

How-to Guides for Specific Audiences: A Corpus and Initial Findings

Nicola Fanton

he/they

Agnieszka Falenska

University of Stuttgart

Michael Roth

Institute for Natural Language Processing

{firstname.lastname}@ims.uni-stuttgart.de

Abstract

Instructional texts for specific target groups should ideally take into account the prior knowledge and needs of the readers in order to guide them efficiently to their desired goals. However, targeting specific groups also carries the risk of reflecting disparate social norms and subtle stereotypes. In this paper, we investigate the extent to which how-to guides from one particular platform, wikiHow, differ in practice depending on the intended audience. We conduct two case studies in which we examine qualitative features of texts written for specific audiences. In a generalization study, we investigate which differences can also be systematically demonstrated using computational methods. The results of our studies show that guides from wikiHow, like other text genres, are subject to subtle biases. We aim to raise awareness of these inequalities as a first step to addressing them in future work.

1 Introduction

How-to guides provide practical instructions that help humans to achieve specific goals. In the past decades, such guides also attracted increasing interest in NLP and AI research (Branavan et al., 2009; Chu et al., 2017; Anthonio et al., 2020). Resources such as wikiHow,¹ a collaboratively edited online platform for instructional texts, make it possible to scale research efforts to hundreds of thousands of articles. By covering an ever-increasing number of guides, including niche topics and articles for minority groups, there is also an increasing risk of perpetuating stereotypes and jeopardizing general accessibility. In fact, we notice that wikiHow already contains articles written for specific target groups as well as articles that exist in different versions for different audiences. As an example, Table 1 shows two articles with the same *title*, “Act Like a Kid Again”, one with the *indicator* ‘(Girls)’ and one with ‘(Boys)’.

¹www.wikihow.com

Act Like a Kid Again (Girls)

Eat well and exercise, but don’t obsess about your body. Be healthy without stressing too much about it. (...) Generally, **go for lots of fruits and veggies**. And even though kids love sugar, **don’t eat too much** of it!

Act Like a Kid Again (Boys)

Eat your childhood favorite food. Recollect **every snack**, chocolates, ice cream, candy bars, cotton candy and **everything that you loved as a kid** or would make you feel pampered. Eat as per your capacity as too much at once may make you feel uncomfortable.

Table 1: Two versions of the same guide in wikiHow.

Among other things, we find that such articles dramatically differ in terms of details. For example, the texts highlighted in Table 1 vary in how much they focus on issues potentially related to body images. As such, the articles reflect *disparate standards*, which ultimately may contribute to discrimination (Prentice and Carranza, 2002). The specific example can also be linked to observations of gender differences in weight concerns from psychology (Dougherty et al., 2022), which might represent a reason for *disparate treatment*. On the surface, it is not always possible to say exactly why there are certain differences in articles for specific audiences. However, through qualitative and quantitative comparisons on the linguistic level, we can at least determine what types of differences are present and to what extent they can be systematically identified. In this sense, we aim to contribute to questions about biases and fairness in data and, at the same time, connect to related research in psychology and other social sciences.

There already exists a large body of research that examines biases and stereotypes in NLP data and, likewise, how-to guides from wikiHow have been

used as training material for a variety of language processing tasks (§2). However, previous studies have not explicitly looked into issues related to bias in the wikiHow data. As a first step towards addressing this gap, we create our own sub-corpora of how-to guides, which let us investigate differences across articles for specific target groups (§3).

We perform two case studies and a generalization study on our collected data: In the first study, we identify a number of articles that exist in multiple variants for different target groups and examine them in terms of distinctive content and linguistic characteristics (§4). As a second case study, we explicitly examine how far topics covered for specific target groups differ from each other (§5). Finally, we investigate whether the qualitative findings from our case studies can be validated quantitatively and generalized to our whole corpus using computational modeling (§6).

In summary, we find systematic differences between articles for specific groups in terms of topic, style, and content. We conclude the paper with a discussion of these findings and point out links to existing work in the social sciences (§7).

2 Related Work

We summarize existing work on the three strains of research that this paper builds on: wikiHow as a data source (§2.1), subtle biases in datasets (§2.2), as well as understanding the characteristics of texts that target specific audiences (§2.3).

2.1 wikiHow as a Data Source

wikiHow is a prominent data source for a variety of tasks, including summarization (Koupaei and Wang, 2018), goal-step inference (Zhang et al., 2020), and question answering (Cai et al., 2022). By exploiting the revision history of wikiHow, Anthonio et al. (2020) created **wikiHowToImprove**, which has been used to better understand phenomena related to the (re-)writing process of how-to guides (Roth and Anthonio, 2021; Anthonio et al., 2022). Writing, but especially revising, instructions should presumably take into account the readers’ context, perspective and knowledge about the domain and the world. The need for clarification stands prominently out as a main purpose of the refinements of wikiHow guides (Bhat et al., 2020). It has been shown that while annotators tend to agree that “revised means better”, the disagreements can be caused by differences in common knowledge

and intuitions (Anthonio and Roth, 2020). As specific phenomena, previous work studied implicit references and lexical vagueness (Anthonio and Roth, 2021; Debnath and Roth, 2021). However, none of the aforementioned studies accounted for audience-specific differences. This work takes a first step to close this gap.

