to do this. Jockers and Kirilloff have shown that a good first step can be to unpack the classification models, noting which factors lead to the most correct gender inferences. In effect, we're looking at which physical descriptions, on average, were most exclusively attributed to one gender category. Since this essay produced fifteen different models of physical description for each decade, we can see bodily descriptors that all fifteen models agreed had a gendered distribution. While this selection process is a more conservative, we can be surer that the resulting physical features tend to be exclusively attributed to either female or male characters. When we perform that analysis, we get too many physical characteristics to list here. Because gender inferences were very accurate in certain decades, it makes some intuitive sense that several bodily descriptors are exclusively deployed along gendered lines. Let's start by taking the terms that the model considered to have the strongest gender associations (calculated, in part, by how often they lead to correct gender inferences). We can then compare them with a sample from the 2000s.

Sample of Ten Words Attributed More Exclusively to Either Fictive Men or Women

1850		2000	
Attributed to Men	Attributed to Women	Attributed to Men	Attributed to Women
Arm	Bare	Chest	Breast
Belly	Was Clasped		
Big	Beautiful		
Breast	Bosom		
Huge	Bright		
Was Bend	Was Kissed		

1850		2000	
Attributed to Men	Attributed to Women	Attributed to Men	Attributed to Women
Red	Blue		
Right	Was Taken		
Hand	Full		
Was Rubbed	Little		

Table 3. Bodily terms that my model considered to be most reliably indicative of either men or women characters.

This isn't a complete list, but that's fine for our purposes. We just want to see which features helped the model infer characters' genders the most, because those features functioned as the most reliable signs of gender for the classifier. In the 1850s model, there are a wide range of features that were consistently associated with either men or women. In the 2000s model, there are far fewer, indicating words like "chest" for men and "breast" for women are some of the only features that the model was confident about. What's fascinating about that result is that "breast" was more exclusively attributed to men in the 1850s, even though that characteristic tends to be associated more exclusively often with women by the 2000s. From a twenty-first century perspective, this result might seem strange, because intuition might lead us to assume that the word "breast" has always been central to configuring a biological female identity.

As one can imagine, this kind of analysis provides a lot of avenues for future research: the gendering of size, the gendering of bodily gestures, etc. But if we want to investigate the gendering of a feature over time, let's say the translocation of the word breast from fictive men to women, we'll need a different kind of method. Drawing from Underwood et. al., we can calculate the difference of a word's usage for men and women across the timeline. Let's say we randomly select 10,000 physical descriptors attributed to women, then we randomly select 10,000

descriptors attributed to men, and then subtract how often each word is used for women by how often it is used for men. We can then ask, “how much more (or less) often was this used for women than men?” Essentially, we’re subtracting the normalized frequency of a term’s use between men and women, scaling that difference to reflect how often we might expect to see that term in 10,000 physical characteristics describing either men or women. The more positive the value, the more often a word was used for women. The more negative, the more it was used for men. In this case, if we take the normalized frequency of “breast,” “chest,” and “bosom” being attributed to women, and subtract the normalized frequency of those words being used from men, will the differences reflect the transfer of the word breast from fictional men to women?




Figure 6. Word use difference between fictional men and women, scaling to see how that difference plays out in 10000 words. Shown for the words *bosom*, *breast*, and *chest*.

While the word “breast” was more frequently attributed to men before the twentieth century, that physical characteristic becomes more exclusively attributed to women as we move towards the twenty-first century. Around the same time, the word “chest” gradually shifts from being attributed to characters of both genders to being more closely attributed to men. “Bosom,” on the other hand, is steadily governed

more often by feminine characters, though not as exclusively as the word breast by the 2000s. If those patterns prove durable, they would add weight to Butler's famous argument that "the gathering of attributes under the category of sex [is] suspect...That penis, vagina, breasts, and so forth, are *named* sexual parts is both a restriction...and a fragmentation of the body".³⁹ For Butler, the difference between biological sex and constructed gender is untenable. Even bodily features that seem like a stable signifier of male or female identity, such as breasts and chests, are not inherently objective signs of natural difference. Rather, as Figure 6 may suggest, the association of these terms with a specific gender is discursively constructed over time.

