

**‘Hey, Listen!’: An Analysis of the Treatment
of ‘Annoying Video Game Helpers’ by
Gender**

Nisha Devasia

Abstract

The idea of woman as man's helper has existed since Biblical times, and this belief has persisted into the modern day. In this paper, I examined this relationship in the context of video games, where women are treated as subordinate and as tools for the male gaze both in the real world and in the games themselves. I scraped statements about male and female 'helper'-type characters from the public wiki TV Tropes and ran them through a sentiment analysis algorithm to determine the magnitude of negativity shown towards each gender. The sentiment analyzer found that on average, male and female characters were treated with the same level of negativity, but females were treated worse in the extremities of negativity. However, due to some seeming inconsistencies in the sentiment analysis model's ratings, I conclude that the sentiment analysis model may be the incorrect tool for an analysis like this.

Introduction

The trope¹ of women as man's helper - submissive, passive, and "[created] for man" - has existed since biblical times (Corinthians 11:9). In Genesis 2:18 from the New King James Version of the Bible: "And the LORD God said, *It is* not good that the man should be alone; I will make him an help meet for him." Soon following are Genesis 2:22 and 2:23: "...Then the rib which the LORD God had taken from man He made into a woman, and He brought her to the man. And Adam said: "This *is* now bone of my bones; And flesh of my flesh; She shall be called Woman, Because she was taken out of Man."

This belief still persists today. Service positions such as secretaries, nursing assistants, and administrative assistants are traditionally occupied by women², and are even dubbed 'pink-collar' work. At historical maximum, women have only ever occupied 6.4% of the CEO positions in the Fortune 500 companies, despite the fact that women are now earning more college degrees than men. While it is unlikely that we can directly prove the correlation between Western society's Biblical foundations and the gender imbalance in the workforce today, it is a suspicious set of occurrences indeed.

To add insult to injury: though the aforementioned positions are societally thought of as 'women's work', according to a survey done by the Bureau of Labor Statistics, men in 'pink-collar' professions still tend to earn more than women. Even if women are performing the work that is 'meant' for them, they still are not treated nearly as well as the men in those professions are³. Women are, in every sense, expected to be helpers. In an article about how women are expected to do "office housework", Sheryl Sandberg - the CEO of one of the biggest companies in the world, Facebook - describes a study "led by the New York University psychologist Madeline Heilman, [in which] participants evaluated the performance of a male or

¹ In this context, the word describes "a common or overused theme or device".

² And if they are taken by men, the men are paid higher wages! (Lowen)

³ Perhaps the best example of this is the culinary industry - although cooking is considered a housewife's duty, there are very few female chefs.

female employee who did or did not stay late to help colleagues prepare for an important meeting. For staying late and helping, a man was rated 14 percent more favorably than a woman. When both declined, a woman was rated 12 percent lower than a man. ...After giving identical help, a man was significantly more likely to be recommended for promotions, important projects, raises and bonuses. A woman had to help just to get the same rating as a man who didn't help".

An interesting space in which to examine these relationships is the rapidly expanding gaming space, in which subordinate roles are clearly defined, both in-game and in real life. The gaming market has traditionally advertised heavily towards males, and even the demographic of the industry itself is heavily male dominated. Games with heroines instead than heroes are few and far between, and it is often argued that even those heroines can serve as "dangerous role-model[s]" (Jones, in Cassell and Jenkins, 2000, 339); despite said female characters being strong and independent women, they are still "solely a tool for the male gaze", (MacCallum-Stewart, 2014)⁴. Even in positions of dominance, these female characters are still portrayed as subordinate to males in some way or another.

The relationship that I want to examine in this paper is that between female helper characters and male helper characters. Since being a helper is considered a female's role, the research question is as follows: how are female helper characters in video games treated by the gaming community relative to how male helper characters are treated? We will investigate this relationship by using sentiment analysis methods on a page about 'Annoying Video Game Helpers' from TV Tropes, a corpus of popular opinions about various forms of media.

