Connotation Frames of Power and Agency in Modern Films

Maarten Sap   **Marcella Cindy Prasetio**   **Ari Holtzman**
Hannah Rashkin   **Yejin Choi**
Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, USA
{msap,mcp21,ahai,hrashkin,yejin}@cs.washington.edu

Abstract

The framing of an action influences how we perceive its actor. We introduce connotation frames of *power* and *agency*, a pragmatic formalism organized using frame semantic representations, to model how different levels of power and agency are implicitly projected on actors through their actions. We use the new power and agency frames to measure the subtle, but prevalent, gender bias in the portrayal of modern film characters and provide insights that deviate from the well-known Bechdel test. Our contributions include an extended lexicon of connotation frames along with a web interface that provides a comprehensive analysis through the lens of connotation frames.

1 Introduction

A viewer's impression of a movie character is influenced by how they are written and portrayed, which can in turn influence how people form stereotypes on gender norms (Behm-Morawitz and Mastro, 2008). A character's actions can be projected with varying levels of power and agency, depending on the specific verbs used. For instance, somebody who "accepts" things is implied to be a passive decision-maker (or of lower agency) than somebody who "assesses" things. While not explicitly stated, these connotative meanings projected by different verbs can influence the assumptions the audience makes about the people being described. These assumptions can have negative consequences if they reinforce negative stereotypes (Walton and Spencer, 2009).

To formalize this implicit information about people projected by actions, we introduce *power* and *agency* connotation frames, two new types of
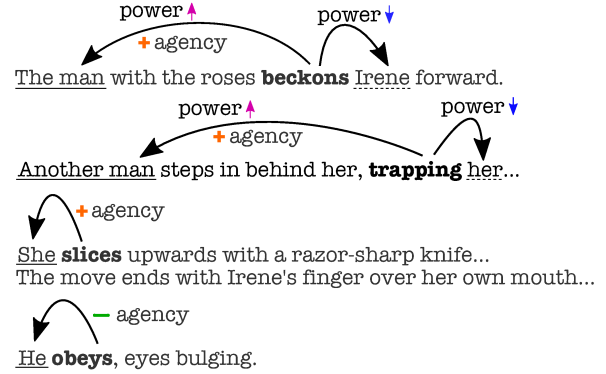

Figure 1: An excerpt from a box-office hit, *Sherlock Holmes* (2009). **Bolded** words are the predicates, solid underlined phrases are the agent of the verb, and dash underlined words are the theme. The full example with additional nuanced discussion is available in Figure 6 in the appendix.

predicate-specific connotative relationships as an extension to Rashkin et al. (2016)'s connotation frame lexicon. For instance, in Figure 1, the verb "beckoning" implies that its theme (Irene) has less power than its agent (the man). In the third line, Irene displays strong agency when she "slices" in self-defense. In contrast, when the man "obeys", the man has low implied agency.

Using the new connotation lexicon, we present a quantitative study to reveal the subtle, but prevalent gender[1] bias in modern films. Going beyond the surface level analysis such as screen time or number of female characters (Google, 2017), our study aims for a more focused and precise analysis of power differentials between fictional men and women.

In summary, our major contributions include the creation and release of a lexicon with two new connotative dimensions: *power* and *agency* and an

---

[1]We acknowledge that gender lies on a spectrum, and reducing it to a male-female binary is simplistic, however our data limits a more complex understanding of gender.

He **implored** the tribunal to show mercy.

power(AG < TH)

The princess **waited** for her prince.

agency(AG) = -

Figure 2: The formal notation of the connotation frames of power and agency. The first example shows the relative power differential implied by the verb **"implored"**, i.e., the agent ("he") is in a position of less power than the theme ("the tribunal"). In contrast, "He **demanded** the tribunal show mercy" implies that the agent has authority over the theme. The second example shows the low level of agency implied by the verb **"waited"**.

interactive demo website of our findings (see Figure 5 in the appendix for a screenshot).[2] Furthermore, as will be seen in Section 4.1, connotation frames offer new insights that complement and deviate from the well-known Bechdel test (Bechdel, 1986). In particular, we find that high-agency women through the lens of connotation frames are rare in modern films. It is, in part, because some movies (e.g., Snow White) accidentally pass the Bechdel test and also because even movies with strong female characters are not entirely free from the deeply ingrained biases in social norms.

