**The Hugo and Nebula Awards for Best Short Story: examining differences in tastes and themes from 1999 to 2018**

~Shana Hadi

The most celebrated American authors take home royalty checks and commemorative plaques aplenty, acknowledgment for their work and its cultural significance. There are awards for most any category of genre, length, theme, and societal relevance, all claiming to honor the "best." But when awards have varying selection criteria and voter composition – often distinguished on the basis of what's "popular" and "prestigious" in a given year – what do these awards really measure? What do these awards say about their stories?

After reading Literary Lab's Popularity/Prestige pamphlet and Techne post on the creation of a "new canon," I started wondering about awards that purposely blend concepts of popularity and prestige. Most, if not all, prestigious awards are judged exclusively by a panel of distinguished jurors (Pulitzer, National Book Award, Booker Prize). Meanwhile, other less highly-vaulted forms of recognition depend on the popular vote (Goodreads Choice) or commercial success (The New York Times Best Seller list).

In speculative fiction, the two most well-known awards – the Hugo and Nebula, established in 1953 and 1966 respectively – offer another perspective to this research question. These awards have an inherent tension in how they annually honor the "best" works of speculative fiction from the previous year, yet have roots to the days of pulp magazines hailed as the successors of penny-dreadfuls and dime-novels. Speculative fiction is a rising genre with increasing relevance to the modern tech-fueled imagination, but in the past it rarely received the same prestige and attention as realist fiction, hence the creation of these awards. Academics notwithstanding, I find some of the most vocal and devout readers – and voters – within this fanbase, who are often responsible for the latest cult following of an often overlooked, but still beloved, work of art.

Presented within a yearly convention, the Hugo and Nebula Awards help unite the larger gathering of fans and authors through a collectively chosen "new canon." Beyond honoring the stories with the greatest print and literary reach, they ultimately reflect and nurture the tastes of their voters through canonizing the "best of the year."

As any reader willing to purchase a Con membership can vote in the Hugos, I wondered if they reward popular works that consciously adopt speculative fiction tropes, while the Nebulas, which require voters to be published authors, would have a literary bent and reward works with more experimental qualities. What are some other factors distinguishing a Hugo from a Nebula story? I am also curious if there would be themes that recur throughout the years (such as the common tropes of "space" or "robots"), or others that rise and fall in success. As the Hugo- and Nebula-nominated works are considered the best of their year, they denote what the overall fanbase values at the time. Would there be an increasing diversity of topics and characters in the more recent selections, especially as more female authors join the lists? And are these themes consistent or volatile?

Footnote: For context on voting procedures: All members of the World Science Fiction Society (WSFS) are eligible to nominate and vote for the Hugo Awards, with the only criteria a purchase of a membership for that year's Worldcon. With two stages of voting, a given voter's five nominations hold equal weight to those of others, with a rank-choice system to select the winner from the final ballot. In contrast, only

Active Members of the Science Fiction and Fantasy Writers of America (SFWA) can nominate and vote for the Nebula Awards; the criteria involves selling at least three short pieces or one long work which meet the minimum word count and prices listed on their website. The six most nominated works form the final Nebula ballot, and the work with the most votes win.

## Background and Corpus: Savoring, Scraping, and Saving Short Stories

Due to my own interests in reading and writing speculative fiction, I decided to focus on the Hugo and Nebula Awards. My first serious attempts at creative writing involved emulating short stories from the nomination lists, and this qualitative introduction to the field inspired me to study them from a computational perspective.

I collected 194 Hugo- and Nebula-nominated short stories spanning 20 years from 1999-2018, with the exclusion of one story from 2001 (Michael Swanwick's Hugo-nominated "Moon Dogs") unable to be located within a reasonable timeframe.[1] This corpus has over 925,000 English-language words from 123 unique authors, with 99 Hugo stories and 126 Nebula stories. Out of these, 30 stories have been cross-nominated for the Hugo and Nebula, with 3 stories having won both.

I scraped the majority of the stories from online magazines and digital Stanford Searchworks resources and scanned the remaining thirty from print (the Nebula Showcase anthologies were particularly helpful), as many stories from the 2000s proved quite scarce.[2]

## Methodology

### Most distinctive words

I first ran several Python scripts to clean the data by removing capitalization, punctuation marks, and special characters, and saved these new files into a different folder. I got word frequencies for my sub-corpora (per short story and category), and then on those I ran Fisher's exact test, which is a statistical test that determines if there are nonrandom associations between two categorical variables. In my case, if certain words are more likely to appear within the Hugo or Nebula stories (the Award subcorpus), or within the stories from a given year between 1999 to 2018 (the Year subcorpus).