Literature Review

Several bodies of academic literature have informed this investigation. First, we will investigate literature regarding the relevant trope: that of the 'helper'. The idea of a helper character stems from Vladimir Propp's observations on the structure of folktales and fairy tales. Propp, a Soviet folklorist and scholar, was interested in breaking down Russian folk tales into its simplest irreducible elements, and identified the *Dramatis Personae*, "a limited set of eight broad character types" that appeared in most folktales. Propp described the Helper as somebody "who appears at critical moments to provide support", or a character whose sole *raison d'être* was to assist the Hero on his quest. Crucially, "the contrast between the limitations of this person and the hero may provide a further elucidation of the hero's defining characteristics such as intelligence, determination, courage, etc" (*Changing Minds*).

Propp's Helper characters were almost exclusively male; they appeared in the form of a wise old man, a magician, or other masculine-typed characters. In Propp's *personae*, the only role that females occupied was that of the Princess, or the damsel in distress. And indeed, this trope was alive and well in the earliest video games; the adventures of the titular character in the

⁴ MacCallum-Stewart actually argues against this particular argument, saying that many women have been empowered by Lara Croft, even though many game scholars claim that she is a dangerous face for female empowerment in video games.

Super Mario franchise primarily centered around rescuing Princess Peach from the evil Villain-type character Bowser, and nearly all the early games in the Legend of Zelda franchise involve rescuing the titular princess in some fashion.

Fortunately, females no longer exclusively occupy the role of the damsel in distress in popular video games. They have taken up a new space of roles in gaming, both in-game and in real life - those that are “supportive [and] subordinate” (Taylor, Jensen, & Castell, 2009). They are objectified, both in the ‘IRL’ case of the ‘booth babe’⁵, as well as the cases of heavily sexualized female characters such as Tifa Lockhart, Lara Croft, and Ivy Valentine, to name a few. This has been shown to have detrimental effects on opinions towards women; an “exploratory study found that a video game depicting sexual objectification of women and violence against women resulted in statistically significant...rape-supportive attitudes for male study participants but not for female participants” (Beck et. al, 2012). Another study showed “the effects of exposure to sex-typed video game characters versus images of professional men and women on judgments and attitudes supporting aggression against women...A significant interaction indicated that men exposed to stereotypical content made judgments that were more tolerant of a real-life instance of sexual harassment compared to controls. Long-term exposure to video game violence was correlated with greater tolerance of sexual harassment and greater rape myth acceptance. This data contributes to our understanding of mass media’s role in socialization that supports violence against women” (Dill et. al, 2008).

To add insult to injury, women have not only been marginalized and objectified in games, but for a long time, the game industry essentially ignored their female audience altogether. “Digital games have traditionally been seen as boys’ toys...[and] male teenagers have been perceived as their main target audience, and have largely continued to be so until recently (Laurel, 2008)...Several attempts have been made at creating games specifically for girls, yet most of these efforts have not succeeded (Graner Ray, 2004)” (Quandt, 2014). “A detailed content analysis of Nintendo Power issues...[showed] that mainstream companies largely ignored the girls’ games movement, instead targeting male audiences through...sexualized female characters, magazine covers featuring men, and predominantly male authors. Given the mutually constitutive nature of representation and reality, the lack of women in consumer press then affected girls’ ability to identify as gamers and enter the gaming community. This shows that, even as gaming audiences diversify, inclusive representations are also needed to redefine gamer as more than just ‘male’ “(Cote, 2015). Indeed, the issue of representation is a crucial one. Time and time again, we have seen fields which are ‘meant’ for males suffer from a serious diversity problem. But while there are huge ‘Women in STEM’ movements, we see no such ‘Women in Games’ movements on the same scale.

