

The Attractive, the Unattractive, and the Ugly: Appearance Adjectives and Gender

Introduction

The concept of “attractiveness” is all but inseparable from the concept of gender. What physical beauty is and how it—or its lack thereof—is perceived is influenced by social context, not limited to but absolutely including gender. What is seen as “attractive” in one gender might be seen as “ugly” in another. Though appearance has no bearing on a person’s character, it can absolutely affect the way their character is perceived. There are archetypes in art such as the “ugly stepsister,” in which moral values are ascribed to appearance, but this can extend past the realm of fiction and into reality. One striking example is that women who dressed in more revealing clothing in the workplace were perceived as less competent, despite clothing having no direct impact upon the quality of their work, suggesting that perceived attractiveness can influence perceptions of unrelated parts of an individual’s character (Wookey et al.).

One of the places where this relationship between attractiveness, gender, and values is especially visible is in the realm of language. For example, words like “handsome” and “beautiful” describe very heavily gendered kinds of attractiveness. These words being so heavily gendered applies expectations about what kind of appearance a person of a certain gender “should” be. A woman should be “beautiful,” but it may be considered odd for her to be “handsome”; a man should be “handsome,” but it may be seen as odd for him to be “beautiful.” In this way, language can carry implications about the people that language is used for.

Along these same lines, discussing the “women as dessert” metaphor, Caitlin Hines suggests that the way women are referred to, in her case via comparison to desserts, carries with it assumptions about what they ought to be: “sweet,” “compliant,” and “attractive.” Quoting Sally McConnell-Ginet, Hines suggests that languages “both encode and perpetuate speakers’

beliefs and attitudes,” and that this is very present with regards to gender (146). This influence of societal attitudes about gender affecting the language used for each gender is also present in discussions of attractiveness. In a study of university students, Fatima Al Qaisiya found that “the language of describing the physical appearance is affected by the genders and the features to be described,” finding that the gender of both the speaker and the subject seemed to correlate with the terms chosen to describe their appearance (52). To understand societal perceptions of gender, understanding the language used in reference to one gender or another can be illuminating.

Cumulatively, previous research in the topic suggests both that how appearances are perceived can affect perception of a person’s character, and that there are differences in the way attractiveness is perceived across genders. However, while there has been a significant amount of research regarding gender, appearance, and language, there appears to be a hole in research where corpora are concerned. The use of corpora is valuable in studying this kind of language due both to their scale and the fact that the samples within a corpora already exist, providing a measure of how people speak and write without the effects of an experimental context. I seek to fill this gap in the research by evaluating corpora for the use of several appearance-related adjectives, “attractive,” “unattractive,” and “ugly,” collocated with gendered nouns. Knowing that there is a body of research to support both the idea that the way people’s appearances are discussed differs by gender and that much of this research has been focused upon women, it follows that in day-to-day speech and writing, the use of adjectives related to appearance may be more commonly used for women than men.

Methods

This experiment was conducted using collocate searches in two of Mark Davies' corpora: The Corpus of Contemporary American English (COCA) and The iWeb Corpus (iWeb). COCA and iWeb were selected in the hopes of surveying a broad set of samples. Though COCA spans a wider timespan from 1990 to 2019, most of its data is from media that has been highly edited, including movies, television shows, and books. In contrast, iWeb draws exclusively from web sources from 2017. These sources, such as forums and blogs, seemed more likely to be personal, informal writing from individuals. iWeb was therefore selected to attempt to provide more unedited, uncensored samples. Together, these corpora were meant to provide a more comprehensive dataset.

As for the words to be studied, "attractive," "unattractive," and "ugly" were selected. These words were chosen for their less gendered connotations compared to other words measuring attractiveness, such as "handsome" or "beautiful," in the hopes of providing results less colored by the target words' connotations. To compare positive and negative descriptors, "attractive" and "unattractive" were chosen. "Ugly" was specifically selected as a less formal synonym.