The script returned individual tokens that had p-values less than 0.05, which denotes there is evidence against the null hypothesis (that there is no correlation between a particular word and the award or year). It also calculated the ratio of observed to expected values for each word in comparison to the larger corpus, with higher values signaling a greater probability of these words being distinct to the subcorpus as they appear more times than would be expected out of all possible words in the corpus.

### Topic modelling

Using the cleaned files from my previous tests, I used Mallet 2.0.8, a package for statistical natural language processing that uses LDA (latent Dirichlet allocation), to build a topic model for my corpora. A topic model is a statistical model frequently used in text-mining, as it produces "topics," or clusters of words that frequently appear together within a text. These can be used as a proxy for themes, as the method's underlying idea is that while a document may contain several topics at given proportions, a topic model can approximately capture these relative balances.[3]

After running over twenty iterations, I settled on 95 topics with an optimization interval of 10 to ensure that I would have a fair balance of topics (approximately 2 stories per topic) as not all results make coherent sense due to the randomness of the method. I ran additional models after I filtered out the most common stopwords and a supplementary list of names compiled from my preliminary results (such as "john" and gehenna"), though I sought to keep the worldbuilding vocabulary since presumably they would cluster around the themes (such as "immerser" and "bridesicle").[4]

**Findings Highlights:**

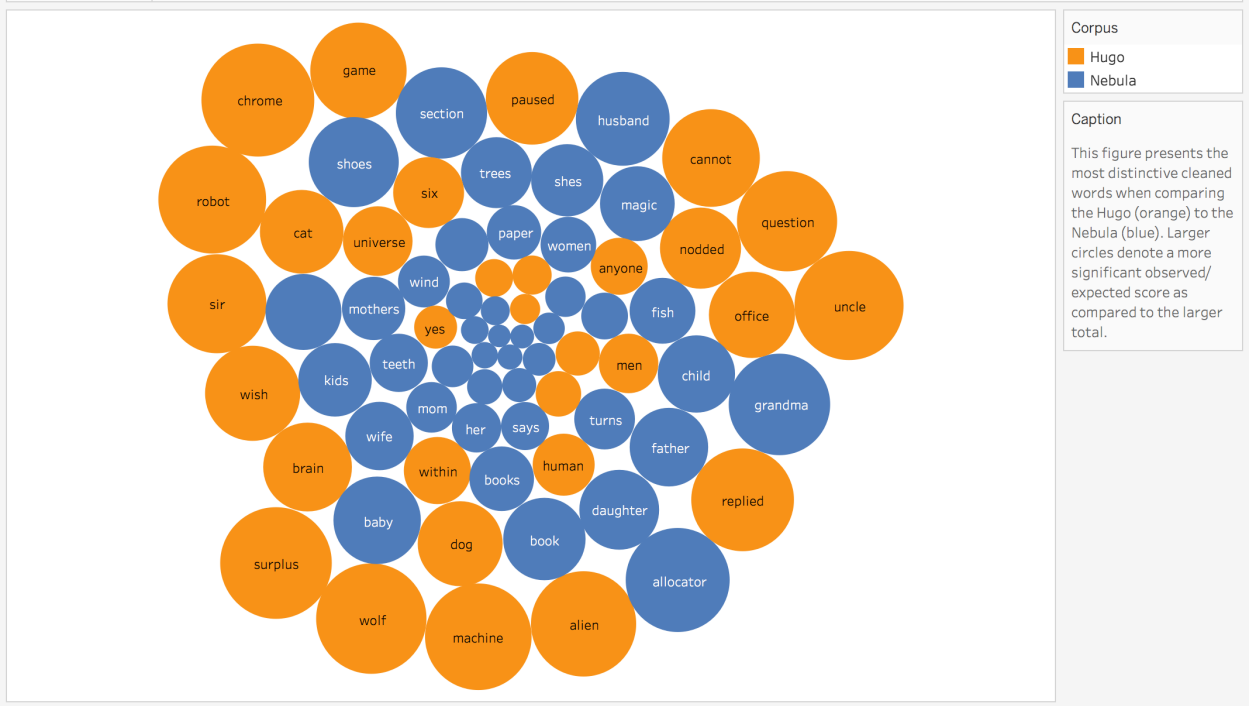



**Most distinctive words by award: the Nebula features more females?**

My script returned 389 words, and to further cut down, I filtered out words that had a p-value at or greater than 0.04 as well as names and unique worldbuilding vocabulary, such as “jeanette” and “gr7.” To reduce the domination of any given short story, I sorted for at least 75 observations and an observation/expected score of at least 1.15 (somewhat arbitrary qualifications). This left 73 words, and my most significant observation was the disparity of female pronouns.