There have, however, been efforts to include women in the space of ‘casual’ gaming. Nintendo in particular made the first breakthrough into this space with the Wii, and the advent of

⁵ Promotional models for video games that appear at booths at large video game conventions.

mobile gaming made it easier to “integrate gameplay into ordinary or mundane segments of everyday life”. “The ‘encroachment’ of women and girls into what was previously a male-gendered space has not happened without incident, and will probably only become worse before it (hopefully) improves” (Consalvo, 2012). Incidents like GamerGate have shown women exactly what men think of their involvement in ‘real’ gaming, which does not encompass casual gaming endeavors such as Wii Sports games.

Studies have also been done on the difference between how men and women are described in literature, media, workplace evaluations, and even by themselves personally. Hoyle et. al trained a machine learning model on the very large corpus of English books (Goldberg and Orwant 2013) and found that “positive adjectives used to describe women are more often related to their bodies than adjectives used to describe men”. “Word embeddings trained on Google News articles exhibit female/male gender stereotypes to a disturbing extent” (Bolukbasi et. al, 2016). Linguist and tech entrepreneur Kieran Snyder conducted a study on performance reviews given in the workplace and found that “women’s reviews are more likely to include critical feedback”, “men are given constructive suggestions [while] women are given constructive suggestions – and told to pipe down”, and “the manager’s gender isn’t a factor” (Snyder).

There are qualitative differences in the way men and women think about the people and the world around them as well. In 1957, Sherriffs and Mckee found that “the stereotype of men as defined by adjectives which are applied to them by *both* sexes appears to involve three general notions...a straightforward uninhibited social style...the notion of rational competence and ability...[and] action, vigor, and effectiveness”. And indeed, while “the picture of men seems forbiddingly cold and devoid of human frailty, the picture of women...more than compensates for it”. The adjectives that both sexes used to describe women “emphasize social skills and grace...imply warmth and emotional support...[and] concern for the significance or spiritual implications of experience”. Furthermore, “men’s vices...seem limited to mild exaggerations of their desirable characteristics”, while for women, undesirable characteristics included “unreasonable emotionality...[and] is pretty well summed up by the term ‘neurotic’”. This study was a strong indicator that men and women both think about men in more positive terms and can find more negative attributes to ascribe to women. Over thousands of years of females being treated as secondary to males, society’s opinion of women - as well as womens’ opinion of women - is understandably lower than society’s opinion of men. For the context of this paper, the study of the use of adjectives is relevant because we are using a sentiment analyzer to look at the specific language used on the TV Tropes page. Since men’s vices “seem limited to mild exaggerations” while women’s vices are numerous and seem to have a high magnitude of negativity, the sentiment analysis algorithm may pick up on this as well, and be more likely to classify the statements about women as negative.

Methodology

I went through TV Tropes' 'Annoying Video Game Helper' page to gather data. The page is a comprehensive user-generated list of characters in video games that "[have] you screaming 'Stop helping me!'. The characters vary from support-type characters who are usually sidekicks or secondary characters in their game, in-game AI helpers, or merely annoying UI features meant to help the player but do the opposite. Some things that are important to note: the TV Tropes statements are all completely based on opinion, a fact which they make clear at the very top of the page. The opinions about the characters on the 'Annoying Video Game Helper' page may not be shared by everybody in the gaming community, but given that TV Tropes is considered to be a corpus of popular opinion, we can assume that a majority of the gaming community feels the same way about these characters.

The page is also only editable with a TV Tropes account, and the account itself must be approved by moderators before any edits to the page can be made. Approval is usually granted within a few hours, and editing permissions come with it. While editing an article, users are granted an 'edit lock' which prevents other users from editing the article simultaneously; this is likely a feature in-built to prevent spamming. By visiting the edit history of the page, one can view which users made which edits to the page. However, inferring anything about the user from their username is a difficult - and likely unproductive - venture. Furthermore, it does not particularly matter; as shown by the Snyder and Sherriffs and Mckee studies, both men and women tend to have less regard for women.