2.2 Subtle Biases in Datasets

Diagnosing the presence of biases in data is one of the crucial steps in diminishing the spread of harmful stereotypes. This work contributes to the research on *subtle biases*, i.e., textual patterns that implicitly reflect societal power asymmetries. Such biases are embedded in specific linguistic phenomena (e.g., masculine generics; Swim et al., 2004) or in inequalities in how people from different demographic groups are represented (e.g., emphasizing the romantic relationships in the bibliographies of women; Wagner et al., 2015). Moreover, they can be frequent even in domains where blatant stereotypes and openly expressing beliefs about social hierarchies is generally considered inappropriate (Cervone et al., 2021). For example, there is a long line of work analyzing subtle stereotypes in Wikipedia (Callahan and Herring, 2011; Reagle and Rhue, 2011; Konieczny and Klein, 2018; Schmahl et al., 2020, among others), where the lack of diversity represents an issue already at the level of the editors’ community (Lam et al., 2011). Beyond notability for representation itself, linguistic aspects in Wikipedia show a remarkable disparity concerning biographies of men and women, both in terms of topics and polarity of abstract terminology (Wagner et al., 2016). Such inequalities do not pertain only to biographies but find systemic correspondence in all domains and across languages (Falenska and Çetinoğlu, 2021).

To the best of our knowledge, the presence of subtle stereotypes in wikiHow has not yet been investigated. However, the guides from this platform are a valuable entry point for studying bias, as they are produced by a community of contributors and by experts² suggesting how to perform activities. In other words, given the different purposes of the platforms, while Wikipedia data is rather descriptive, wikiHow data features instructional texts that potentially differ depending on the audience.

²<https://www.wikihow.com/Experts>

2.3 Different Audiences

The mind of the readers features a priori goals that affect the understanding of written texts (Fum et al., 1986). However, the goals and knowledge of different (groups of) people may vary. An example of work that considers different readers’ expertise regards title generation (Senda and Shinohara, 2002). In that work, less expert readers were found to be tentatively more influenced by effective titles. Consequently, a system for revising titles accounting for the readers’ expertise has been proposed (Senda et al., 2004). As such, that contribution indicates the importance of considering the target audience for efficient communication. Additionally, different audiences can understand to different extents technical terminology (Senda et al., 2006; Elhadad and Sutaria, 2007) and causation (Siddharthan and Katsos, 2010). Previous contributions accounted for different target groups also in the controllable text generation tasks of paraphrasing (Kajiwara et al., 2013), text simplification (Scarton and Specia, 2018; Sheang and Saggion, 2021), machine translation (Agrawal and Carpuat, 2019), and dictionary examples generation (He and Yiu, 2022).

3 Corpus Construction

As introduced in §2.1, wikiHowToImprove is a well-established data set derived from wikiHow and consisting of more than 246,000 how-to guides. In general, each guide consists of multiple revisions of an *article*, a fixed goal that is named in the *title*, and (optionally) an *indicator* that follows the title in parentheses (cf. Table 1). As we are interested in how-to guides for different target groups, we filter the data for indicators that specify a group of people as targets, which we also refer to as the *audience*. Table 2 lists the 20 most frequent indicators extracted from wikiHowToImprove.

Based on a manual grouping of these indicators, we find that 15 out of 20 indicators refer to attributes of performative gender and age (the remaining five are underlined in Table 2). Apart from their high frequency, both of these attributes are of interest to studies in the social sciences, in which they are often used as independent variables (Cortina et al., 2013; Cha and Weeden, 2014; Palència et al., 2014). Following a traditional binary setup, we distinguish two audiences based on gender, women (**W**) and men (**M**), and two audiences based on age, kids (**K**) and teens (**T**).³ For each type of audience,

³Note that while the selected audiences follow discrete

Rank	Indicator	#	Rank	Indicator	#
1	Girls	370	11	Guys	35
2	for Girls	284	12	for Women	35
3	for Kids	182	13	Women	34
4	Kids	114	14	<u>UK</u>	34
5	Teens	110	15	for Men	31
6	Teen Girls	100	16	<u>Christianity</u>	31
7	for Teens	73	17	Men	29
8	<u>USA</u>	49	18	<u>for Beginners</u>	29
9	for Guys	42	19	Boys	25
10	<u>Windows</u>	38	20	Teenage Girls	25

Table 2: Counts of the 20 most frequent indicators.

	W	M	K	T
Indicators	29	13	23	16
Articles	993	209	499	411
Sentences per article	40	50	29	43
Words per article	509	682	352	544

Table 3: The distribution of the indicators and of the articles for the target audience groups. Sentences and words are indicated via their median values by article.

we create a set of all indicators used and collect all corresponding guides by extracting the latest article versions from wikiHowToImprove.

Statistics of our corpus with audience-specific how-to guides are provided in Table 3. We note that there is a much higher number of indicators and articles for W than for M. In comparison, the number of articles and indicators for K and T are similar. With only 2,112 how-to guides in total, the corpus seems relatively small. However, the average length of articles ranges from 352 to 682 words, which adds up to a corpus size of more than one million words. Throughout this work, we refer to this dataset as wikiHowAudiences.⁴ Next, we approach it in its entirety with two case studies.

4 Case Study: Same Title, Different Audience

Our starting example from Table 1 includes two guides with the same title but different target indicators. Such guides outline the ultimate instances of instructions that are written for different audiences.

categories, we explicitly caution that individual readers can only be represented on a continuum.

⁴<https://github.com/mnfanton/wikiHowAudiences>

Women – Men		
BODY	11	Lose Belly Fat
INTERACT	11	Act on a Date
PRESENT	13	Dress Like a CEO
Kids – Teens		
GROWN-UP	3	Look Older
ADVICE	4	Balance School and Life
ACTIVITY	10	Apply Makeup

Table 4: Frequencies and examples of topical categories.