To measure the number of usages for each term, a collocate search was performed for each word, tracking nouns that commonly appeared after each adjective. After this, all gendered nouns among the top 100 most frequent collocates (e.g. "man," "woman," "girl," "boy," etc.) and their number of usages were recorded. These usages were added together for each gender and compared to find which gender, if any, these words were used more often for.

All searches limited results to collocates within two words after the target. While this was necessary to capture most instances of the word describing a person while excluding edge-

cases, it also ended up including phrases like “she is attractive to men,” where the subject of the sentence is feminine, but the search counts “men” as the collocate. To attempt to capture this nuance, a sample of 100 (or, up to 100 when the search returned fewer than 100 entries) of the most common masculine and feminine terms for each search were examined, and the number of instances in which the gender of the recipient of the adjective was the opposite of the collocate was counted. For ease, these will be referred to as “false masculines” for examples where results were mislabeled as masculine while actually referring to a female subject, and “false feminines” when an example was mislabeled as feminine while referring to a male subject. While it is out of the scope of this paper to correct every single example, this should provide insight into the margin of error of this method of searching.

In addition to quantitative analysis, the number of phrases making reference to sex or sexuality were also counted from this same sample of up to 100 in an attempt to capture possible qualitative differences in discussion of gender. The same searches were repeated across both corpora for each word measured. They were then compiled to make comparisons between usages of appearance-related adjectives for women and men.

Results

Gendered Usages of “Attractive”

In COCA, the top 100 collocates of the word “attractive” contained 10 gendered nouns, 8 of which were feminine and 2 masculine (fig. 1). In total, there were 1636 usages of “attractive” with feminine collocates, and 159 with masculine collocates, a difference of over 1000%. 91.11% of the total gendered collocate usages for the word were feminine and 8.86% were masculine.

Fig. 1: Gendered Collocates of “Attractive” in COCA

Word	Frequency	Ranking in Search	Gender
Woman	787	1	Feminine
Women	437	2	Feminine
Girl	135	7	Feminine
Men	134	8	Masculine
Female	101	9	Feminine
Girls	70	13	Feminine
Lady	51	20	Feminine
Females	40	29	Feminine
Males	25	38	Masculine
Ladies	15	60	Feminine

In this case, “women” was compared to “men” rather than “woman.” This was done due to “woman” being a singular noun, which could affect the number of instances where the collocate’s gender differed from the recipient of the adjective, resulting in misleading data. Out of a sample of 100, 12 of the examples of “women” as a collocate of “attractive” were in contexts that were referring to men, typically with the construction “attractive to women.” From the sample of “men,” 43 examples were actually referring to women, typically with the construction “attractive to men.” This means that 43% of the sample from “men” were false masculines and 12% of the examples from “women” were false feminines. This implies that the true difference in the number of usages for each gender may actually be more extreme than it initially appears. In addition, 8 references to sex and sexuality were found in the sample of “women,” and 6 from the sample of “men.”

In iWeb, there were 7 gendered nouns in the top 100 collocates of “attractive,” all of which were feminine, totalling to 11519 usages (fig. 2). Of the sample of 100 examples containing “women” as a collocate of “attractive,” 18 of them were false feminines, and 6 examples mentioned sex or sexuality. Though there were more false feminines here than in the results for “women” in COCA, there were still significantly fewer false feminines than there were false masculines for “men” in COCA.

Fig. 2: Gendered Collocates of “Attractive” in iWeb			
Word	Frequency	Ranking in Search	Gender
Women	4485	2	Feminine
Woman	3360	4	Feminine
Girl	1055	22	Feminine
Female	909	30	Feminine
Girls	880	31	Feminine
Lady	530	48	Feminine
Females	300	79	Feminine

Gendered Usages of “Unattractive”

The top 100 collocates of “unattractive” in COCA included 8 gendered nouns, 7 of which were feminine. In total, there were 85 usages of “unattractive” with feminine collocates, and 15 with masculine (fig. 3). Due to a lack of examples, the sample from “women” consisted of 35 examples, and the sample from “men” consisted of 15. Of these, none of the masculine examples mentioned sex or sexuality, but 6 of the feminine examples did. There were no false feminines, but 6 false masculines, representing 40% of the dataset, again suggesting that the numbers may be even more skewed towards feminine usages.