I separated all the statements made about different characters and sorted them by their gender (or lack thereof). I evaluated the gender by the pronouns used in the statement - if 'she' or 'her' was used, I classified them as female, and if 'he', 'his', or 'him' was used, I classified them as male. If no pronouns were used, I first looked up the character on Google and searched the wiki⁶ for their corresponding game to see if any pronouns were used in their wiki description. If there were pronouns used them, I classified them as female or male accordingly. If there were no pronouns used, or ungendered language such as 'it' or 'they' was used to describe the character, I classified the character as ungendered. Some statements also referred to large groups of non-player characters of both genders, and I classified these as ungendered as well. Some statements very clearly referred to certain UI features about the game that were annoying, or to AI personas that were non-gendered, and all of these statements were classified as ungendered.

Since I wanted to investigate whether the female characters were treated more poorly than the male characters, I decided to use a sentiment analyzer. Sentiment analysis models use machine learning algorithms to report the magnitude of the calculated sentiment of the statement that it receives. They are primarily used by companies to "[help] a business to understand the social sentiment of their brand, product or service while monitoring online conversations" (Gupta). The models are able to predict whether a block of text has 'positive' or 'negative'

⁶ A wiki is a website that allows visitors to make changes, contributions, or corrections, and they are typically used for various forms of popular media.

sentiment with relative accuracy, as well as the magnitude of positivity or negativity in each statement, which ranges from -1 to 1.

Sentiment analysis models are trained on different corpuses. The first sentiment analyzer I considered using was trained on a dataset comprising of 50,000 IMDb movie reviews. However, I decided against this because the IMDb reviews, upon first glance, seemed to be all written by a certain demographic that did not seem relevant to the data that I wanted to run through the sentiment analyzer. Furthermore, the model only returned ‘positive’ or ‘negative’ to the user, and I wanted the magnitudes of positivity or negativity rather than the binary answers.

I then tried to find a model that would classify the adjectives used in the TV Tropes page. I first tried to use the model proposed in Hoyle et. al to classify gendered adjectives by their ‘supersenses’, or lexicographer categories. Hoyle et. al used Tsvetkov et. al (2014) to determine the supersense of each adjective, and although I attempted to adapt Tsvetkov’s model for my own use, I did not have adequate training data to widen the range of adjectives from the ones originally used in Tsvetkov’s dataset.

I eventually decided on IBM’s Natural Language Understanding sentiment analyzer. This decision was made for a few reasons. First, the sentiment analyzer returned the magnitudes of the sentiment rather than just 1 or -1, so I could acquire more granular results from using it. Second, the sentiment analysis model itself is a black box. I personally do not know what data IBM trained it on, which disallows me from forming biases about what the results might be beforehand.

I then ran the ‘female’ and ‘male’ statements through IBM’s Natural Language Understanding sentiment analyzer, which calculated the sentiment of the statement. Sentiments are decimal values between -1 and 1, where a negative number is a calculated negative sentiment, a positive number is a calculated positive sentiment, and a 0 is a neutral sentiment. The model also gave ‘targeted emotion’ values for the categories: joy, anger, disgust, sadness, and fear, which I calculated for each statement as well.

	A	B	C	D	E	F	G	H	I	J
1	Memetic Muta	-0.9								
2	Navi's Suspici	-0.89								
3	Fi in The Lege	-0.65								
4	Roll tended to	-0.62								
5	Taken to extre	0.41								
6	In Spider-Man	-0.88								
7	Angie from Tr	-0.78								
8	Zone of the Er	-0.86								
9	Your AI ally frc	0								
10	The Journeyr	-0.66								
11	A less extremi	0.56								
12	Intentionally a	-0.75								
13	The anonymo	-0.9								

Figure 1. Each statement from the TV Tropes page, along with its calculated sentiment.