Because it's not intuitively clear what that broad historical transition might mean at the level of text, it's worth working through a few concrete examples. Let's take novels from each end of the timeline that most frequently mention a man's breast and compare them. Starting at the 1850s-1860s, it's not difficult to find various scenes where descriptions of men's breasts are used to dramatize sensational fiction. On an anecdotal register, several of those novels turn out to be penny dreadful adventure books, such as Edward Viles's immensely popular *Black Bess; or, The Knight of the Road* and John Mills's *Too Fast to Last*. In *Black Bess*, mentions of men's breasts often appear in theatrical scenes involving swashbuckling combat – the kind where an unfortunate character announces his demise by melodramatically placing "his hands upon his breast and [falling] backwards."⁴⁰ In these scenes, the mentions of men's breasts are quite literal. There are other scenes in *Black Bess*, however, where that body part dramatically focalizes the narrative around more internal conflicts. In one scene, for instance, a stable boy agrees to shelter rogue highwaymen but then debates betraying their presence to the authorities while they sleep. The reader is then told that "a severe struggle took place in his breast. By giving up the highwaymen to justice, he would be merely doing his duty to society." Where readers might expect such ethical conflicts to take place in the "head," "mind", or "heart" of a character, these novels locate such conflicts in the breasts of men.

Anecdotally, when we fast forward to the late 90s and early 2000s, the books that most often mention men's chests, as well as women's breasts, are sometimes fantasy novels with a similar adventure element. In fantasy adventures such as Terry

Goodkind's *The Stone of Truth* or Stephen R. Donaldson's *The Wounded Land*, we see some similarly literal mentions of these body parts. Adventuring men get into medieval skirmishes and come out with "a handful of arrows that stuck out of his chest."⁴¹ When it comes to describing these sorts of melees, the use of the word breast versus chest is arguably quite interchangeable. What's very different, however, is that the use of the word chest to describe men sets up scenes where women can then contrastingly be described as having breasts. At the very least, in Goodkind's novel, there's a much more noticeable heteronormative arrangement where women are depicted as intimately "rubbing a hand on his chest" or scenes of a man "complimenting a woman on her breasts."⁴² Unlike Viles's nineteenth century penny dreadful, the more physically intimate representations of men and women tend to produce these sorts of gendered contrasts. That's not to say that this heteronormative arrangement is the only reason why the word chest gets increasingly aligned with men. Just that Goodkind's novel presents some examples of how that alignment of men with chests and women with breasts could have happened.

From these close readings and the evidence in Figure 6, we may be able to conclude that the bodies of men and women are now quite different from what they were in the 1850s. The gender-body associations built up in the nineteenth century have collapsed (or, in some cases, have become re-associated with a different gender). To advance that thesis, we could find other bodily characteristics that gradually move from one gender to another across the timeline.




Figure 7. Word use difference between fictional men and women, scaling to see how that difference plays out in 10000 words. Shown for the words: *hip*, *skin*, and *was bitten*.

These particular gender instabilities might turn out to be more important to literary scholars, as they more immediately lead to questions that could use literary and social forms of evidence. At the very least, they point to a kind of instability not covered in my classification experiment (Figure 4). We saw that gender gradually diminishes as an organizing force from the 1960s forward, but that's not the only form instability can take. Instead, body language that was once used to configure one gender identity can be redeployed to construct different gender identities altogether.

Conclusion

This essay has argued that gender divisions became increasingly central to a growing aspect of character. From the mid-nineteenth century, to the early twenty-first, bodily features continually become a growing aspect of all characters. This

expansion, however, has had a disproportionate effect on the characterization of women. In fact, the proportion of words used to describe characters' bodies has consistently been greater for women than men, suggesting that the body has been more central to developing female characters. At the same time, as the body continued to become an increasingly conventional mode of representing people, bodily characteristics were increasingly becoming distributed along a masculine-feminine axis. This continues to be the case until around the 1960s, when inferring gender from characters' physical features starts to become more and more difficult. Taken altogether, there's a lot of evidence to support the long historical relationship between representations of the body and binary gender. And there's also evidence suggesting that those embodied binaries have played a more prominent role in books by men until the past four decades.

Each of these conclusions enticingly maps onto existing theories about the broader culture, but there are several methodological questions to explore along the way. An important question, for instance, is how best to negotiate the quantitative evidence presented in this essay with that of previous research. If we consider the different results between this essay and Underwood et al., for instance, one could reasonably conclude that certain forms of gender differentiation diminished while others were on the rise. I want to take that statement just one step further, suggesting that we may want to be more sensitive to the *varying weight* of those different forms of characterization. If we're going to argue that the strength of gender stereotypes concretizes in some modes of characterization and diminishes in others, then we're really going to want some way to compare the relative significance of those different modes. Given that bodily description, for instance, is nearly one-fifth of the words used in producing characters, it doesn't seem completely unfair to suggest that certain forms of gender differentiation were more visible than others. This isn't at all to say that some configurations of gender are more significant than others. Rather, there should be a conversation about whether certain forms of differentiation had a disproportionate effect on the gendering of characters.