2 Connotation Frames of Power and Agency

We create two new connotation relations, *power* and *agency* (examples in Figure 3), as an expansion of the existing connotation frame lexicons.[3] Three AMT crowdworkers annotated the verbs with placeholders to avoid gender bias in the context (e.g., X *rescued* Y; an example task is shown in the appendix in Figure 7). We define the annotated constructs as follows:

Power Differentials Many verbs imply the authority levels of the agent and theme relative to


Figure 3: Sample verbs in the connotation frames with high annotator agreement. Size is indicative of verb frequency in our corpus (bigger = more frequent), color differences are only for legibility.

one another. For example, if the agent "dominates" the theme (denoted as $power(\text{AG}>\text{TH})$), then the agent is implied to have a level of control over the theme. Alternatively, if the agent "honors" the theme (denoted as $power(\text{AG}<\text{TH})$), the writer implies that the theme is more important or authoritative. We used AMT crowdsourcing to label 1700 transitive verbs for power differentials. With three annotators per verb, the inter-annotator agreement is 0.34 (Krippendorff's α).

Agency The agency attributed to the agent of the verb denotes whether the action being described implies that the agent is powerful, decisive, and capable of pushing forward their own storyline. For example, a person who is described as "experiencing" things does not seem as active and decisive as someone who is described as "determining" things. AMT workers labeled 2000 transitive verbs for implying high/moderate/low agency (inter-annotator agreement of 0.27). We denote high agency as $agency(\text{AG})=+$, and low agency as $agency(\text{AG})=-$.

Pairwise agreements on a hard constraint are 56% and 51% for power and agency, respectively. Despite this, agreements reach 96% and 94% when moderate labels are counted as agreeing with either high or low labels, showing that annotators rarely strongly disagree with one another. Some contributing factors in the lower KA scores include the subtlety of choosing between neutral

Figure 4: Label distributions for *power* and *agency* based on the crowdsourced annotations.

and positive/negative as well as the skews in the distributions of labels (i.e. more positive than negative labels, see Figure 4). Note that a similar difference between KA scores and soft percent agreement was found in our previous connotation frame work (Rashkin et al., 2016).

3 Bias in Movie Scripts

We use 772 movie scripts from (Gorinski and Lapata, 2015) as a test bed to validate our new connotation frames. Scripts have distinct structure, which allows us to easily parse narrations, dialogues and character names.

We automatically extract 21K male/female characters, using a name-gender list[4] along with gender specific lexicons (e.g., "actor"/"actresses", "duke"/"duchess") to automatically assign gender based on their first three narrations. To identify verbs with characters as their agent, we dependency parse the narratives using the SpaCy[5] parser. Power and agency label distributions in our corpus are consistent with the annotation distribution (Figure 4), and there is little variance across movies (see Figure 8 in the appendix).

In our dataset, there are nearly twice as many men as there are women (34.6% women), in line with previous findings by Smith et al. (2015) and Radford and Gallé (2015). Women are also less present on screen and speak less in movies (Google, 2017). We control for that disparity in all subsequent analyses by including the number of words for each character (standardized) as a covariate. Findings in all the following sections hold when controlling for movie genre (as retrieved from IMDB.com), as well as when controlling for effects from individual movies.

---

[4] http://www.cs.cmu.edu/Groups/AI/util/areas/nlp/corpora/names/0.html
[5] https://spacy.io/

Frame	β	gender
$agency(\text{AG})=+$	-0.951	**M****
$power(\text{AG}>\text{TH})$	-0.468	**M****
$agency(\text{AG})=-$	0.277	**F****
$power(\text{AG}<\text{TH})$	*not sig.*	

Table 1: Power and agency connotation frames for male and female narratives, controlled for length of narrative text. β represents the change in log-odds of a character being male/female were the corresponding frame to change by one unit. Significant results ($**: p<.001$) are in bold. "Male" was coded as 0, "Female" as 1.