Therefore, we start our investigation by analyzing how often such cases occur in wikiHowAudiences, which topics they cover, and what differs between versions for specific target groups.

4.1 Guides Selection

First, we identify titles that occur more than once in wikiHowAudiences: 32 unique titles for W–M and 15 for K–T. Next, we group guides with the same title but different target audiences into pairs. A complete list of article titles in this subset can be found in Appendix A.1.

4.2 Guides Analysis

To understand which goals require audience-specific adaptations, we analyze the topics and articles of the filtered guides.

Topics. We start by manually investigating titles of the filtered pairs of guides. For this purpose, we assign each of them to one of three content-related categories. The categories were designed to cover all the titles while being as concrete as possible. An overview of all the categories and their examples is listed in Table 4.

We find that W–M instructions cover a relatively wide range of topics, from body-related activities (BODY), over interacting with other people (INTERACT), to self-presentation (PRESENT), which is the most frequent category. In contrast, among titles in K–T, we notice one clear pattern: all topics focus on issues that require different steps depending on the age of the target. Among them, we distinguish and report in ascending order of frequency articles about learning how to do activities for grown-ups or concerning the urge to grow old (GROWN-UP), advice related to the life of young people (ADVICE), and activities about oneself or the relation of oneself to others (ACTIVITY).

Length. Next, we check whether there are significant differences in terms of how detailed the instructions are for different target groups. We quantify this by simply measuring the length per article in words and sentences. We notice a considerable difference between K and T: the median length of articles for K is only 30 sentences and 346 words, while articles for T contain 98 sentences and 1081 words. In the case of W and M, we do not find such large differences in terms of average word (785 vs. 856) and sentence counts (59 vs. 62). Overall, the numbers reflect the patterns shown in Table 3 for the whole wikiHowAudiences data.

Content. Finally, we switch our attention to the actual content of the articles. As a simple measure of how similar two guides are, we consider their word overlap in both directions using BLEU score (Papineni et al., 2002).

Table 5 presents the articles with the lowest and highest word overlap in both analyzed groups. Interestingly in the case of W–M, both articles cover concepts related to BODY, namely clearing skin and recognizing an infection. Manual inspection of their content reveals that even in the case of the least overlapping articles, “Get Clear Skin”, slight differences can be noticed: W article includes more specific information as well as different usage of punctuation. In the case of most overlapping articles, “Recognize Chlamydia Symptoms”, the main difference comes from the vocabulary related to different body parts from body types. The high word overlap of these two versions is likely related to their introductions, which provide an interchangeable overview to the topic.

In the case of K–T, the least and most overlapping articles come from two different categories: ACTIVITY and GROWN-UP. The least overlapping pair, “Flirt”, is a case of two instructions that treat the same goal with different levels of complexity. For example, the matter of eye contact is described with one step in K and more than ten in T. The most overlapping articles, “Make Money”, can be an example of a content stalemate – for both target audiences, babysitting is the first suggested activity to achieve the profit goal. However, it is possible to notice differences in how this concept is contextualized for two groups: either in a list of activities or discussed with its implications and advantages.

W–M	Get Clear Skin (0.02 BLEU)	Recognize Chlamydia Symptom (0.69 BLEU)
W	Gently pat your face dry with a clean towel. Don't rub your face! This can irritate your skin more.	Chlamydia is a dangerous yet common and curable sexually transmitted infection (...)
M	Dry your face – but not roughly.	Chlamydia, specifically chlamydia trachomatis, is a common and curable but dangerous sexually transmitted infection (...)
K–T	Flirt (0.05 BLEU)	Make Money (0.59 BLEU)
K	Make eye contact. Both girls and boys love eye contact.	There are the traditional jobs like babysitting, shoveling snow, and doing chores around the house.
T	Make eye contact. Body language is a big part of flirting, and a big part of that is eye contact. Eye contact conveys intimacy (...)	Babysit for friends and family. One of the best ways for teenagers to make money and help out in the community is babysitting.

Table 5: Excerpts from the article pairs with the lowest (left) and highest (right) word overlap.

4.3 Summary

We exemplified three characteristics that can distinguish guides written for different audiences. First, the instructions written for K–T significantly differed in *length*. Next, we saw pairs of guides that varied in *style* (such as punctuation) and *content* (e.g., vocabulary in BODY articles). Some of the presented examples suggest that considering only simple content features could be enough to distinguish articles written for different audiences. However, such an approach could be insufficient in more complex cases, such as pairs of guides with high word overlap (see “Make Money”). We discuss these articles again in our generalization study (§6).

5 Case Study: “How To Be” Guides

In the previous section, we looked at how-to guides that occur in different versions for specific audiences. Such guides might concern particular goals that *require* being addressed in distinct ways. In this section, in contrast, we broaden the scope of analysis to explore other cases of differences in audience-specific instructions.

5.1 Guides Selection

The initial example from the introduction (see Table 1) explain how to perform like somebody the reader is presumably not. Inspired by this example, we investigate what other guides instruct their readers “how to be”. Concretely, we filter titles starting with the word ‘be’, which gives us 118 guides for W, 20 for M, 32 for K, and 30 for T.

	Completion(s)	Title
W	Popular Cute	Be <u>Popular</u> and Athletic Be <u>Cute</u> at School
M	Cool More	Be <u>Cool</u> in High School Be <u>More</u> Physically Attractive
K	Good	Be <u>Good</u> With Money
T	Good	Be a <u>Good</u> Friend

Table 6: The most frequent target-specific completions of “how to be” guides and examples of respective titles.