I also don't want to make it seem like I have provided a complete solution. I've merely sketched out one way of analyzing the varied relationship between character and gender. It's entirely possible, and very likely, that dividing some of the words describing characters and labelling it "anatomical description" proves to be an

incomplete approach. One can imagine, for instance, that if I had included non-anatomical forms of embodiment, such as “fear,” “nervousness,” etc., the results could look very different. Moreover, as I discussed earlier, there’s also the open question of how to relate the results of two models conditioned on different feature sets. My model of body language contains a subset of the features used in Underwood et al.’s experiments, but there’s an important open question about what the discrepancies mean for each other. I suggest that my model disaggregates Underwood et al.’s model, unfolding patterns of literary change that would otherwise be suppressed. However, I ultimately conclude that we’ll want both models, as they shed light on different scales of characterization. These methods are still relatively new to the study of literature, so there’s still room to discuss how best to contextualize the different results of predictive models. At the very least, this essay hopes to instigate conversation on that methodological front. And, if we expect that more models of character and gender will be developed in the coming years, that conversation will have to happen eventually.

As we sort out the different approaches, there are still some theoretical debates about gender that will need to be addressed. To what degree does the evidence lend itself to Mulveyian theories of sexual objectification? We have seen that fictive women, at least since the 1850s, have been comprised of more bodily description than men, meaning that we have some evidence that women have historically been more associated with the body. But we also saw that a binary conception of embodied gender was much more prominent in novels written by men? What’s the gap between those patterns and a more complete mapping of objectifying body language? If cultural analytics is invested in engaging with that longer history of gender and sexual criticism, that question seems like a good place to start.

Taken altogether, there’s a lot more work to be done in order to understand embodied conceptions of gender. My hope is that this essay sheds light on the longer literary history about the body and gender. Future work can hopefully move past a binary model of gender, analyzing how more diverse forms of gender identity have been embodied throughout history. In this case, those binaries were necessary for rendering the literary effect of particular history of body language. More specifically, treating characters as either men or women elucidates the relationship between binary gender and anatomical language, analyzing the imbrication of

physiological words and conventional representations of gender. Moving forward, however, it will be exciting to remove the binary entirely and study more androgynous or gender fluid modes of embodiment.

Appendix: Corpus Metadata

Corpus	Year Range
Stanford Literary Lab	1750 - 1900
Chicago Text Lab	1880 – 2000

Year Range	Percent of Total Dataset
1850 – 1900	29 %
1900 – 2000	71 %

Year Range	Percent of Novels by Women	Percent of Novels by Men
1850 – 1900	43 %	57 %
1900 – 2000	38 %	62 %

Decade	Percent of Novels by Women	Percent of Novels by Men
1850	42%	58%
1860	42%	58%
1870	40%	60%
1880	50%	50%
1890	40%	60%

Decade	Percent of Novels by Women	Percent of Novels by Men
1900	30%	70%
1910	30%	70%
1920	30%	70%
1930	34%	66%
1940	40%	60%
1950	24%	76%
1960	23%	77%
1970	28%	72%
1980	38%	62%
1990	49%	51%
2000	51%	49%

Notes

¹ Matthew Jockers and Gabi Kiriloff, “Understanding Gender and Character Agency in the 19th Century Novel,” *Cultural Analytics* (December, 2016).

² See also Andrew Piper, *Enumerations: Data and Literary Study* (Chicago: University of Chicago Press, 2018), 118-146. Piper’s fifth chapter, “Characterization,” explores the relationship between gender and different parts of speech.

³ Using numbers to study the representation of gender in fiction also extends far beyond these two examples. See also, Eve Kraicer and Andrew Piper, “Social Characters: The Hierarchy of Gender in Contemporary English-Language Fiction,” *Cultural Analytics* (January 2019). Or, Sean G. Weidman and James O’Sullivan, “The Limits of Distinctive Words: Re-Evaluating Literature’s Gender Marker Debate,” *Digital Scholarship in the Humanities* 33, no. 2 (2017): 374-90.

⁴ Ted Underwood, David Bamman, and Sabrina Lee, “The Transformation of Gender in English-Language Fiction,” *Cultural Analytics* (February, 2018). For further discussion of that article, see also: Ted Underwood, *Distant Horizons: Digital Evidence and Literary Change* (Chicago: University of Chicago Press, 2019), 111-142. The chapter “Metamorphoses of Gender” draws from their article in *Cultural Analytics*, discussing some of the open questions that come out of their study.

⁵ Laura Mulvey, “Visual Pleasure and Narrative Cinema,” *Screen* 16, no. 3 (1975): 6–18.