I found that the Nebula stories had more occurrences of “she,” “shes” (she’s), “her,” “herself,” “woman,” and “women,” with a high degree of certainty. The word “shed” in this case would most likely refer to “she’d” (with the apostrophe removed), but the remaining words suggest that the Nebula short stories are more likely to feature female main characters, especially with the collective results of these words and their high observed/expected scores.

| token_ | Corpus | Observations | p_value | Obs/Exp |
|---|---|---|---|---|
| shes | Nebula | 361 | 0.0003148243035 | 1.317518248 |
| women | Nebula | 143 | 0.03483532739 | 1.265486726 |
| herself | Nebula | 379 | 0.001362573897 | 1.263333333 |
| her | Nebula | 8518 | 2.96E-41 | 1.240425222 |
| woman | Nebula | 431 | 0.006367385345 | 1.197222222 |
| she | Nebula | 7855 | 1.30E-25 | 1.187991531 |
| shed | Nebula | 353 | 0.03337882897 | 1.157377049 |

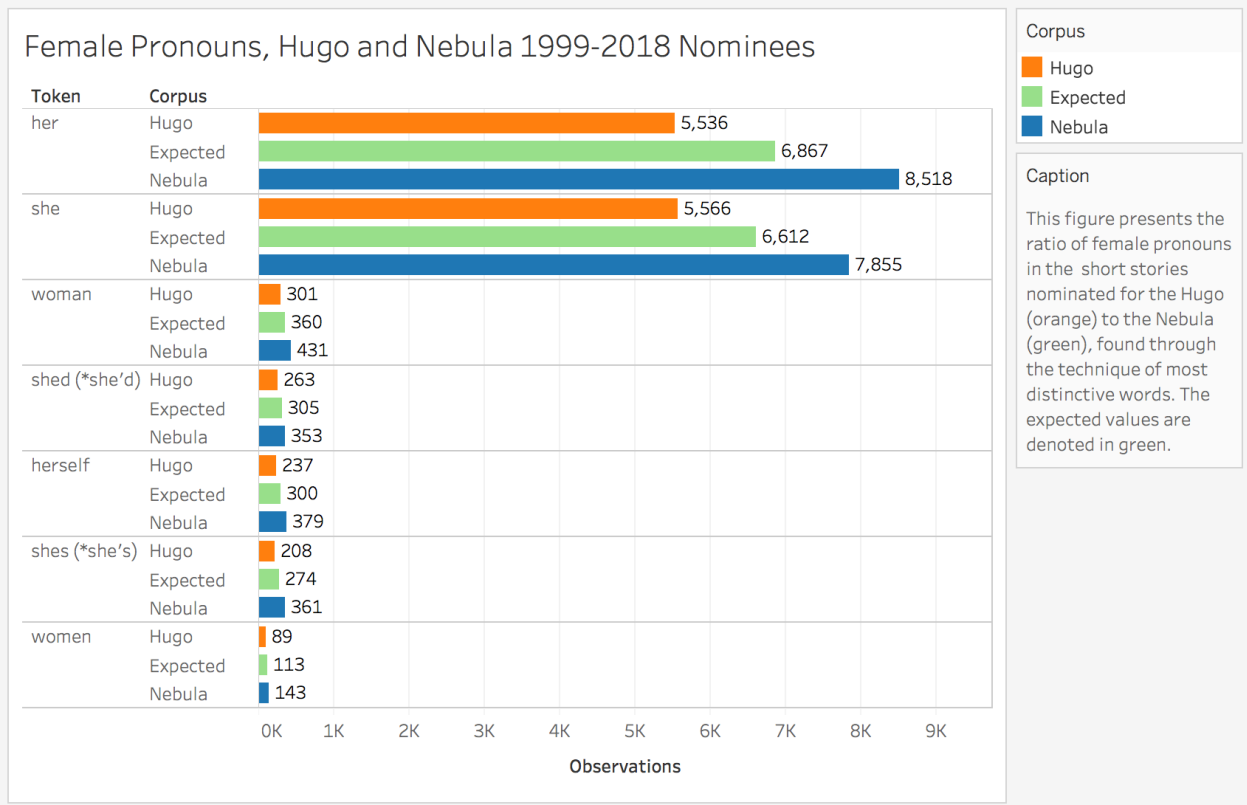


When I factor in occurrences of "i" and "you" (associated with first- and second-person), which are slightly more common in the Hugo stories but tragically failed to meet my Obs/Exp baseline despite their significant p-values (likewise, "me," myself," and "your" didn't even make an appearance on the list), the Nebula-features-females angle becomes more compelling.

| token_ | Corpus | Observations | p_value | Obs/Exp |
|---|---|---|---|---|
| i | Hugo | 9710 | 2.28E-09 | 1.088931255 |
| you | Hugo | 5497 | 0.0005305687983 | 1.065310078 |

Short stories with only one character, even when told with the intimacy of first-person, are "very hard to make interesting" (to quote a former TA from English 146S: Secret Lives of the Short Story). Hence, the many observations of "she" and "her" at least suggest that a larger variety of female characters – whether in the foreground or background – populate these short stories. Nonbinary and nonhuman characters most likely do as well, though their descriptors are not as easy to find through most distinctive words.