	A	B	C	D	E	F	G	H	I	J
1	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy
2	0.43	0.08	0.06	0.06	0.44	0.01	0.07	0.04	0.2	0.41
3	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger
4	0.15	0.62	0.16	0.17	0.4	0.53	0.63	0.56	0.17	0.53
5	Disgust	Disgust	Disgust	Disgust	Disgust	Disgust	Disgust	Disgust	Disgust	Disgust
6	0.1	0.1	0.22	0.02	0.03	0.2	0.1	0.03	0.1	0.02
7	Sadness	Sadness	Sadness	Sadness	Sadness	Sadness	Sadness	Sadness	Sadness	Sadness
8	0.37	0.32	0.71	0.58	0.22	0.21	0.57	0.47	0.37	0.12
9	Fear	Fear	Fear	Fear	Fear	Fear	Fear	Fear	Fear	Fear
10	0.06	0.05	0.06	0.4	0.02	0.46	0.6	0.12	0.2	0.02
11	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy	Joy
12	0.09	0.55	0.58	0.02	0.01	0.12	0.02	0.6	0.12	0.03
13	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger	Anger
14	0.15	0.62	0.16	0.17	0.4	0.53	0.63	0.56	0.17	0.53

Figure 2. Targeted emotion values calculated by the model (female values in the top ten rows, male values in rows 11-21).

After calculating all the values for each statement, I took the average of all the calculated values - i.e., the average of the calculated sentiment values and each of the averages of the targeted emotion values by gender value.

Results

	Female	Male
Total number of statements	52	94
Average sentiment value	-0.534	-0.418
Average joy value	0.171	0.213
Average anger value	0.423	0.388
Average disgust value	0.101	0.116
Average sadness value	0.404	0.350
Average fear value	0.175	0.153

Discussion

First of all, we must address the wide discrepancy in the magnitude of female helper characters versus the magnitude of male helper characters. Contradictory to our discussion in the introduction, we observe that there are far more male helpers than there are female helpers. Another data point that was not included in the above summary was that count of the ungendered helpers, which was almost equal to that of the male helpers. This may imply that perhaps there just are not many female characters in games in general, which has been shown to be true to some extent (Jeroen and Martis, 2007).

If we look at the average sentiment value of the female helpers versus the average sentiment value of the male helpers, we can observe that the value for the female helpers is lower than that of the males, which indicates that on average, the sentiment analysis model classified the statements about the female helpers to be more negative than the statements about the male helpers. However, before we start making assumptions about this particular piece of data, let us take a look at the accuracy of the model itself.

Naturally, the model was not 100% accurate. A likely reason for this is the sarcastic⁷ nature of some of the statements. Sentiment analysis models essentially look for ‘buzzwords’ that they consider negative to classify a statement as negative. Sarcastic statements usually do not use any deliberately insulting language, and therefore the sentiment analyzer doesn’t pick up on the negativity of the statement and classifies it as positive. We can consider this to be a false positive. Similarly, some statements were not negative; rather, they listed a few annoying characteristics of the character, but finished by saying that the character was good overall. Some statements were about how a particular character ‘subverted’ the trope, which means that the trope in question - that of the ‘Annoying Video Game Helper’ does not apply to them. We can calculate the accuracy of the model by adding the number of true negative statements and true positive statements, and dividing this by the total number of statements the model classified.

	True positives	True negatives	False positives	False negatives	Overall model accuracy
Female	2	41	7	2	82.7%
Male	1	74	18	1	76.6%

The accuracies per gender are comparable, and the slightly lower accuracy for the male statements can probably be attributed to the slightly higher incidence of false positive statements for males - 19.1% versus 13.5%. The fairly large number of false positive statements inflated the average sentiment value for males to a number larger than it should have been. If we assume the same percentage of false positives for males as we did for females, and assume that a false positive on average is about 0.5 and a true negative is on average about -0.5, we can recalculate the average sentiment value of the males to be approximately -0.47, a number not significantly different from the average sentiment value of the females.

The targeted emotion values were also not significantly different. While male helpers seemed to ‘spark’ slightly more joy, and female helpers caused slightly more anger and sadness, these numbers are not different in any statistically significant manner.