5.2 Guides Analysis

To understand which topics the “how to be” guides cover, we group them according to the first word that occurs after ‘be’ (henceforth the *completion*).⁵ Table 6 shows the most frequent completions for each target group and respective example titles.

Regarding K–T guides, we notice no clear pattern that would distinguish instructions based only on their titles. There is roughly the same number of how-to articles for K and T (32 vs. 30). Moreover, among the most frequent completions we commonly find the word ‘good’, followed by words such as ‘comfortable’, ‘less’, or ‘safe’.

In contrast, we find substantial differences for W–M. Specifically, we note that “how to be” guides are more common for W (12% of all articles for this target group) and for both audiences we find differing frequencies of completions: While W articles focus

⁵We ignore the articles ‘a’, ‘an’, and ‘the’.

on being ‘cute’ and ‘popular’ (9 guides), M articles put more emphasis on being ‘cool’ and ‘more’ (6 guides). Even though all the how-to guides refer to similar contexts (mostly related to school), we do not find mutual correspondence—there are no instructions for how to “be cool at school” for W and no guide for how to “be cute at school” for M.

5.3 Summary

In this section, we looked at a particular subset of wikiHowAudiences, namely guides with titles starting with the word ‘be’. We found that, in the case of W–M targets, the differences in instructions occur already at the level of goals that these guides describe. In other words, we saw examples of instructions where the information for which audience they were intended could be deduced strictly from their *titles*.

6 Generalization Study: Computational Approach

Our case studies show that, depending on the audience, there exist examples of articles that differ in terms of topic, length, style, and/or vocabulary. However, an open question is whether these are only individual cases or if such differences occur systematically. In this study, we investigate this question computationally and attempt to verify our observations on the basis of a larger dataset. For this purpose, we implement tentative characteristics in the form of features and models (§6.1), evaluate in a setting with our full sub-corpora (§6.2), discuss quantitative results (§6.3), and analyze qualitative findings (§6.4).

6.1 Models

Based on the findings from the two case studies, we define majority and length-based baselines and several simple logistic regression classifiers with different sets of features.

Baselines. We use a simple majority baseline that always assigns the most frequent class. We also implement two length-based baseline models that use the number of words in a title (or article) as the only feature for classification.

Content (title/article). The words and phrases used in a text can be potential indicators of its target group. Thus, we make use of the most common⁶

⁶Note that we could have used all n-grams, but due to the small size of our data (see §6.2), we decided to limit the number of features via an additional hyperparameter.

	W	M	K	T	Total
TRAIN	805	172	416	337	1,730
DEV	94	23	45	37	199
TEST	94	14	38	37	183
Total	1,202		910		2,112

Table 7: Number of articles for each target group and data split, as well as for each task in total.

uni-grams and bi-grams, excluding stop words, as a feature representation for the content of a how-to guide. We evaluate two variants: features derived from the articles and from the titles.

Style (article). We represent style using two sets of established features from authorship attribution (Sari et al., 2018), namely *lexical* style: average word length, number of short words, vocabulary richness in terms of hapax-legomena and dislegomena, % of digits, % of upper case letters; and *syntactical* style: occurrences of punctuation, frequencies of POS tags, and stop-word frequencies.

combined (article). Content and style can potentially provide complementary information. We test whether a model can leverage a combination of information from different sources. For this purpose, we simply concatenate the article-level features for content, style, and length.

RoBERTa (article). As an alternative to manually selected features, we further test features derived from a large language model, RoBERTa (Liu et al., 2019). Specifically, we encode the article’s text, truncated to the first 512 tokens, and extract the representation of the special classification token from the last hidden layer as a set of feature values.

6.2 Experimental Setup

In order to find out whether and to what extent articles for different target groups can be distinguished computationally, we define two classification tasks in which specific articles, based on their characteristics, are to be assigned to one target group each. We distinguish between articles for women and men (W–M) and between articles for kids and teenagers (K–T). For all four classes, we use the full wikiHowAudiences, which we divide into TRAIN, DEV, and TEST sets following the article-level partition of the original wikiHowToImprove corpus (Anthonio et al., 2020). Statistics for each

Model	W-M	K-T
Baselines		
Majority baseline	0.47	0.34
Length (title)	0.47	0.44
Length (article)	0.47	0.61
Content & Style		
Content (title)	0.57	0.57
Content (article)	0.59	0.78
Style (article)	0.58	0.67
“Full” models		
combined (article)	0.71	0.78
RoBERTa (article)	0.68	0.74

Table 8: Macro F_1 -scores on the test sets.

class and set are shown in Table 7. For the style features, the texts are lemmatized with spaCy.⁷

We train each model on the TRAIN set and evaluate in terms of macro F_1 -score on the TEST set. We compute F_1 -score per class as the harmonic mean between precision (ratio of correct predictions) and recall (ratio of correctly classified instances). As our data is imbalanced, we use macro F_1 instead of a weighted/micro score to treat each class (rather than each instance) as equally important.

A number of hyperparameters are optimized on the DEV set: We try different values for the logistic regression classifiers’ L1 and C terms, sampled from 10 instances between $1e - 5$ and 100. For the content features, we optimize the number of k most common n-grams ($k = 200$). We also made use of the DEV set to determine the best language model for our tasks, which we found to be roberta-large (results of other models are shown in Appendix A.2).⁸

6.3 Results

The results are summarized in Table 8. As conjectured based on the K-T articles from the first case study, we find that the length-based baselines indeed outperform the majority baseline⁹ in that setting. As the further results show, content and stylistic features can indeed be used to correctly assign a specified target group to many how-to guides. According to the evaluation scores, features calculated

⁷<https://spacy.io/>

⁸We used HuggingFace Transformers (Wolf et al., 2020).