⁶ Judith Butler, *Gender Trouble: Feminism and the Subversion of Identity* (New York: Routledge, 1990).

⁷ Several parts of this project were only doable with the help of others. Funding for this project came through a research assistantship funded by NovelTM, directed by Andrew Piper. The Chicago Novel Corpus was collected by Hoyt Long and Richard Jean So and metadata was provided by Teddy Roland. The Stanford Novel Corpus was provided by Matthew Jockers and required the contributions of the Stanford Literary Lab.

⁸ For two recently analogous projects that have reached adjacent arguments, see also: (Piper, 105-115). Piper’s fourth chapter, “Fictionality,” compares a corpus of fiction and nonfiction texts, providing evidence that words about sensory experience and the observed human body are characteristic of fiction. See also (Underwood, 24-26). Underwood’s first chapter, “Do We Understand the Outlines of Literary History?” uses logistic regression to model the relationship between biography and fiction. Then, he sees which features correlate with which genre. Similar to Piper’s findings, body parts turn out to be an important feature of fiction.

⁹ In “The Transformation of Gender in English-Language Fiction,” Underwood et. al. show how words like “hair” have become more associated with women and “chest” with men.

¹⁰ David Bamman, Ted Underwood, and Noah Smith, “[A Bayesian Mixed Effects Model of Literary Character](#),” *ACL* 2014.

¹¹ This preprocessing was primarily performed using spaCy libraries. <https://spacy.io/>.

¹² For a more thorough explanation of dependency parsing, see Stanford’s page on their Neural Network Dependency Parser. <https://nlp.stanford.edu/software/nndep.html>.

¹³ The list of body language to extract was created by scraping several different ESL lists and scraping synonyms for those parts. The full list of body words is as follows: [body, physique, build, figure, anatomy, shape, flesh, skin, frame, skeleton, appearance, hair, hairs, head, skull, forehead, brow, temple, face, countenance, physiognomy, eye, eyes, pupil, pupils, lash, lashes, eyelash, eyelashes, eyebrow, eyebrows, iris, irises, nose, nostril, nostrils, lip, lips, mouth, tooth, teeth, tongue, ear, ears, eardrum, eardrums, earlobe, earlobes, neck, cheek, cheeks, dimple, dimples, cheekbone, cheekbones, chin, jaw, jawbone, jowl, jowls, dimple, dimples, cheekbone, cheekbones, moustache, beard, mole, freckle, freckles, shoulder, shoulders, armpit, armpits, chest, bust, torso, breast, breasts, bosom, nipple, nipples, arm, arms, limb, limbs, elbow, elbows, forearm, forearms, wrist, wrists, bicep, biceps, triceps, stomach, abdomen, belly, waist, hip, hips, thigh, thighs, crotch, groin, butt, buttocks, leg, legs, knee, knees, calf, calves, ankle, ankles, heel, heels, foot, feet, toe, toes, toenail, toenails, sole, soles, hand, hands, fist, fists, finger, fingers, digit, digits, thumb, thumbs, fingernail, fingernails, forefinger, palm, palms, knuckle, knuckles, wrist, wrists, brain, brains, nerve, nerves, join, joints, tendon, tendons, ligament, ligaments, lung, lungs, diaphragm, pharynx, larynx, trachea, bronchi, kidney, kidneys, bladder, uterus, esophagus, gut, liver, intestine, intestines, bowel, bowels, ovaries, vagina, prostate, penis, testicle, testicles, heart, artery, arteries, vein, veins, capillary, capillaries, bone, bones, marrow, lock, locks, curls, sight, eyesight, vision, perception, view, hearing, taste, appetite, smell, touch.] One major oversight of this approach is that it does not include slang terms for these bodily terms. Intuitively, we can imagine literature that relies heavily on such language (e.g. novels with a significant regional dimension), and this word-list *could* systematically diminish anatomical description in those novels.

¹⁴ Stephen King, *Salem’s Lot* (New York: Anchor Books, 2013).

¹⁵ For a critique of these labels, see Miriam Posner, “What’s Next: The Radical, Unrealized Potential of Digital Humanities,” in *Debates in the Digital Humanities*, ed. Matthew K. Gold and Lauren F. Klein (Minneapolis: University of Minnesota Press, 2016).

¹⁶ Inferring gender from names was done using Lincoln Mullen, *Gender: Predict Gender from Names Using Historical Data*. (version R package version 0.5.2), 2018. <https://github.com/ropensci/gender>.

¹⁷ A complete spreadsheet of the metadata is provided in the data repository.