And since words such as "he," "his," "him," and "himself" didn't make the most distinctive words list, this means that their proportions are roughly equal between the Hugo and Nebula stories. So while they may have a similar number of male characters, more female characters join the cast of the Nebula short stories.

Several short stories in this corpus don't have an identifiably-gendered protagonist, which may challenge this observation. Ted Chiang's "Exhalation" (2009 Hugo-winner) exclusively features the firsthand scientific observations of a nonhuman lifeform. There are other outliers, such as Ursula K. Le Guin's "The Bones of the Earth" (2002 Hugo-nominee) which uses she/her pronouns, but is set in the same world of *The Left Hand of Darkness* which used he/his descriptors and thematically explored the fluidity of gender. That's the caveat of studying speculative fiction, which often embraces breaking existing boundaries and conventions.

These results support my hypothesis that there are differences between the Hugo and Nebula, and these short stories reflect the growing diversity of authors and fans who vote on them. Could we say the same 50 years ago?

Now for the idle *speculation* phase: why the Nebulas over the Hugos? Perhaps it's the composition of the fanbase. It's possible that WorldCon members may skew towards men raised on the "Big Three" in science fiction (with Robert A. Heinlein's documented "female troubles"). Most pre-2010 Hugo-nominees come from three old magazines established in the mid-twentieth century, such as *Asimov's Science Fiction*, *Fantasy and Science Fiction*, and *Analog Science Fiction*, which have varying generations of readers. Likewise, any fan willing to purchase a membership can vote, and the fanbase has historically leaned towards white men interested in science, similar to the composition of older STEM professionals. (The mere existence of the 2015 and 2016 Rabid Puppies campaigns are not helping the optics.) In contrast, the composition of active members of the SFWA may be more well-balanced, with author-voters who will pay greater attention to stories that experiment with the craft of writing. At the very least, the recent Nebula-nominees have a significant proportion of female authors, and five female authors have won the Nebula in a row from 2013-2018, and 15/20 Nebula-winners were written by women. Then again, the 2019 Hugos were swept by women, so my mileage may vary.

**Most distinctive words by year: an abundance of male/female pronouns**


My script returned 2885 words. With fewer stories per year (9-12), I had to be more careful on what constituted a significant result, so I additionally filtered out names, words that had a p-value at or greater than 0.03, an observed/expected score of at least 1.2, and over 50 observations (so unique words from one story could make the list, but would not dominate it). This left 323 words, and with quite a number of stopwords and worldbuilding vocabulary, I am hesitant to extrapolate further (how much can you really say from the occurrences of one word?), but these could be an interesting starting point for further research.

Linked to my earlier discovery from finding the most distinctive words from the Hugo and Nebula subcorpora, I observed the years where male and female pronouns were distinctive. I am tentatively using this as a proxy to discuss which years might have had more male- or female-identifying characters, out of the 12 years that made the list.

| token_ | Corpus | Observations | p_value | Obs/Exp |
|---|---|---|---|---|
| his | 2000 | 371 | 2.67E-04 | 1.315602837 |
| her | 2001 | 619 | 5.20E-07 | 1.348583878 |
| she | 2001 | 639 | 8.83E-10 | 1.445701357 |
| he | 2002 | 813 | 2.61E-18 | 1.619521912 |
| him | 2002 | 220 | 2.43E-03 | 1.341463415 |
| his | 2002 | 571 | 2.18E-13 | 1.622159091 |

| | | | | |
|---|---|---|---|---|
| she | 2003 | 562 | 5.18E-07 | 1.370731707 |
| he | 2004 | 862 | 9.00E-08 | 1.308042489 |
| his | 2004 | 554 | 2.06E-03 | 1.199134199 |
| he | 2005 | 735 | 2.51E-09 | 1.392045455 |
| him | 2005 | 231 | 1.87E-03 | 1.343023256 |
| his | 2005 | 527 | 6.39E-08 | 1.428184282 |
| he | 2006 | 705 | 4.18E-03 | 1.157635468 |
| he | 2007 | 685 | 1.43E-04 | 1.229802513 |
| her | 2007 | 798 | 6.33E-08 | 1.327787022 |
| she | 2007 | 730 | 1.48E-05 | 1.260794473 |
| her | 2008 | 938 | 2.15E-14 | 1.465625 |
| she | 2008 | 764 | 3.80E-05 | 1.238249595 |
| her | 2010 | 575 | 2.46E-07 | 1.378896882 |
| she | 2010 | 674 | 2.19E-17 | 1.680798005 |
| her | 2011 | 714 | 4.13E-10 | 1.425149701 |
| she | 2011 | 611 | 5.49E-05 | 1.265010352 |
| her | 2013 | 479 | 2.93E-03 | 1.206549118 |