3.1 Bias in Narratives

Narratives describe what characters are *doing*. We investigate how they vary in terms of *power* and *agency*, using our connotation frames. We measure how each standardized frame metric is associated with the gender of the character through a logistic regression, controlling for the total number of words that the character said, and correcting for multiple comparisons using Holm's correction (Holm, 1979).

Listed in Table 1, our results show that male characters are portrayed with higher level of agency compared to women. Men are also portrayed to have more authority than women as they are more often the agent of powerful verbs.

This suggests that screenwriters tend to have female characters contribute more to the aesthetic of the movie through low-agency verbs, rather than the plot, which is reminiscent of existing gender bias tests for movies (Yehl, 2013).

3.2 Bias in Character Expression

To further our validation of the new connotative dimensions, we look at how characters *express* themselves in movies. Using our connotation frames and LIWC (Tausczik and Pennebaker, 2016), we compile metrics for every character's dialogue. As in subsection 3.1, metrics were standardized for better β interpretability. LIWC results that are not discussed below can be found in the appendix (Tables 4 and 5).

From Table 2, it seems male characters display more power and authority through their speech than their female counterparts do. Specifically, women are written to use more hedges (*# Hedges*) whereas men are written to use more imperative sentences (*# Imperative Sent.*), a finding that re-

Frame/metric	β	gender
$agency(\text{AG})=-$	0.968	**F****
$agency(\text{AG})=+$	−2.177	**M****
$power(\text{AG}>\text{TH})$	−0.542	**M****
$power(\text{AG}<\text{TH})$	0.236	**F****
# Imperative Sent.	−0.232	**M****
# Hedges	0.165	**F****
I	0.835	**F****
they	−0.160	**M****
we	−0.361	**M****
you	0.405	**F****
assent	0.202	**F****
space	−1.136	**M****
discrep	0.423	**F****
inhib	−0.171	**M****

Table 2: Gender association with our connotation frames and a subset of LIWC metrics for characters' dialogue, controlled for number of words spoken. All results are significant (** : $p<.001$).

flects real-world dialogues (Prabhakaran et al., 2014). The usage of imperatives tends to convey power and dominance according to the findings of Bramsen et al. (2011). Along with the fact that female characters tend to agree (*assent*) more than male characters, this corroborates the finding in subsection 3.1 that male characters are generally given more power and agency. Furthermore, male characters use inhibitory language more (*inhib*), which contains words pertaining to blocking or allowing, suggesting that these characters are in positions of power.

Further evidence of power imbalances is found through function words. Women tend to use *I* and *you* pronouns more, whereas men use *we* and *they* pronouns more, echoing real life (Schwartz et al., 2013). Kacewicz et al. (2014) found an association between using "I" pronouns and being lower status, and conversely between "we" pronouns and being higher status. This corroborates the theory that women in movies are generally portrayed with a lower status than men.

Men in movies tend to mention more physical actions (*space* category) whereas women tend talk about what "could" be but isn't (*discrep*; e.g.,"should", "would"). This evokes more commanding connotations for male characters and subordinate connotations for female ones, reinforcing gender stereotypes.

These findings, rooted in previous research, confirm that our connotation frames capture exist-

Metric/Frame		β	P/F
F dial.	# Words	10.02	pass**
F dial.	$agency(\text{AG})=+$	−9.65	fail**
F dial.	$power(\text{AG}>\text{TH})$	2.05	pass*
F narr.	$power(\text{AG}>\text{TH})$	−1.19	fail*

Table 3: Significant correlates of passing the Bechdel test. *F*: metric for female characters, computed on the dialogues (*dial.*) or on the narratives (*narr.*). * : $p<.05$; ** : $p<.001$.

ing bias in how male and female characters display different levels of power and agency in their dialogue.