⁹Note that the F_1 -score for the majority baseline lies below 0.5 because we calculate the *macro average* over both classes and the score does not reach 1.0 for either class.

at the article level are particularly suitable for this purpose: The combined model, which uses content, style and length features on the article level, achieves the best result with macro- F_1 scores of 0.71 and 0.78 for W-M and K-T, respectively. Features generated based on the roberta-large language model achieve competitive scores (0.68 and 0.74), but fall short of the combined model.

The large differences in result between the baselines and our models show that the target audience of many articles can be determined simply from the vocabulary and style of an article. Next, we take a closer look at model features and errors.

6.4 Analyses

For our analyses, we focus on the combined model because it achieves the best results and its features are easily interpretable.

Features. For each target group, we analyze what features are most important to the model. Since our model uses independent features in a binary classification task, we can simply check the highest positive and negative feature weights for this purpose. A selection from the ten most predictive features¹⁰ and example sentences are shown in Table 9. As the examples illustrate, some of the strongest features are, again, based on stereotypes (e.g., ‘cute’, ‘makeup’ for W) or reflect heteronormative assumptions (‘hers’ for M). Interestingly, we also see characteristics of gender-inclusive language (‘theirs’ for M) and direct address of the reader in terms of their group membership (‘kid’ for K and ‘teen’ for T). We further find negations (e.g., ‘wasn’t’) as part of strong features for W, which is particularly worrying in light of sociopsychological findings that have shown negations to serve a stereotype-maintaining function across languages (Beukeboom et al., 2010, 2020).

Same title articles. As examples of particularly hard cases, we return to the how-to guides from the first case study, which consisted of article versions for different audiences (§4). Following the data partition from previous work, we identify 16 such articles in the DEV and TEST splits. We find that the combined model classifies 12 of them correctly (75%). In the remaining 4 cases, the prediction errors could have been caused by superficial features that are predictive for the opposite audience. We note for each of these 16 articles that the version

¹⁰Appendix A.2 lists all top-10 most predictive features.

	Feature(s)	Example	Title
W	cute, makeup wasn't	Do <u>cute</u> <u>makeup</u> . She most likely <u>wasn't</u> wearing the right colors for her skin <u>tone</u> .	Look Cute Go from Ugly to Popular
M	hers theirs	Slowly move your hand towards <u>hers</u> ... Being a good partner is all about ... adjusting your style to suit <u>theirs</u> .	Know if Your Crush Likes You Back Grind
K	name kid	Think of your blog's <u>name</u> even if you're a <u>kid</u> , there are ways to bank a few extra bucks.	Write a Blog Make Money
T	dress teen	<u>Dress</u> up, make it look important. <u>When</u> you're a <u>teen</u> with a busy schedule, it can be difficult to find time to be active.	Know What to Wear on Dates Stay Active After School

Table 9: Sample of the top-10 most predictive features and example sentences from articles of each target group.

for the opposite audience is part of the TRAIN split. Therefore, the topics of the guides are generally not specific to one audience, and a correct classification of the majority of cases demonstrates that the model indeed captures characteristics of content and style that seem specific to the audience itself.

7 Discussion and Conclusion

In this paper, we assessed differences across how-to guides written for specific audiences. In the construction of sub-corpora for four target groups, we already noticed inequalities on the level of who is being instructed in wikiHow: as a target audience, women are mentioned more than four times more frequently than men, and teens receive about 50% more instructions per article than kids. In two case studies, we investigated and provided examples of target-related differences on the levels of topic, style, and content.

The differences observed in our case studies inspired feature sets of shallow classifiers for predicting the target audience of a given guide. Using these classifiers, we showed that it is, in many cases, indeed possible to automatically predict for which audience an article was written. In an analysis of our results, we found that this success is not merely based on different topics covered for each target group but that the articles for each group systematically differ in terms of content and style.

Each of the aforementioned observations presents a tiny, seemingly insignificant piece of a puzzle. But taken together, these pieces reveal a surprisingly clear picture: there are noticeable

differences in what topics are covered for each target group, how many articles and instructions are provided for each audience, and how these articles are written. Even though the audience-specific characteristics used in our studies are by no means exhaustive, our straightforward approach allowed us to identify, qualitatively and quantitatively, debatable differences in how wikiHow guides present particular topics to specific target groups. While there is an inevitable need for differences in vocabulary when speaking about physical features or body parts, it is at best unclear in which ways how-to guides about human interactions or self-presentation should cast significant differences.

Some of the observed differences have already been critically discussed in the context of social science research. For example, it is well-known that labels such as ‘cute’ are used pejoratively as a form of social control (Talbot, 2019) and that prescriptive components of gender stereotypes in education contribute to discrimination (Kollmayer et al., 2018). However, exposing readers to cultural messages and beliefs about age, gender or other factors cannot be avoided entirely, especially on a collaboratively edited online platform. In fact, it seems to be a challenge for any pluralistic society to find a balance between communicating traditional values and empowering everyone. It is therefore all the more important for a comprehensive understanding to determine when and in what form social norms are conveyed. As such, we view the contributions of this paper, namely our data set of audience-specific guides, wikiHowAudiences,

and our mixed-methods approach for identifying and verifying differences, as a valuable connecting point to raise awareness of potential issues and to foster interdisciplinary dialogue for future research.

Limitations

Our studies focus on the differences in how-to guides written for specific audiences only in one language, namely English. A major limitation is therefore that we do not consider other languages.