If we count the 4 years with only one distinguishing pronoun (2000, 2003, 2006, 2013) there are 5 male-dominated years and 6 female-dominated years, and when those are excluded, there are 3 male-dominated years and 4 female-dominated years. 2007 is an exception, as while 3 stories use first-person (two male, one female perspective), the remaining 8 use third-person, so perhaps the prevalence of “he” and “she” signal the third-person-perspective rather than the larger array of characters. The female pronouns “she” and “her” both occur more often than the male pronoun “he” and have higher p-values and observed/expected scores, but the differences are too slight to make a ruling.

When in the context of the Nebulas having significantly more female pronouns than the Hugos, I hypothesized that the distinctiveness of female pronouns would be a more recent phenomenon, tied with my perception of speculative fiction becoming especially more inclusive in recent years. There seems to be more male-dominated years pre-2007, and more female-dominated years post-2008, but the results only lukewarmly support my claim. The more recent five-year period of 2014-2018 don’t have any distinctive pronouns at all, and though several stories used second-person perspective, the absence of distinctive pronouns strikes me as quite odd.

**Topic modelling: toppled theories**

Unfortunately for my topic modelling ambitions, the results here are ambiguous. While I have several beautiful graphs, from a statistical level, there are few significant claims I can make that cannot be intuited from just reading the stories. Like my theory on the varying results from most distinctive words by year, I wonder if the variety of worldbuilding vocabularies prevent Mallet from accurately cohering the themes together, having to sort through so many seemingly one-off words, even after names and stopwords were removed.

Using the same short story "Immersion" for an example, topic modelling may be able to connect "immerser" with "gadget" or "technology" when interpreted very broadly, but it is unlikely to connect it to "culture," "identity," or "loss," which are significant elements from the short story symbolized through the word. And frankly, a theme on "technology" does not say very much, considering the genre.

**Space and the universe:** At least supporting my faith in the model, space-related words showed up in 4 of my topics, though the bulk of topic 6 recognizably comes from Ted Chiang's "Exhalation" (2009 Hugo-winner) and Ken Liu's "Mono no Aware" (2013 Hugo-winner), topic 23 from Ken Wharton's "Aloha" (2005 Nebula-nominee), topic 47 from Robert Reed's "Eight Episodes" (2007 Hugo-nominee), with varying lesser percentages from the other stories, and topic 25 from a variety of sources.

| Topic_Number | Weight | Topic_Words |
|---|---|---|
| 6 | 0.08471 | air thoughts universe pressure leaves brain engine brains lungs source end gold chamber filling continue sheets hope patterns clocks flow |
| 25 | 0.25763 | space time light stars universe earth energy planet future ship long solar orbit ancient star black galaxy species planets humans |
| 33 | 0.00791 | aloha universe krenn protocol galaxy future temporal library existence locus converted love seela alohas beach device zone time-reversed hannahs visit |
| 47 | 0.10692 | tiny series audience final invasion research world universe scene worlds aliens episode student scientific students production seek shows complex event |