Fig. 3: Gendered Collocates of “Unattractive” in COCA			
Word	Frequency	Ranking in Search	Gender
Women	35	1	Feminine
Woman	27	2	Feminine
Men	15	3	Masculine
Girls	10	5	Feminine
Female	5	15	Feminine
Girls	4	22	Feminine
Females	2	53	Feminine
Ladies	2	66	Feminine

In iWeb, the results were far more even. Of the 12 gendered noun collocates of “unattractive,” 6 were feminine and 6 were masculine (fig. 4). However, feminine examples comprised 567 usages, 60.38%, and masculine usages totaled 372, 39.62%. In the sample from “women,” there were 22 false feminines and 7 mentions of sex or sexuality. In the sample of “men,” there were 32 false masculines and 3 mentions of sex or sexuality. The fewer false feminines than false masculines is in line with the rest of the data and continues to suggest that the true number of feminine usages may be higher than what the rankings initially suggest.

Fig. 4: Gendered Collocates of “Unattractive” in iWeb			
Word	Frequency	Ranking in Search	Gender
Women	280	1	Feminine
Men	138	2	Masculine
Woman	135	3	Feminine

Man	106	6	Masculine
Girl	64	7	Feminine
Guy	48	9	Masculine
Guys	40	14	Masculine
Female	35	18	Feminine
Girls	35	19	Feminine
Male	24	28	Masculine
Females	18	41	Feminine
Males	16	46	Masculine

Gendered Usages of “Ugly”

The top 100 collocates for “ugly” from COCA contained 11 gendered collocates, 8 of which were feminine and 3 of which were masculine (fig. 5). In total, feminine nouns accounted for 149 usages and 71.98% of all gendered collocates, whereas masculine nouns represented 58 usages and 28.02% of the data.

Fig. 5: Gendered Collocates of “Ugly” in COCA			
Word	Frequency	Ranking in Search	Gender
Bitch	58	14	Feminine
Bastard	34	32	Masculine
Stepsister	20	52	Feminine
Sisters	14	77	Feminine
Bitches	13	80	Feminine
Dudes	12	88	Masculine
Bastards	12	89	Masculine

Hag	11	97	Feminine
Chicks	11	98	Feminine
Whore	11	99	Feminine
Chick	11	100	Feminine

Due to the small sample size, all masculine examples were compared against the examples from “bitch,” for a total of 58 examples for each gender. There were no instances of either false feminines or false masculines. Of the sample of 58, 5 of the feminine examples mentioned sex or sexuality. There were no mentions of sex or sexuality in the masculine examples. On the whole, the gendered collocates for “ugly” in COCA were significantly different from those of “attractive” and “unattractive,” having more vulgar results such as “bitch,” “bastard,” and “whore.”

For “ugly,” iWeb contained 10 gendered nouns within the top 100 collocates, 9 of which were feminine, and 1 of which was masculine (fig. 6). Feminine collocates represented 95.96% of the total gendered nouns with 2066 usages. Masculine collocates made up the remaining 4.04%, with 87 usages. Due to the lack of examples of masculine collocates, the 87 usages of “dude” were compared to 100 uses of “woman.” There was 1 false masculine, but no false feminines. There were 3 references to sex or sexuality in the feminine sample and 9 in the masculine examples, a significant difference from the rest of the data where there tended to be more references to sex or sexuality for feminine collocates.