The perspectives provided by the data source we rely on, wikiHow, allow us to identify specific phenomena and peculiarities. Yet, contemplating only one data source lets us generalize only to a limited extent. For example, the audiences considered in this work depended on the target groups portrayed in the data. They are neither exhaustive nor representative of the diversity of humankind, especially of marginalized social groups. Therefore, a wider variety of data sources will be needed to test generalizations.

Finally, a further limitation of our studies concerns intersectionality. While it seems possible that guides can be tuned by contemplating one specific attribute of the audience at a time, this does not hold with regard to the actual attributes of the readers. Such attributes are per se coexistent, and consequently, they are not separable.

Ethics Statement

We acknowledge that the content that emerged from the data is narrow in terms of cultural perspectives, mainly addressing western cultures. Moreover, the analysis of the audiences is not exhaustive of the diversity of humankind, especially not exhaustively accounting for queer identities in particular trans and non-binary identities. With the present research, we do not intend to reinforce representational biases, rather to highlight them.

References

Sweta Agrawal and Marine Carpuat. 2019. [Controlling text complexity in neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1549–1564, Hong Kong, China. Association for Computational Linguistics.

Talita Anthonio, Irshad Bhat, and Michael Roth. 2020. [wikiHowToImprove: A resource and analyses on](#)

[edits in instructional texts](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5721–5729, Marseille, France. European Language Resources Association.

Talita Anthonio and Michael Roth. 2020. [What can we learn from noun substitutions in revision histories?](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1359–1370, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Talita Anthonio and Michael Roth. 2021. [Resolving implicit references in instructional texts](#). In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 58–71, Punta Cana, Dominican Republic and Online. Association for Computational Linguistics.

Talita Anthonio, Anna Sauer, and Michael Roth. 2022. [Clarifying implicit and underspecified phrases in instructional text](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3319–3330, Marseille, France. European Language Resources Association.

Camiel J. Beukeboom, Christian Burgers, Zsolt P. Szabó, Slavica Cvejic, Jan-Erik M. Lönnqvist, and Kasper Welbers. 2020. [The negation bias in stereotype maintenance: A replication in five languages](#). *Journal of Language and Social Psychology*, 39(2):219–236.

Camiel J. Beukeboom, Catrin Finkenauer, and Daniël H. J. Wigboldus. 2010. [The negation bias: When negations signal stereotypic expectancies](#). *Journal of Personality and Social Psychology*, 99(6):978–992.

Irshad Bhat, Talita Anthonio, and Michael Roth. 2020. [Towards modeling revision requirements in wikiHow instructions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8407–8414, Online. Association for Computational Linguistics.

S.R.K. Branavan, Harr Chen, Luke Zettlemoyer, and Regina Barzilay. 2009. [Reinforcement learning for mapping instructions to actions](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 82–90, Suntec, Singapore. Association for Computational Linguistics.

Pengshan Cai, Mo Yu, Fei Liu, and Hong Yu. 2022. [Generating coherent narratives with subtopic planning to answer how-to questions](#). In *Proceedings of the 2nd Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 26–42, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ewa S. Callahan and Susan C. Herring. 2011. [Cultural Bias in Wikipedia Content on Famous Persons](#). *Journal of the American society for information science and technology*, 62(10):1899–1915.

- Carmen Cervone, Martha Augoustinos, and Anne Maass. 2021. [The language of derogation and hate: Functions, consequences, and reappropriation](#). *Journal of language and social psychology*, 40(1):80–101.
- Youngjoo Cha and Kim A Weeden. 2014. [Overwork and the slow convergence in the gender gap in wages](#). *American Sociological Review*, 79(3):457–484.
- Cuong Xuan Chu, Niket Tandon, and Gerhard Weikum. 2017. [Distilling task knowledge from how-to communities](#). In *Proceedings of the 26th International Conference on World Wide Web*, pages 805–814.
- Lilia M Cortina, Dana Kabat-Farr, Emily A Leskinen, Marisela Huerta, and Vicki J Magley. 2013. [Selective incivility as modern discrimination in organizations: Evidence and impact](#). *Journal of management*, 39(6):1579–1605.
- Alok Debnath and Michael Roth. 2021. [A computational analysis of vagueness in revisions of instructional texts](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 30–35, Online. Association for Computational Linguistics.
- Elizabeth N. Dougherty, Andrea B. Goldschmidt, Nicole K. Johnson, Krystal Badillo, Scott G. Engel, and Alissa A. Haedt-Matt. 2022. [Gender differences in the relation between interpersonal stress and momentary shape and weight concerns in youth with overweight/obesity](#). *Body Image*, 40:249–255.
- Noemie Elhadad and Komal Sutaria. 2007. [Mining a lexicon of technical terms and lay equivalents](#). In *Biological, translational, and clinical language processing*, pages 49–56, Prague, Czech Republic. Association for Computational Linguistics.
- Agnieszka Falenska and Özlem Çetinoğlu. 2021. [Assessing gender bias in Wikipedia: Inequalities in article titles](#). In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 75–85, Online. Association for Computational Linguistics.
- Danilo Fum, Giovanni Guida, and Carlo Tasso. 1986. [Tailoring importance evaluation to reader’s goals: A contribution to descriptive text summarization](#). In *Coling 1986 Volume 1: The 11th International Conference on Computational Linguistics*.
- Xingwei He and Siu Ming Yiu. 2022. [Controllable dictionary example generation: Generating example sentences for specific targeted audiences](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 610–627, Dublin, Ireland. Association for Computational Linguistics.
- Tomoyuki Kajiwara, Hiroshi Matsumoto, and Kazuhide Yamamoto. 2013. [Selecting proper lexical paraphrase for children](#). In *Proceedings of the 25th Conference on Computational Linguistics and Speech Processing (ROCLING 2013)*, pages 59–73, Kaohsiung, Taiwan. The Association for Computational Linguistics and Chinese Language Processing (ACLCLP).
- Marlene Kollmayer, Barbara Schober, and Christiane Spiel. 2018. [Gender stereotypes in education: Development, consequences, and interventions](#). *European Journal of Developmental Psychology*, 15(4):361–377.
- Piotr Konieczny and Maximilian Klein. 2018. [Gender gap through time and space: A journey through Wikipedia biographies via the Wikidata Human Gender Indicator](#). *New Media Soc.*, 20(12).
- Mahnaz Koupaee and William Yang Wang. 2018. [Wikihow: A large scale text summarization dataset](#). *CoRR*, abs/1810.09305.
- Shyong (Tony) K. Lam, Anuradha Uduwage, Zhenhua Dong, Shilad Sen, David R. Musicant, Loren Terveen, and John Riedl. 2011. [WP:Clubhouse? an exploration of Wikipedia’s gender imbalance](#). In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration, WikiSym ’11*, page 1–10, New York, NY, USA. Association for Computing Machinery.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach](#).
- Laia Palència, Davide Malmusi, Deborah De Moortel, Lucía Artazcoz, Mona Backhans, Christophe Vanroelen, and Carme Borrell. 2014. [The influence of gender equality policies on gender inequalities in health in europe](#). *Social science & medicine*, 117:25–33.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Deborah A Prentice and Erica Carranza. 2002. [What women and men should be, shouldn’t be, are allowed to be, and don’t have to be: The contents of prescriptive gender stereotypes](#). *Psychology of women quarterly*, 26(4):269–281.
- Joseph Reagle and Lauren Rhue. 2011. [Gender bias in Wikipedia and Britannica](#). *International Journal of Communication*, 5:21.
- Michael Roth and Talita Anthonio. 2021. [UnImplicit shared task report: Detecting clarification requirements in instructional text](#). In *Proceedings of the 1st Workshop on Understanding Implicit and Underspecified Language*, pages 28–32, Online. Association for Computational Linguistics.