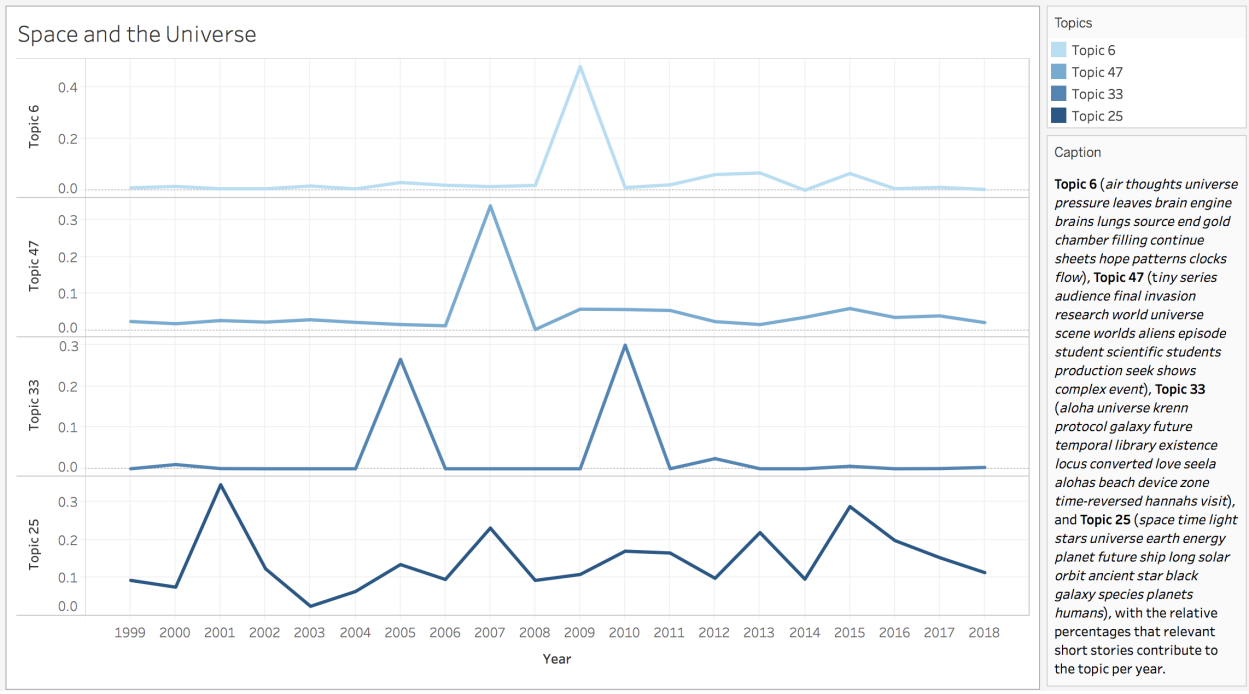


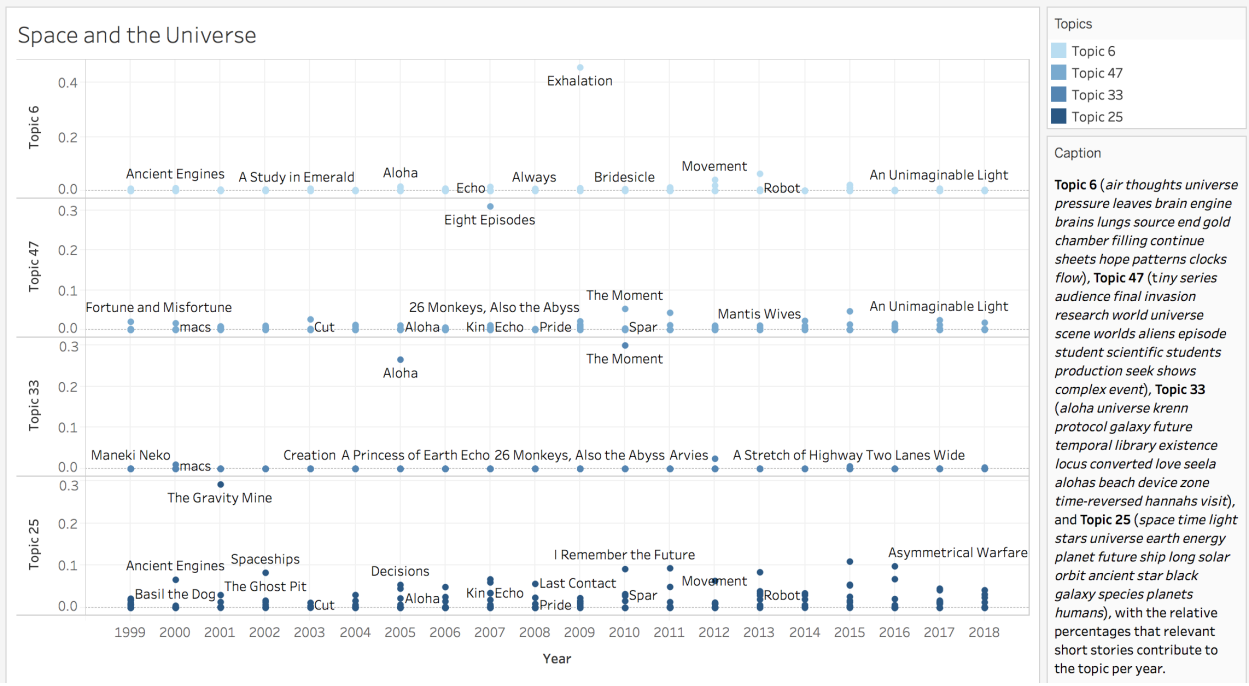


**Writing and books:** While story writing is a recurring occupation for many characters (writers truly write what they know), I find it interesting to see how this fares across time and which stories share this theme.

| Topic_Number | Weight | Topic_Words |
|---|---|---|

| | | |
|---|---|---|
| 17 | 0.10399 | story writer fiction stories email writing computer read science write idea wrote future writers subject ideas reading written fact fan |
| 44 | 0.09477 | book books read write written pages reading words library war paper tea ink fathers cats letters page closure author travels |