¹⁸ The appendix provides information on the gender and historical distribution of my dataset. One aspect of the 1900-2000 portion of my corpus (from Chicago Text Lab), that is not apparent in the metadata, is that it is somewhat dependent upon popular genre forms in certain periods (Romance, Detective, Social, Fantasy, Sci – Fi, etc.). The 1850-2000 portion of the corpus comes from multiple sources (Project Gutenberg, Chadwyck-Healey, etc.) and is less informed by notions of genre.

¹⁹ In fact, according to their study, there is a steady decline from about 1800 to 1970. At which point, around only 30% of characterizing words are used to describe women.

²⁰ Ryan Heuser and Long Le-Khac, "A Quantitative Literary History of 2,958 Nineteenth-Century British Novels: The Semantic Cohort Method." Stanford Literary Lab Pamphlet Series, 2012.

²¹ Herbert L. Sussman, *Victorian Masculinities: Manhood and Masculine Poetics in Early Victorian Literature and Art* (Cambridge: Cambridge University Press, 2008).

²² If there were, hypothetically, 1000 words describing women and 200 of them were describing their physical features, we'd say that 20 percent of the words characterizing women had something to do with their physicality.

²³ There are many ways of interpreting Pearson's correlation coefficient. In psychological research, there are some conventions to interpret r 's effect size. Small: 0.10, Moderate: 0.30, Large: 0.50. Jacob Cohen, "A Power Primer," *Psychological Bulletin* 112, no. 1 (1992): 155-59.

²⁴ (Piper, 128-138).

²⁵ Bill Burgwinkle, "Medieval Somatics," In *The Cambridge Companion to the Body in Literature*, ed. David A. Hillman and Ulrika Maude (Cambridge: Cambridge University Press, 2015), 10–23.

²⁶ Fredric Jameson, *The Antinomies of Realism* (London: New York: Verso, 2015), 27-44.

²⁷ Catherine Gallagher, *The Body Economic: Life, Death, and Sensation in Political Economy and the Victorian Novel* (Princeton: Princeton University Press, 2006).

²⁸ Lucy Hartley, *Physiognomy and the Meaning of Expression in Nineteenth-Century Culture* (Cambridge: Cambridge University Press, 2005).

²⁹ (Butler, 12).

³⁰ Binary classification was used in Jockers's and Kirilloff's study of gendered action. They modelled the extent that character-verbs were attributed along gendered lines. Similarly, Underwood et al. use this method to model all words attributed to fictive men and women.

³¹ The classification algorithm was nearest shrunken centroid. For an application and thorough discussion of that classifier, see Robert Tibshirani, Trevor Hasti, Balasubramanian Narasimhan, and Gilbert Chu, "Class Prediction by Nearest Shrunken Centroid, with Application to DNA Microarrays," *Statistical Science* 18, no. 1 (February 2003): 104-17.

³² In Underwood et al., inferring gender peaks at around 76% and steadily dips to sub-65% accuracy. See Figure 7.

³³ For a longer discussion of unintended machine learning biases, including aggregation bias, see: Harini Suresh and John V. Guttag, "A Framework for Understanding Unintended Consequences of Machine Learning," *arXiv* (Jan, 2019). [arXiv:1901.10002v1](https://arxiv.org/abs/1901.10002v1).

³⁴ Kate Flint, *The Victorians and the Visual Imagination* (Cambridge: Cambridge University Press, 2000), 26.

³⁵ Peter J. Capuano, *Changing Hands: Industry, Evolution and the Reconfiguration of the Victorian Body* (Ann Arbor: University of Michigan Press, 2015), 100.

³⁶ More specifically, Capuano demonstrates how those bodily conflicts often surface in William Thackeray's *Vanity Fair* where Becky "'Sharp put out her right fore-finger – And gave him a little nod, so cool and so killing, that Rawdon Crawley...could hardly contain his laughter.'"

³⁷ As with the previous models, we randomly sampled an even number of characters from each decade (225 men and 225 women). I continued to use the nearest shrunken centroid classifier. Future replications will want to test other classifiers, such as logistic regression.

³⁸ Underwood will present evidence to aligns with this possibility. Ted Underwood, "A Measured Perspective," *PMLA* (forthcoming). Further research will then be needed to see whether this improvement is more attributable to the gendered perspectives of authors themselves or the types of gendered fiction they produce.

³⁹ (Butler, 146).

⁴⁰ Terry Goodkind. *Stone of Tears*. (New York: Tor Books, 1995).

⁴¹ Stephen R. Donaldson. *The Wounded Land*. (New York: Del Rey, 2012).

⁴² (Donaldson, 82). (Donaldson, 464).