- Yunita Sari, Mark Stevenson, and Andreas Vlachos. 2018. [Topic or style? exploring the most useful features for authorship attribution](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 343–353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Carolina Scarton and Lucia Specia. 2018. [Learning simplifications for specific target audiences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 712–718, Melbourne, Australia. Association for Computational Linguistics.
- Katja Geertruida Schmahl, Tom Julian Viering, Stavros Makrodimitis, Arman Naseri Jahfari, and Marco Tax, David andj Loog. 2020. [Is Wikipedia succeeding in reducing gender bias? assessing changes in gender bias in Wikipedia using word embeddings](#). In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 94–103, Online. Association for Computational Linguistics.
- Yasuko Senda and Yaushi Shinohara. 2002. [Analysis of titles and readers for title generation centered on the readers](#). In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Yasuko Senda, Yasusi Sinohara, and Manabu Okumura. 2004. [A support system for revising titles to stimulate the lay reader’s interest in technical achievements](#). In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 155–161, Geneva, Switzerland. COLING.
- Yasuko Senda, Yasusi Sinohara, and Manabu Okumura. 2006. [Automatic terminology intelligibility estimation for readership-oriented technical writing](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Kim Cheng Sheang and Horacio Saggion. 2021. [Controllable sentence simplification with a unified text-to-text transfer transformer](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 341–352, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Advaith Siddharthan and Napoleon Katsos. 2010. [Reformulating discourse connectives for non-expert readers](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1002–1010, Los Angeles, California. Association for Computational Linguistics.
- Janet K Swim, Robyn Mallett, and Charles Stangor. 2004. [Understanding subtle sexism: Detection and use of sexist language](#). *Sex roles*, 51(3):117–128.
- Mary Talbot. 2019. *Language and gender*. John Wiley & Sons.
- Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. 2015. [It’s a man’s Wikipedia? assessing gender inequality in an online encyclopedia](#). In *Proceedings of the international AAAI conference on web and social media*, volume 9, pages 454–463.
- Claudia Wagner, Eduardo Graells-Garrido, and David García. 2016. [Women through the glass ceiling: gender asymmetries in Wikipedia](#). *EPJ Data Science*, 5:1–24.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Li Zhang, Qing Lyu, and Chris Callison-Burch. 2020. [Reasoning about goals, steps, and temporal ordering with WikiHow](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4630–4639, Online. Association for Computational Linguistics.