4 Power, Agency and the Bechdel test

A movie passes the Bechdel test (Bechdel, 1986) if it (1) has two (named) female characters, (2) who talk to each other, (3) about something other than a man. While this is a low bar, a surprising number of movies fail at least one of the three criteria. In particular, as many as 42% of the movies in our dataset fail the test according to an online database of the Bechdel scores.[6]

4.1 Beyond the Bechdel Test

We provide comparative insights between the analysis based on connotation frames and the Bechdel test. First, we aggregate our connotation frames, both on dialogue and narration, into movie-level averages per gender. Then, we add features capturing presence of female/male characters (e.g., # F/M words, # F/M characters). Table 3 shows the correlation between passing the Bechdel test and our movie-level connotation frame features using a multivariate logistic regression.

As expected, a movie with more female speaking time is more likely to pass the Bechdel test since it mostly captures female representation. We also find that female characters using agent-empowering verbs, which tend to be more assertive, slightly increases the odds of passing the Bechdel test. Female speakers who use empowering verbs, regardless of the verb's agent, tend to go against the gender-norms of hedging and being less assertive (as we showed in subsection 3.2).

Unexpectedly, movies where women talk with high agency are much less likely to pass the

Bechdel test. Perhaps these movies typically only show scenes of women interacting in a male-dominated setting. Similarly, the use of more agent-empowering verbs in female narratives decrease the odds of passing the Bechdel test. Chances of two powerful women talking to each other might be lower because movies are less likely to have a lot of powerful women.[7]

Power and Agency of Princesses We further provide a qualitative analysis using Wikipedia plot summaries for movies that are not in our script dataset. Bechdel-passing movies with female protagonists, such as *Frozen* (2013) or *Cinderella* (1950), still perpetuate negative female stereotypes. In *Frozen*, Elsa is portrayed as the only high agency, high power woman, as seen below.[8] Anna and Cinderella, despite also being protagonists, display significantly less power and agency.

Power and Agency levels of Disney princesses

The Bechdel test is limited, either by being too inclusive of movies who portray women in non-authoritative, passive positions or by excluding movies that have strong women with agency, who just happen not to talk to each other about something besides men. Our extensions to the connotation frame lexicon provide finer grained information about how women are portrayed through their expression and their actions, which can act complementary to measures of their presence.

5 Related work

There is much prior research focus on bias in social media (Garcia et al., 2014; Prabhakaran et al., 2014; Ratkiewicz et al., 2011; Yano et al., 2010; Srivastava and Sahami, 2009), complementing our investigation on movies. Fast et al. (2016) examine the stereotypes present in fan-fiction us-

ing a lexicon-based strategy that focus on commonly gender-biased attributes (e.g., *emotional* for women) rather than the overall power dynamics of the story. In a similar vein, Ramakrishna et al. (2015) learn word-level "gender ladenness" features by looking at the neighbors of 925 manually annotated words.

There exist various sets of high-level criteria to assess gender bias of character portrayal in fiction (Yehl, 2013; Romano, 2013; Powers, 2016). Agarwal et al. (2015), in particular, automate the Bechdel test using social network features, finding that women are less central to the plot in movies that fail it. We compare our linguistic analysis of power and agency with the Bechdel test, demonstrating the need for more fine-grained analysis of how gender is depicted in movies.

Close in spirit to our investigation, Schofield and Mehr (2016) train a number of classifiers over movie scripts for determining the gender of individual (and pairs) of speakers as well as the expected length of their relationships. In contrast, we focus on understanding how the gender of a given character implicitly relates to features that track their control over their own path (agency) and the world around them (power).

6 Conclusion

We created and released new connotation frames of *power* and *agency*, allowing for more nuanced writing analysis than previously possible. We validate our new frames through a case study on movie scripts. Specifically, we analyze differences in power and agency for male and female characters, and compare these dimensions to the Bechdel test. Our connotation frames confirm evidence of imbalances in gender portrayal in movies.