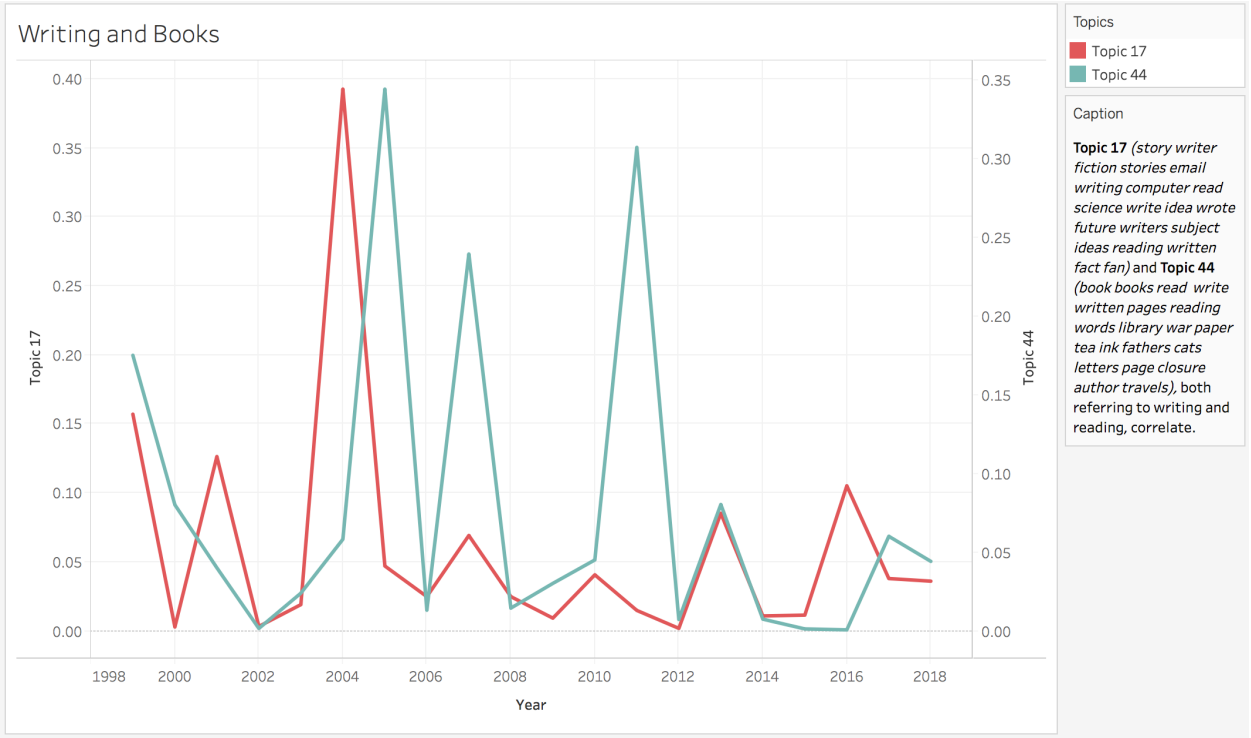


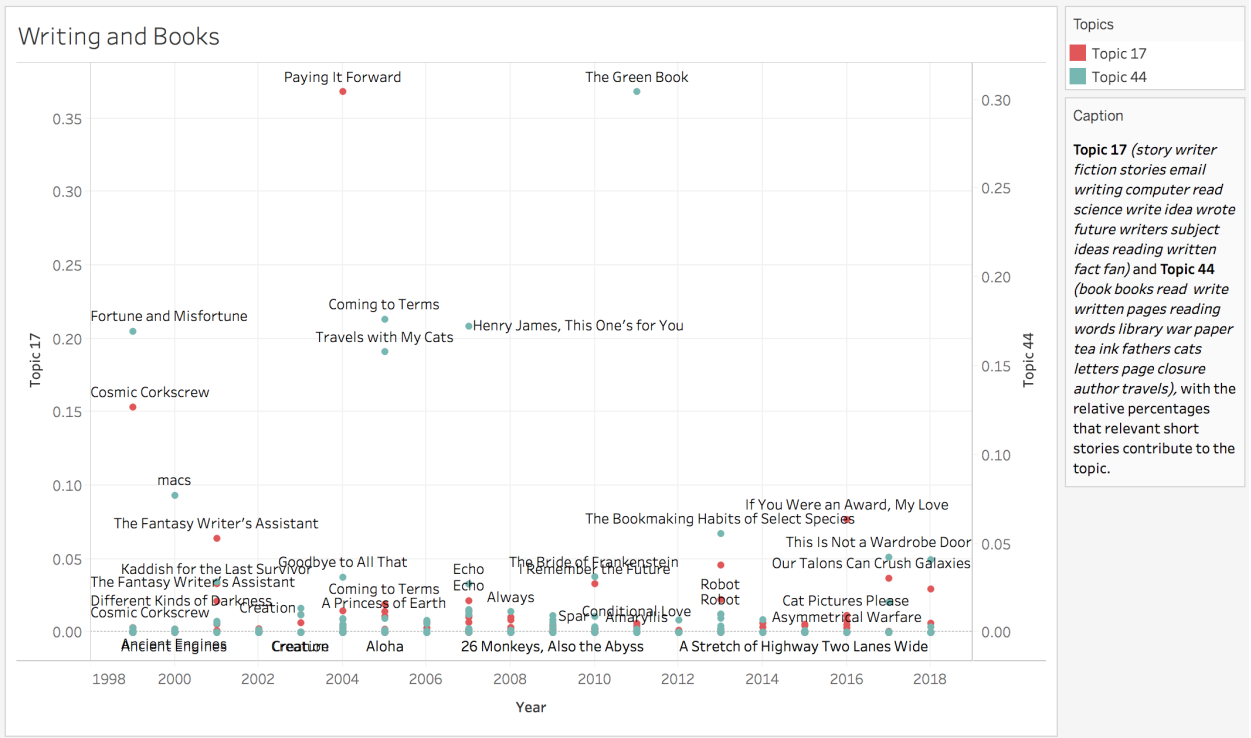

Do these topics reveal much that I can't intuit from reading a random selection of short stories across years? Perhaps not. But the visualizations are interesting, and across a longer period of time (starting from the dawn of the awards), more themes may be revealed.

**Next Steps: Shooting Past the Moon, Stars, and Genre**

The observed distinctions between the Hugo and Nebula Awards for Best Short Story may reflect different compositions of their audiences and tastes. While this corpus has more insights to uncover, another frontier would be exploring how these stories – the speculative fiction "canon" – would compare to those canonized from the realist tradition, such as nominated works for the O. Henry Prize or Pen/Hemingway Award. How blurry are the lines between genre expectations – speculative and realist fiction – and critical and commercial success?

As most prestigious literary awards are decided through juries, perhaps there would be interesting comparisons on the basis of lingering effects from the body of "tastemakers." When comparing award-nominated speculative fiction works to a larger whole, and realist-nominated works to a larger whole, do the different voting processes affect which stories win? Since the Hugo- and Nebula-awards depend on nominations to make the voter ballot, are these stories already commercially popular before they are nominated, as compared to realist works, which may be more critically lauded before they catch the eye of a jury and sell more as a result?

Which work is the "best" continues to be up for question, especially as tastes and themes change to reflect and challenge society. We may not know what makes the canon until fifty years later, but until then, at least we'll have lots of excellent reading.