A Appendix

A.1 Case Study

Category	Word Overlap	Same Title	Indicator for W	Indicator for M
BODY	0.02	Get Clear Skin	for Middle School Girls	Guys
BODY	0.05	Burn Fat	for Girls	for Men
PRESENT	0.07	Get Ready for School	for Girls	Guys
PRESENT	0.09	Look Rich Without Being Rich	Teen Girls	for Guys
PRESENT	0.16	Get Ready for School	Teen Girls	Guys
INTERACT	0.16	Catch Your Crush's Eye	for Girls Only	Boys
INTERACT	0.16	Dance at a School Dance	for Girls	for Guys
PRESENT	0.21	Act Like a Kid Again	Girls	Boys
PRESENT	0.21	Look Like an Abercrombie Model	for Girls	Boys
PRESENT	0.21	Dress Emo	for Girls	Guys
BODY	0.24	Have Good Hygiene	Girls	Boys
PRESENT	0.25	Prepare for a School Dance	for Girls	for Guys
PRESENT	0.27	Pack for Soccer Practice	Girls	Boys
INTERACT	0.28	Be in a Female Led Relationship	Women	Men
INTERACT	0.30	Act on a Date	for Girls	for Boys
PRESENT	0.31	Dress Cool	for Girls	Guys
BODY	0.31	Lose Belly Fat	Teen Girls	for Men
PRESENT	0.33	Look Hot on Club Penguin	Girls	Guys
PRESENT	0.35	Be Awesome	for Girls	for Boys
INTERACT	0.36	Have Fun with Your Friends	Teen Girls	Guys
PRESENT	0.36	Dress Like a CEO	Women	Men
INTERACT	0.37	Cradle a Lacrosse Stick	Girls	Men
INTERACT	0.41	Get Your Crush to Like You	Girls	Guys
INTERACT	0.43	Practice Changing Room Etiquette	Girls	Men
INTERACT	0.43	Practice Changing Room Etiquette	Women	Men
BODY	0.44	Recognize Trichomoniasis Symptoms	Women	Men
PRESENT	0.45	Be Popular in Middle School	for Girls	for Boys
BODY	0.47	Lose Belly Fat	for Women	for Men
BODY	0.49	Gain Weight Fast	for Women	for Men
BODY	0.52	Be Indie	for Girls	for Guys
INTERACT	0.54	Grind	for Girls	for Guys
INTERACT	0.55	Host a Sleepover	Teen Girls	for Boys
BODY	0.57	Treat Acne	Teenage Girls	Teen Boys
BODY	0.60	Prevent HIV Infection	Women	Men
BODY	0.69	Recognize Chlamydia Symptoms	for Women	for Men
			Indicator for K	Indicator for T
ACTIVITY	0.05	Flirt	Middle School	for Teens
ACTIVITY	0.10	Redo Your Bedroom	Preteen Girls	Teen Girls
GROWN-UP	0.11	Look Older	Preteen Girls	Teenage Girls
ACTIVITY	0.14	Enjoy Summer Vacation	for Kids	for Teens
ACTIVITY	0.18	Clean Your Room	Kids	Teens
ADVICE	0.19	Enjoy a Plane Ride	for Grade School Kids	Teen Girls
ADVICE	0.20	Be Less Insecure	Preteens	for Teen Girls
ACTIVITY	0.28	Clean Your Room	Tween Girls	Teens
ACTIVITY	0.29	Pack for a Vacation	Preteen Girls	Teen Girls
ADVICE	0.30	Get a Boy to Like You	Pre Teens	Teens
ACTIVITY	0.31	Apply Makeup	Preteens	for Teen Girls
ADVICE	0.33	Balance School and Life	Middle School	Teens
ACTIVITY	0.36	Host a Girls Only Sleepover	for Preteens	Teens
ACTIVITY	0.39	Get Ready for Bed	Tween Girls	for Teenage Girls
GROWN-UP	0.46	Get Fit	for Kids	Teenage Girls
ACTIVITY	0.53	Apply Makeup	Preteens	for Teens
GROWN-UP	0.59	Make Money	for Kids	for Teenagers

Table 10: All “Same Title, Different Audience” guides.

Be (...)	X	(...)	indicator
Be	Popular	and Athletic	(for Girls)
Be	Popular	in Grade 6.	(for Girls.)
Be	Popular	in Middle School	(for Girls)
Be	Popular	in a School Uniform	(Girls)
Be	Popular	in Secondary School	(for Girls)
Be a	Cute	Teen	(Girl)
Be	Cute		(Tween Girls)
Be the	Cute	and Hot Teen	(Girls)
Be	Cute	at School	(Girls)
Be	Cool	Around Your Crush	(for Boys)
Be	Cool	in High School	(Boys)
Be a	Cool	Christian	(Teen Guys)
Be	More	Attractive to Girls	(for Boys)
Be	More	Physically Attractive	(Men)
Be	More	Socially Open	(Men)
Be a	Good	Hamster Owner	(for Kids)
Be a	Good	Stuffed Animal Mom	(for Kids)
Be	Good	With Money	(for Kids)
Be a	Good	Friend	(Teens)
Be a	Good	Writer	(Teens)

Table 11: The most common completions in the titles for “how to be”.

A.2 Classification tasks

model-name	W-M	K-T
bert-base-uncased	0.57	0.64
roberta-base	0.81	0.73
bert-large-uncased	0.73	0.74
roberta-large	0.82	0.75

Table 12: The performance on the DEV set of the classification tasks with optimized LR using the [CLS] token representations from the different LMs.

Most predictive features of the combined model:

W: hadn’t - wasn’t - cute - makeup - ourselves - bag - skirt - outfit - move - sleep

M: man - product - boy - yourselves - o - dance - theirs - shoe - hers - person

K: kid - the - adult - name - are - step - were - else - probably - mean

T: teen - without - than - dress - next - her - want - buy - everyone - ADJ

A.3 Confusion Matrices

	DEV		TEST	
	W	M	W	M
W	0.83	0.17	0.87	0.13
M	0.48	0.52	0.36	0.64
W	78	16	82	12
M	11	12	5	9

Table 13: The confusion matrix for the dev set (left) and the confusion matrix for the test set (right).

	DEV		TEST	
	K	T	K	T
K	0.78	0.22	0.87	0.13
T	0.35	0.65	0.30	0.70
K	35	10	33	5
T	13	24	11	26

Table 14: The confusion matrix for the dev set (left) and the confusion matrix for the test set (right).