⁷ Sentiment analysis models have been shown to be bad at sarcasm. Humans reading textual data also have a hard time recognizing sarcasm because a large part of indicating sarcasm is done with tone of voice and body language.

It seems that the sentiment analysis has not yielded anything significant regarding this paper’s original question. However, we can dive deeper to find an interesting result. I decided to define a ‘very negative’ sentiment classification as one that was less than or equal to -0.85. This number is, for all purposes⁸, arbitrary. I then calculated the number of statements about female and male helpers that were ‘very negative’.

	Female	Male
Number of ‘very negative’ statements	16	9
Total number of statements	52	94

Already, these numbers are interesting. Despite having about half as many statements about them in total, the female helpers have almost twice as many ‘very negative’ statements compared to the male helpers.

We can take this analysis further by conducting a chi-square test upon these proportions. A chi-square test is “a statistical test applied to sets of categorical data to evaluate how likely it is that any observed difference between the sets arose by chance” (*Wikipedia*).

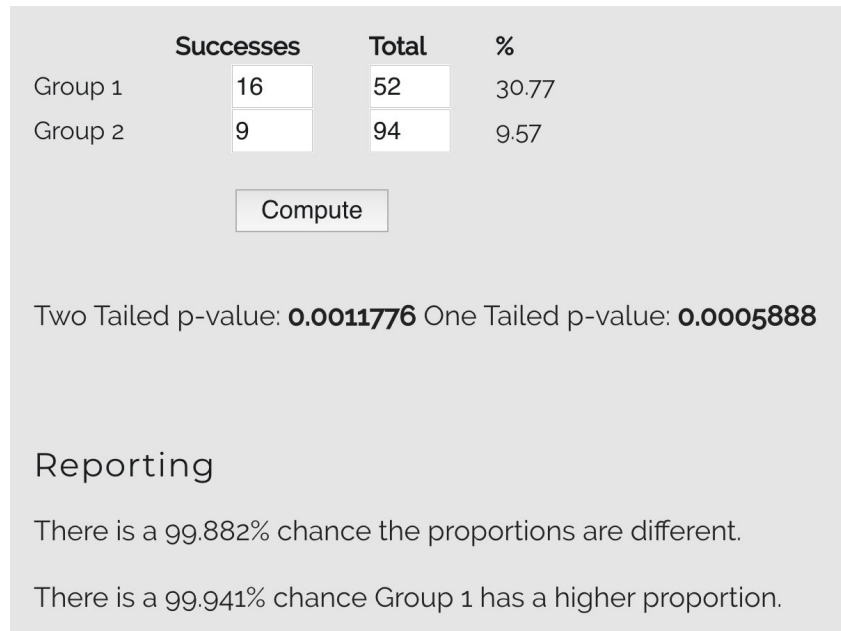


Figure 3. The results of a chi-square test on the proportions of ‘very negative’ statements classified per gender.

This is a startling result. Essentially, the chi-square test predicts that the female helpers have almost a 100% chance of having a higher proportion of ‘very negative’ statements. We can

⁸ An ‘A’ grade at my university is typically defined as 85 or above, and this is admittedly the source of this number.

perhaps infer from this that while the female helpers may not be treated significantly worse on average, they receive more verbal abuse in the extremities of the spectrum.

However, perhaps we should not be too quick to jump to this conclusion. Let's take a look at some of the statements that the classifier flagged as 'very negative'.

In Dragon's Dogma, Pawns will repeatedly remind you of things you already know from experience, like wolves roam in packs, goblins hate fire, a fall from high up means death, and that treading in waist high water will kill you, repeating these "suggestions" to you over and over again, sometimes, in a row.

Figure 4. The specific statement that the classifier found to be the most negative for males, with a score of -0.98.

**Her attacks with a weak bow that only inconveniences the majority of your foes, has a very low firing rate, and is extremely liable to shoot you in the back (and when that happens, the arrows take out a significant chunk of your health bar).
Lampshaded in dialog by the Prince:**

Figure 5. The specific statement that the classifier found to be the most negative for females, with a score of -0.94.