**Acknowledgements**

Many thank yous to Literary Lab's Dr. Laura McGrath and Dr. J.D. Porter for their generous mentorship, guidance, and fascinating research stories. I also want to thank Green Library curator Dr. Rebecca Wingfield for helping me find several scarce short stories, Mark Algee-Hewitt and Nichole Nomura for their thought-provoking "Novel World-Building: Science Fiction" presentation, and CESTA and Literary Lab for the opportunity to work on exciting projects throughout the summer.

**Notes:**

[1] As Michael Swanwick has 10 other short stories within the corpus, I thought it was an acceptable omission – he is one of the two most heavily nominated authors within this period, the other being Mike Resnick.

[2] I used Abbyy FineReader, a software capable of optical character recognition, to convert my PDF scans into text. I saved each story as an individual .txt file and collected its metadata information, with each story linked via filename in my table.

[3] Most distinctive words Python scripts seeded from Dr. J. D. Porter (thank you!), Mallet software from UMass-Amherst's Andrew Kachites McCallum, and topic modelling tutorial from Programming Historian's Shawn Graham, Scott Weingart, and Ian Milligan.

[4] I did include the 2015 and 2016 Hugo Rabid Puppies slate, as technically they did receive the nominations. This may have skewed more recent results, especially Chuck Tingle's 2016 Hugo-nominee.

**Stanford-protected links:**
Master spreadsheet: https://docs.google.com/spreadsheets/d/1NeTqinZsWCoFM65v-Tlr8OU5FXfhg9fczCIjbFEbKz4/edit#gid=2138875826
MDW by award and year: https://docs.google.com/spreadsheets/d/1C0vVhQOVPb6fpJ3wFj1uDrlKQ7hQvALrWRZ7uDnGuzI/edit#gid=1783141889
Topic model: https://docs.google.com/spreadsheets/d/1llKlrNtLv_NQ-pgr4YfzyKyZumdqQUE6dv--1MOxyAA/edit#gid=1947959982