Fig. 6: Gendered Collocates of “Ugly” in iWeb			
Word	Frequency	Ranking in Search	Gender
Woman	466	15	Feminine

Girl	406	22	Feminine
Girls	316	25	Feminine
Sisters	275	30	Feminine
Sister	195	42	Feminine
Bitch	150	57	Feminine
Stepsister	95	76	Feminine
Dude	87	83	Masculine
Chicks	85	89	Feminine
Chick	78	99	Feminine

Totals

Across both corpora, there were a total of 16663 usages of gendered noun collocates of “attractive,” “unattractive,” and “ugly.” Of these, 16022 instances were of feminine nouns, and 641 were masculine. Cumulatively, feminine usages represented 96.15% of the dataset, and masculine usages represented 3.85% (fig. 7). Feminine collocates were even more common in the combined dataset for “attractive” across both corpora, representing 98.8% of the total gendered noun usage within the top 100 collocates. When combined, the negative terms, “unattractive” and “ugly,” presented far closer results, but feminine uses still made up 84.35% of the data, and masculine only 15.65%. This suggests that while there is a feminine lean to discussion of both attractiveness and unattractiveness, there is a stronger lean in the former than the latter.

Fig. 7: Gendered Collocates of Positive and Negative Appearance-Related Words by Percentage



Despite the different make-up of each corpus's sources, the results remained relatively consistent across both COCA and iWeb, with feminine collocates being severely overrepresented (fig. 8, fig. 9). COCA had proportionately more instances of masculine usages for "attractive" and "ugly," but proportionately fewer examples of masculine usages for "unattractive." By far, the more equivalently used terms were negative, with iWeb's gendered collocates for "unattractive" being the most equivalent of the entire dataset (fig. 9).

Fig. 8: Gendered Collocates of Appearance-Related Words in COCA


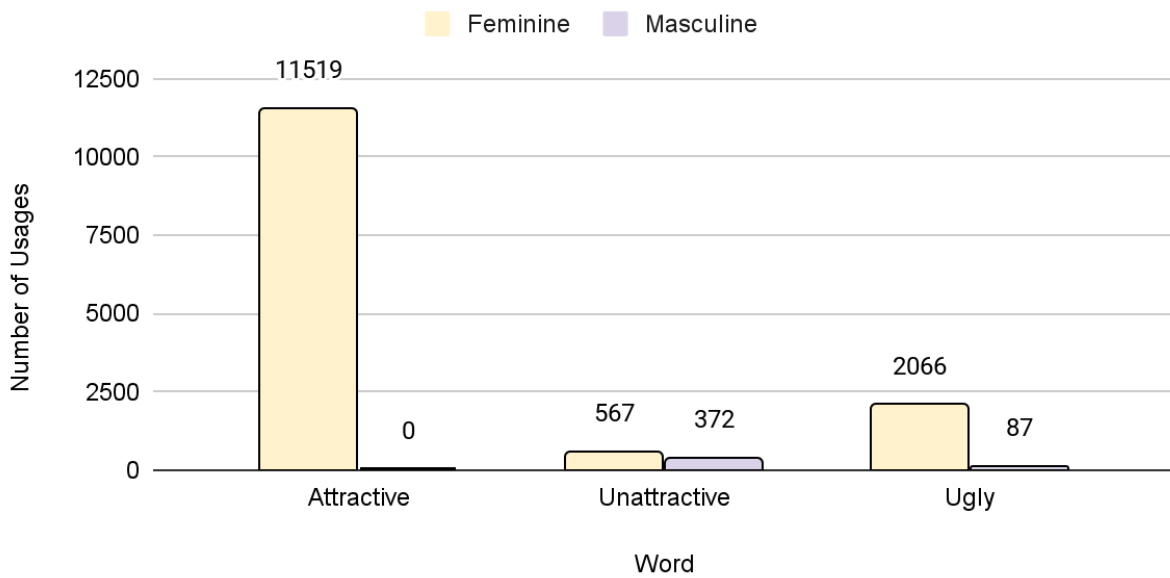



Fig. 9: Gendered Collocates of Appearance-Related Words in iWeb



As for the content of the samples, on the whole there was a higher percentage of sexual references in the feminine collocates, though “ugly” in iWeb was a very prominent outlier (fig.