These statements are unquestionably negative. However, do they deserve an almost perfect score on the negativity spectrum? I would argue that compared to other statements that received a less extreme negativity score, they do not. Let's take a look at the statements that I personally found to be the most negative.

Unfortunately in Metal Gear Solid 2: Sons of Liberty, the person you need to call to save your game is Rose: Raiden's clingy, obnoxious, emotional, whiny, and loquacious significant other who feels the need to inflict an entire novel of sitcom-level relationship drama on you every single time you want to save. While you can skip it, they're so ridiculously long that this still takes longer than it did to actually listen to Mei Ling's proverbs in the previous game. If that's not enough, there are also many, many, many times where she just radios you up for the hell of it or butts into a conversation you're having with Colonel Campbell. It will get to the point where just hearing her say "Jack" will make you twitch.

Figure 6. The specific statement that the author found to be the most negative for females, with a score of -0.63.

Patch the doctor is annoying for a different reason. He heals your sick pinata (probably after they've gotten into a fight courtesy of Leafos). However, another character, Dastardos, will "euthanize" a sick animal if Patch doesn't make it in time. If both doctor and Dastardos appear at the same time though, Patch will stop what he's doing and let Dastardos kill your animal – right under the good doctor's nose. Feel free to beat Patch with your spade at this point; it's the least he deserves.

Figure 7. The specific statement that the author found to be the most negative for males, with a score of -0.75.

First, we will analyze the statement in Figure 6. I flagged this statement as particularly negative because of the liberal use of phrases that I found to be somewhat gendered: 'clingy', 'emotional', 'whiny', 'an entire novel of sitcom-level relationship drama', and 'just hearing her say "Jack" will make you twitch'. However, despite the numerous negative phrases, the sentiment analyzer gave this statement a score of -0.63, nowhere near the most negative statement in the set. While this is subjective, it seems reasonable to believe that more humans, when presented with the statements from Figures 5 and 6, would agree that the latter is a more negative statement.

Similarly, for the statement in Figure 7, I flagged it because of the mention of physical violence against the character: "Feel free to beat Patch with your spade at this point; it's the least he deserves". This was the most blatant mention of physical violence on the whole page, and notably, there were no mentions of physical violence in the statements made about female characters. While the phrase itself is probably supposed to be dry humor in the context of the statement, subjectively, most people would likely agree that it is still more violent, and therefore more negative than any subset of the statement in Figure 4, despite the fact that it was scored -0.75 by the sentiment analyzer.

Although these are subjective observations, from a human perspective, it calls the accuracy of the sentiment analyzer's scoring system into question. The sentiment analyzer does not necessarily represent human opinion; certainly, it is based on an algorithm that returns an objective result, but objectivity is not always the same as truth. So while it is certainly interesting that based on the scores the sentiment analyzer returned, female characters are guaranteed to receive a higher proportion of 'very negative' comments, the model's definition and the human definition of 'very negative' seem to vary significantly.

Conclusion

While on first glance, the sentiment analysis model seemed to perform with about an 80% accuracy rate on the TV Tropes corpus, we saw that on closer observation, the scoring of the statements did not seem to accurately reflect their true negativity. We therefore could not provide a concrete answer to our initial question regarding how female helper characters in video

games are treated relative to male helper characters. This may seem like method failure, but I believe that it serves as a warning against using 'one size fits all' machine learning models to make assumptions. Just as facial recognition algorithms are poor at classifying people of color, a standard use sentiment analysis model is perhaps not the best tool to use on a corpus full of sarcasm and colloquialisms.

Bibliography

Beck, Victoria Simpson, et al. "Violence Against Women in Video Games." *Journal of Interpersonal Violence*, vol. 27, no. 15, 2012, pp. 3016–3031., doi:10.1177/0886260512441078.

Bolukbasi, Tolga, et al. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*, pages 4349–4357.