---

[7]Similar observations may have inspired the Mako Mori test (Romano, 2013), which looks at whether there's a female character with a story arc that doesn't support a man's.

[8]Note that plot summaries are more biased toward active verbs, which explains the low negative agency for all characters.

References

Apoorv Agarwal, Jiehan Zheng, Shruti Kamath, Sriramkumar Balasubramanian, and Shirin Ann Dey. 2015. Key female characters in film have more to talk about besides men: Automating the bechdel test. In *HLT-NAACL*.

Alison Bechdel. 1986. *Dykes to watch out for*. Firebrand Books.

Elizabeth Behm-Morawitz and Dana E Mastro. 2008. Mean girls? the influence of gender portrayals in teen movies on emerging adults' gender-based attitudes and beliefs. *Journalism & Mass Communication Quarterly*, 85(1):131–146.

Philip Bramsen, Martha Escobar-Molano, Ami Patel, and Rafael Alonso. 2011. Extracting social power relationships from natural language. In *ACL*.

Ethan Fast, Tina Vachovsky, and Michael S Bernstein. 2016. Shirtless and dangerous: Quantifying linguistic signals of gender bias in an online fiction writing community. In *ICWSM*.

David Garcia, Ingmar Weber, and Venkata Rama Kiran Garimella. 2014. Gender asymmetries in reality and fiction: The bechdel test of social media. In *ICWSM*.

Google. 2017. Using technology to address gender bias in film. https://www.google.com/about/main/gender-equality-films/index.html.

Philip John Gorinski and Mirella Lapata. 2015. Movie script summarization as graph-based scene extraction. In *NAACL*.

Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70.

Ewa Kacewicz, James W Pennebaker, Matthew Davis, Moongee Jeon, and Arthur C Graesser. 2014. Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33(2):125–143.

Kelsey Powers. 2016. The furiosa test developed from female presence in mad max: Fury road. http://www.calvin.edu/chimes/2015/11/09/the-furiosa-test-developed-from-female-presence-in-mad-max-fury-road.

Vinodkumar Prabhakaran, Emily E Reid, and Owen Rambow. 2014. Gender and power: How gender and gender environment affect manifestations of power. In *EMNLP*.

Will Radford and Matthias Gallé. 2015. Roles for the boys?: Mining cast lists for gender and role distributions over time. In *WWW*.

Anil Ramakrishna, Nikolaos Malandrakis, Elizabeth Staruk, and Shrikanth S Narayanan. 2015. A quantitative analysis of gender differences in movies using psycholinguistic normatives. In *EMNLP*.

Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A Data-Driven investigation. In *ACL*.

Jacob Ratkiewicz, Michael Conover, Mark R Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011. Detecting and tracking political abuse in social media. *ICWSM*.

Aja Romano. 2013. The mako mori test: 'pacific rim' inspires a bechdel test alternative. https://www.dailydot.com/parsec/fandom/mako-mori-test-bechdel-pacific-rim/.

Alexandra Schofield and Leo Mehr. 2016. Gender-distinguishing features in film dialogue. In *CLfL@NAACL-HLT*.

H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin E P Seligman, and Lyle H Ungar. 2013. Personality, gender, and age in the language of social media: the open-vocabulary approach. *PLoS One*, 8(9):e73791.

Stacy L Smith, Marc Choueiti, Katherine Pieper, Traci Gillig, Carmen Lee, and Dylan DeLuca. 2015. Inequality in 700 popular films: Examining portrayals of gender, race & lgbt status from 2007 to 2014. *Media, Diversity & Social Change Initiative, USC Annenberg School for Communication and Journalism*.

Ashok N Srivastava and Mehran Sahami. 2009. *Text mining: Classification, clustering, and applications*. CRC Press.

Yla R Tausczik and James W Pennebaker. 2016. The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.*

Gregory M Walton and Steven J Spencer. 2009. Latent ability: Grades and test scores systematically underestimate the intellectual ability of negatively stereotyped students. *Psychological Science*, 20(9):1132–1139.