10). However, these results are skewed by the comparatively few instances of masculine collocates across the dataset.

Fig. 10: Sexual References by Percentage



Finally, there were consistently a higher percentage of false masculines than false feminines when there were sufficient masculine examples to examine, with the exceptions of “ugly” in COCA, where there were no false masculines or false feminines, and “attractive” in “iWeb,” where there were no masculine collocates in the top 100 results to examine (fig. 11).

Fig. 11: False Feminines and False Masculines by Percentage



Discussion

Overall, the results of this experiment were extremely consistent with prior research on gendered language. Women were not only referenced more often than men for all three terms but appeared more often by almost 25x. This is a massive difference, and one that is statistically unlikely to be explained away by instances where the recipient of the target adjective did not match the gender of the collocate noun. In fact, the comparatively large number of false masculine samples across the majority of the datasets examined suggests that the totals could be even more drastically slanted towards feminine usages than these numbers suggest.

None of the words measured were used for only men or women, implying that “attractive,” “unattractive,” and “ugly,” are not necessarily words with gendered connotations, but rather words that are applied inequitably across genders. Extrapolating from these results, it seems likely that the topic of appearances is seen as relevant in discussion of women in a way it generally is not for men. The fact that results were consistent between the positive descriptor of

appearance, “attractive,” and the two negative descriptors implies that both negative and positive appearances of women are something seen as worth commenting on in a way that isn’t nearly as extreme for men.

Interestingly, negative descriptors of appearance were used proportionally more often for men than positive descriptors. This could suggest that, in general, appearances perceived negatively are seen as worth commenting on across all genders, but that it is seen as more significant when a man is “unattractive” or “ugly” as opposed to “attractive,” though further investigation would be needed to come to a decisive conclusion. Deeper examination of the samples for “unattractive” and “ugly” may reveal qualitative differences that were out of the scope of this experiment.

In addition to the quantitative results, it would be disingenuous to not acknowledge the character of many of the examples. In particular, the results from iWeb were less edited, and were often especially negative in their discussion of women in a way that the generally more polished examples from COCA were not. To demonstrate this, a few especially grim examples from iWeb (Davies):

- “Fat, unattractive women, especially the younger ones, are some of the most prideful, insufferable [...]”
- “Also, all unattractive women with IQs below 90 could be irreversibly sterilized also.”
- “being a bitch, in-and-of-itself, is actually very unattractive to men.”
- “Being equitable or fair doesn't make you more attractive to women it just makes you cheap, and, likely, still single.”
- “Jobs and education make women ugly and unattractive”

iWeb's data is from 2017. Many of these examples are statements that would have been absolutely unacceptable in face-to-face, polite conversation at the time they were written. Though the sources of each example were not analyzed in detail, a significant portion of them were from sites dedicated to pick-up artistry and dating. It's likely the difference between the character of the two corpora was due to the sources each draws from. Because iWeb draws from online sources only, there are more examples from forums and blogs, which are by their nature less edited and more casual. Digital sources grant speakers some degree of anonymity, and may begin to account for how inflammatory some of the examples were. COCA, on the other hand, is a balanced corpus drawing from an equivalent number of sources across several different mediums, including ones that would undergo multiple rounds of editing, including TV/film scripts and books. Even then, the COCA's web samples weren't immune to similar kinds of rhetoric (Davies):

- 2012, WEB: "I don't know how likely it is I wouldn't be inclined to have sex with even the ugliest and most unattractive of women. After all, below the waist, it's all the same."
- 2012, WEB: "Young women are a hotter commodity in the workplace because having attractive young women around is good for many different types of businesses"

The character of many of these samples, combined with the higher rate of sexual references for the feminine examples, creates a sense of objectification that just was not nearly as present in the masculine examples. Further examples from COCA and iWeb make this clear:

- COCA, 2012, WEB: "Some very sexually attractive men are unconsciously waving great big red flags about their attitudes that are offputting"
- iWeb: [in reference to male pornography actors] "just regular guys off the street. They are not sleazy, unattractive men, generally they are young, good-looking, and built."