Cassell, Justine and Jenkins, Henry. (2000). *From Barbie to Mortal Kombat: Gender and Computer Games*. Massachusetts: MIT Press.

Consalvo, Mia. "Confronting Toxic Gamer Culture: A Challenge for Feminist Game Studies Scholars." *Ada: A Journal of Gender, New Media, and Technology*, no. 1, 15 Nov. 2012, doi:10.7264/N33X84KH.

Cote, Amanda C. "Writing 'Gamers.'" *Games and Culture*, vol. 13, no. 5, 2015, pp. 479–503., doi:10.1177/1555412015624742.

Dill, Karen E., et al. "Effects of Exposure to Sex-Stereotyped Video Game Characters on Tolerance of Sexual Harassment." *Journal of Experimental Social Psychology*, vol. 44, no. 5, Sept. 2008, pp. 1402–1408., doi:10.1016/j.jesp.2008.06.002.

Grant, Adam, and Sheryl Sandberg. "Madam C.E.O., Get Me a Coffee." *The New York Times*, The New York Times, 6 Feb. 2015, https://www.nytimes.com/2015/02/08/opinion/sunday/sheryl-sandberg-and-adam-grant-on-women-doing-office-housework.html?_r=2.

Gupta, Shashank. "Sentiment Analysis: Concept, Analysis and Applications." *Medium*, Towards Data Science, 19 Jan. 2018, <https://towardsdatascience.com/sentiment-analysis-concept-analysis-and-applications-6c94d6f58c17>.

Hoyle, Alexander Miserlis, et al. "Unsupervised Discovery of Gendered Language through Latent-Variable Modeling." *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 11 June 2019, doi:10.18653/v1/p19-1167.

Jansz, Jeroen, and Raynel G. Martis. "The Lara Phenomenon: Powerful Female Characters in Video Games." *Sex Roles*, vol. 56, no. 3-4, Feb. 2007, pp. 141–148., doi:10.1007/s11199-006-9158-0.

Lowen, Linda. "Which Jobs Are Most Often Performed by Women?" *ThoughtCo*, ThoughtCo, 17 Sept. 2017, <https://www.thoughtco.com/women-work-traditional-careers-3534385>.

Kanter, Rosabeth Moss. *Men and Women of the Corporation*. Basic Books, 2010.

MacCallum-Stewart, Esther. "‘Take That, Bitches!’ Refiguring Lara Croft in Feminist Game Narratives." *Game Studies*, vol. 14, no. 2, Dec. 2014, <http://gamestudies.org/1402/articles/maccallumstewart>.

"Median weekly earnings of full-time wage and salary workers by detailed occupation and sex." 2016. Bureau of Labor Statistics, U.S. Department of Labor.

"Pearson's Chi-Squared Test." *Wikipedia*, Wikimedia Foundation, 8 Dec. 2019, https://en.wikipedia.org/wiki/Pearson's_chi-squared_test.

"Propp's Dramatis Personae." *Changing Minds*, changingminds.org/disciplines/storytelling/characters/propp_personae.htm.

Quandt, Thorsten. *Multiplayer: the Social Aspects of Digital Gaming*. Routledge, Taylor & Francis Group, 2014.

Snyder, Kieran. "Women Should Watch out for This One Word in Their Reviews." *Fortune*, Fortune, 3 Dec. 2014, <https://fortune.com/2014/08/26/performance-review-gender-bias/>.

Taylor, Nicholas, Jenson, Jen & de Castell, Suzanne (2009): Cheerleaders/booth babes/ Halo hoes: progaming, gender and jobs for the boys, *Digital Creativity*, 20:4, 239-252

The Bible. Authorized King James Version, Oxford UP, 1998.

ZamTheWriter. "Just 24 Female CEOs Lead the Companies on the 2018 Fortune 500-Fewer than Last Year." *CNBC*, CNBC, 21 May 2018, <https://www.cnbc.com/2018/05/21/2018s-fortune-500-companies-have-just-24-female-ceos.html>