Tae Yano, Philip Resnik, and Noah A Smith. 2010. Shedding (a thousand points of) light on biased language. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*.

Joshua Yehl. 2013. Kelly sue deconnick talks captain marvel, pretty deadly, and the sexy lamp test. http://www.ign.com/articles/2013/06/20/kelly-sue-deconnick-talks-captain-marvel-pretty-deadly-and-the-sexy-lamp-test.

Figure 5: Online connotation frame visualization tool, available at `http://homes.cs.washington.edu/~msap/movie-bias/`. Interactively visualize a movie from our corpus for power and agency breakdowns by gender.

The man(+,↑) with the roses **beckons** Irene(↓) forward. Another man(+,↑) steps in behind her, **trapping** her(↓). He(+,↑) closes in to **grab** her(↓)————and WHAP!WHAP!WHAP! Gets three cosh-strikes across the chin, dropping him immediately. Still holding the roses, the man(=,↕) **finds** himself(↕) smacked up against the bricks with Irene's hand over his mouth. Quick as a snake, she(+) **slices** up-wards with a razor-sharp knife, cutting his belt, then his clothes, all the way to his collar. The move ends with Irene's finger(+,↕) over her own mouth, **signalling** the mugger(↕) to be quiet. He(−) **obeys**, eyes bulging. Irene(+,↑) **frisks** him(↓) expertly. She's(+,↑) **mugging** the mugger(↓). There is excitement in her eyes; this turns her on.

Figure 6: Unabridged text from Figure 1. Taken from *Sherlock Holmes* (2009), which grossed $524 million. **Bolded** words are verbs being examined, solid underlined phrases are the agent of the verb, and dashed words denote the theme. Parentheticals are connotation frame annotations, with ↑, ↓, and ↕ denoting a power gradient (or lack thereof) and +,−,= denoting high agency, low agency, and neutral agency, respectively. Note that *Sherlock Holmes* did not pass the Bechdel test.

Task Description: For each verb, determine whether the subject or the object seems to have more authority (higher status) relative to each other.
Examples:
- X has more authority:
  - "X vetoes Y" --> X is clearly presumed to outrank Y.
- X and Y have similar authority:
  - "X loves Y" --> X and Y are mutually involved and appear to be similar status
- Y has more authority:
  - "X idolizes Y" --> Y is presumed to have some power over X.

More Examples:

X has more authority	X and Y have similar authority	Y has more authority
X overrules Y, X vetoes Y	X fights Y, X marries Y	X idolizes Y, X salutes Y

For the following verbs, which has higher authority:
1) X **rescues** Y: ◯ X has more authority ◯ similar ◯ Y has more authority
2) X **serves** Y: ◯ X has more authority ◯ similar ◯ Y has more authority

Figure 7: Example annotation task for authority. Each verb was annotated by three crowdworkers. Each verb was rated using placeholders (e.g., X,Y), to prevent context biasing rater's perception of the verb.

Figure 8: Label distributions for power and agency frames. Each datapoint represents a movie's aggregated frame counts (aggregated over all characters's dialogue and narrations). The distributions are consistent with the annotation distributions in Figure 4, and variance across movies is relatively low. Note that our analyses in Section 3 are performed at the character level, but these distributions show that our results are likely not swayed by specific outlier movies.

metric	β	gender
Presence		
# Words	-0.126	**M****
# Scenes Present	-0.054	
# Talk Turns	0.308	**F****
# Sentences	0.203	
Spoken features		
assent	0.202	**F****
filler	0.231	**F****
nonfl	0.070	
Function words		
adverb	0.904	**F****
article	-1.557	**M****
auxverb	1.844	**F****
conj	0.571	**F****
funct	3.124	**F****
future	-0.154	**M***
i	0.835	**F****
ipron	0.180	
negate	0.746	**F****
number	-0.397	**M****
past	0.418	**F****
ppron	1.97	**F****
preps	-1.188	**M****
present	1.174	**F****
pronoun	2.131	**F****
quant	-0.225	**M***
shehe	0.327	**F****
they	-0.16	**M****
verb	2.215	**F****
we	-0.361	**M****
you	0.405	**F****