- iWeb: “attractiveness rarely matters. I can have the best sex with the most unattractive men”

The magnitude of difference in both a qualitative and quantitative analysis of the results from both iWeb and COCA across all three terms have concerning implications for both the societal expectations upon women regarding appearance, but also attitudes towards women in general, removed from the concept of attractiveness. The consistency of the results, though pulled from two different corpora, supports the idea that this gender inequity in discussion of appearance is not merely an issue of one group or time period.

Limitations

Given the tools used to conduct this study, there are some limitations to the data gathered. Most pressingly, I was unable to exclude usages of the target word where the adjective was applied to the opposite gender of the collocate it was listed under. Further studies could refine the search method to exclude these usages, or simply look more thoroughly through all results, rather than a sample. Due to the use of a random sample to attempt to capture this nuance, it is also possible to have ended up with samples that misrepresent the whole. Further research into the numbers could involve looking at a larger number of examples from a broader set of terms. Additionally, this study limited results to the top 100 most frequent collocates of the target words. Expanded lists could provide more accurate results.

Beyond the methodology of the searches themselves, deciding whether or not an example referenced sex or sexuality was done purely by the author. These observations and some additional observations made using individual examples were qualitative and therefore subjective, and may not reflect the opinions of all readers.

Finally, it is pertinent to state that while this study focused solely on men and women, this is not to imply that this binary is all there is. There is a lack of representation of nonbinary identities across all sorts of sources, making the way gendered language affects anyone other than men and women difficult to measure via corpora. This is a large hole in the current understanding of gendered language. It would be especially interesting to compare how and when terms for attractiveness with especially gendered connotations, such as “handsome” or “beautiful,” are applied to nonbinary individuals. Further research into the intersection of gender and language has plenty of room to complicate these results.

Conclusion

The gendering of language is often more complex than words that directly reference gender. Even the terms we see as “ungendered” can still be applied inequitably. Seemingly positive terms can, in this way, come to represent wider differences in societal ideals applied to different social groups. This data leads us to ask: if women are talked about as “attractive” so much more often than men, what does this mean for “unattractive” women? If it is so important for women to be “attractive,” what does that mean when some is unable or simply does not want to follow that social norm?

These results have real-world implications, especially for women. I feel the data makes it clear that there is a significantly higher amount of conversation surrounding attractiveness and unattractiveness specifically regarding women, implying further scrutiny is applied to women regarding attractiveness than men. Whether or not someone personally believes appearances matter, the sheer inequity in the amount the topic is discussed for one gender over another can lead to the internalization of these ideas in even the best-intentioned people. Recognition of these

sorts of patterns is an essential step towards understanding where societal ideals lie, and begin to change them for the better.

Works Cited

- Al Qaisiya, Fatima. "Language and Gender in Describing Physical Appearance." *International Journal of Language and Linguistics*, vol. 4, no. 1, Mar. 2017, pp. 49-54.
https://www.ijllnet.com/journals/Vol_4_No_1_March_2017/7.pdf
- Davies, Mark. *The Corpus of Contemporary American English*. 2008, www.english-corpora.org/coca/.
- Davies, Mark. *The iWeb Corpus*. 2007, <https://www.english-corpora.org/iWeb/>.
- Hines, Caitlin. "Rebaking the Pie: The WOMAN as DESSERT Metaphor." *Reinventing Identities : The Gendered Self in Discourse*, edited by Mary Bucholtz, et al.. Oxford University Press, 1999, pp. 145-162. *EBSCOhost*, search.ebscohost.com/login.aspx?direct=true&db=e000xna&AN=144010&site=ehost-live&scope=site.
- Wookey, Melissa L., et al. "Effects of a Sexy Appearance on Perceived Competence of Women." *Journal of Social Psychology*, vol. 149, no. 1, Feb. 2009, pp. 116–18. *EBSCOhost*, <https://doi.org/10.3200/SOCP.149.1.116-118>.