Table 4: Gender association of various standardized metrics in dialogue. All metrics are LIWC categories, except those starting with "#". β represents the change in log-odds of a character being male/female were the corresponding frame to change by one unit. Significant results (**** : $p<.001$) are in bold. "Male" was coded as 0, "Female" as 1.

metric	β	gender
Affective processes		
affect	0.601	**F****
anger	-0.463	**M****
anx	0.2	**F****
negemo	-0.178	**M***
posemo	0.758	**F****
sad	0.253	**F****
swear	-0.392	**M****
Biological processes		
bio	0.016	
body	-0.172	**M****
health	0.093	
ingest	0.157	**F****
sexual	0.026	
Cognitive mechanisms		
# Hedges	0.165	**F***
# Imperative Sent.	-0.232	**M****
cause	0.077	
certain	0.178	**F***
cogmech	0.968	**F****
discrep	0.423	**F****
excl	0.731	**F****
incl	-0.46	**M****
inhib	-0.171	**M****
insight	0.41	**F****
tentat	0.219	**F***
Perceptive processes		
feel	0.039	
hear	0.15	**F***
percept	0.238	**F***
see	0.022	
Personal concerns		
achieve	-0.365	**M****
death	-0.125	**M***
home	0.264	**F****
leisure	-0.011	
money	-0.269	**M****
relig	-0.033	
work	-0.344	**M****
Relativity		
motion	-0.075	
relativ	-0.679	**M****
space	-1.136	**M****
time	0.559	**F****
Social processes		
family	0.231	**F****
friend	0.088	**F***
humans	-0.317	**M****
social	0.592	**F****

Table 5: continuation of Table 4

Frames		Pearson r
agency(AG)=−	*agency*(AG)==	0.067 **
agency(AG)=−	*agency*(AG)=+	−0.209 **
agency(AG)=−	*power*(AG>TH)	−0.023 *
agency(AG)=−	*power*(AG<TH)	0.399 **
agency(AG)=−	*power*(AG=TH)	0.086 **
agency(AG)==	agency(AG)=+	*not sig.*
agency(AG)==	*power*(AG>TH)	0.023 *
agency(AG)==	*power*(AG<TH)	0.043 **
agency(AG)==	*power*(AG=TH)	0.069 **
agency(AG)=+	*power*(AG>TH)	0.406 **
agency(AG)=+	*power*(AG<TH)	−0.095 **
agency(AG)=+	power(AG=TH)	*not sig.*
power(AG>TH)	*power*(AG<TH)	0.024 *
power(AG>TH)	*power*(AG=TH)	−0.043 **
power(AG=TH)	*power*(AG<TH)	0.049 **

Table 6: Partial correlation between connotation frames in the narrations for characters in our dataset (controlled for number of words). Significant results (* : $p<.05$; ** : $p<.001$, Holm corrected) are in bold. Most frames have low correlations with others, signifying that the dimensions captured are different. We find mild correlations between negative agency and theme empowering verbs, as well as for positive agency and agent empowering verbs.

	# Imper. Sent.	# Hedges
agency(AG)=+	0.268 **	−0.146 **
agency(AG)=−	−0.056 **	0.245 **
agency(AG)==	0.134 **	*not sig.*
power(AG>TH)	*not sig.*	−0.062 **
power(AG<TH)	0.02 *	0.019 *
power(AG=TH)	−0.084 **	0.087 **

Table 7: Partial correlation on our movie corpus between frames and select dialogue features (controlling for number of words). Holm p-value correction is applied (* : $p<.05$; ** : $p<.001$). Most power frames have low correlations with the two dialogue features, though the correlations are all in the expected directions. We find that agency frames are moderately correlated with imperatives and hedges, as expected.