Chapter 1

Linguistic Bias in Crowdsourced Biographies
A Cross-lingual Examination

Jahna Otterbacher*, Ioannis Katakis† and Pantelis Agathangelou‡

Research Centre on Interactive Media, Smart Systems and Emerging Technologies &
Cyprus Center for Algorithmic Transparency
Open University of Cyprus
Nicosia, Cyprus
jahna.otterbacher@ouc.ac.cy

Computer Science Department
University of Nicosia
Nicosia, Cyprus
katakis.i@unic.ac.cy, pandelisagathangelou@gmail.com

Biographies make up a significant portion of Wikipedia entries and are a source of information and inspiration for the public. We examine a threat to their objectivity, *linguistic biases*, which are pervasive in human communication. Linguistic bias, the systematic asymmetry in the language used to describe people as a function of their social groups, plays a role in the perpetuation of stereotypes. Theory predicts that we describe people who are expected – because they are members of our own in-groups or are stereotype-congruent – with more abstract, subjective language, as compared to others. Abstract language has the power to sway our impressions of others as it implies stability over time. Extending our monolingual work, we consider biographies of intellectuals at the English- and Greek-language Wikipedias. We use our recently introduced sentiment analysis tool, DidaxTo, which extracts domain-specific opinion words to build lexicons of subjective words in each language and for each gender, and compare the extent to which abstract language is used. Contrary to expectation, we find evidence of gender-based linguistic bias, with women being described more abstractly as compared

*Open University of Cyprus
†University of Nicosia
‡University of Nicosia

to men. However, this is limited to English-language biographies. We discuss the implications of using DidaxTo to monitor linguistic bias in texts produced via crowdsourcing.

1. Introduction

Wikipedia continues to be one of the world's most popular websites, and is often described as being the largest collaboratively-edited source of free information. In addition to providing a platform for both informal[1] and formal[2] learning amongst citizens, Wikipedia has become a rich data source for researchers. For instance, corpora built from Wikipedia entries are often used by natural language processing[3,4] and machine learning[5] researchers. In addition, Wikipedia is often considered a prime case study for those researching information diffusion[6] and the growth of social networks.[7] Given Wikipedia's influence across a number of domains, it is not surprising that many have raised concerns as to its quality and reliability.

A number of researchers has attempted to develop generalized automated methods for detecting articles of high and low quality. Some of these methods rely on article metadata, including the edit history and profiles of contributing editors[8] or the patterns over time in article activity and overall lifecycles.[9] Others have attempted to exploit the textual characteristics of articles, such as linguistic and stylistic features[10] or even simple wordcounts.[11] Nonetheless, Wikipedia itself does not provide users with any such metrics, maintaining that quality is ensured by its unique collaborative editorial control processes. Indeed, it has been reported that given a critical mass of contributors to a given article, high levels of accuracy, completeness and clarity are reached.[12]

The current work concerns a specific type of Wikipedia entry that is particularly sensitive to issues of information quality – biographies of persons, living or deceased. Pentzold, in characterizing Wikipedia as a place where collective memories are negotiated and archived, notes that Wikipedia biographies detail the public view of a person's character and lifetime accomplishments.[13] Not surprisingly, Wikipedia maintains a page on "Biographies of living persons," in which it emphasizes the sensitive nature of such entries, as well as the need for participants to adhere to its three core policies of conveying a neutral point of view (NPOV), verifiability of all information, and no original research[a]. These guidelines are necessary not only because of the sensitive nature of biographies, but also because they

[a] https://en.wikipedia.org/wiki/Wikipedia:Biographies_of_living_persons

are a very common type of entry at Wikipedia.

In fact, Flekova et al.[14] found that over one-fifth of all Wikipedia articles describes persons. Therefore, they argued for the development of automated methods for ensuring the quality of biographies. To this end, they developed a means for scoring a given article with respect to four dimensions of quality: completeness, writing style, trustworthiness and, closely related to our work, objectivity. In gauging the extent to which a given biography reflects the subjectivity/objectivity of its authors, they used both textual features (e.g., words expressing sentiment) and Wikipedia features (i.e., article metadata). Given the importance of the textual features in their predictive model, their results demonstrated that the manner in which Wikipedia participants describe others not only conveys information about the target persons, but also about the authors as information sources. In contrast to this approach, we focus on a particular phenomenon that might negatively impact the objectivity of Wikipedia biographies, *linguistic biases*, which are known to be not only persistent in human communication, but also very subtle. As will be explained in detail, we shall study how famous scientists and intellectuals, both women and men, are described in their Wikipedia biographies in two language editions, English and Greek. In particular, we shall examine which types of linguistic biases are likely to pose a threat to objectivity in biographies.

1.1. *Linguistic bias and social stereotypes*

Social psychologists have long been convinced that the stylistic features of language play a key role in the transmission and maintainance of social stereotypes.[15,16] In other words, when we are describing others, it is not only what we say about them (i.e., the content of our message) but also how we say what we say, which reveals the stereotypes that influence our perceptions of others. Beukeboom[17] defines the term *linguistic bias* as:

> A systematic asymmetry in the way that one uses language, as a function of the social group of the person(s) being described.

Given Wikipedia's NPOV policy, and its extensive editorial control processes, it is not likely that we would find explicit indicators of bias, such as the use of racial slurs or sexist language being used in biographies. However, we may find that there are systematic asymmetries in the manner in which social groups (e.g., gender- or ethnicity-based) are described. Consider the following three statements:

(1) Thomas Edison invented several devices.
(2) Thomas Edison was an American inventor and businessman.
(3) Thomas Edison was America's greatest inventor.

The first of the three descriptions is the most concrete and objective; it contains no subjective words and details a specific action. In contrast, the second description is a bit more abstract, since it describes an enduring characteristic of the target person. Finally, the third description is the most abstract, as it makes a general statement about Edison using a subjective adjective ("greatest"). The question of interest is whether such subtle differences in biographies are systematic, with respect to three characteristics: (1) the gender of the target person (i.e., men versus women intellectuals), (2) the ethnic background of the target person (i.e., individuals hailing from the English-speaking world versus others), and (3) the language in which the biography is written (i.e., Greek versus English). Theory holds that such systematic differences are likely to reinforce social stereotypes. For instance, if men intellectuals were consistently described more abstractly and positively, as compared to women, this would reinforce the notion that men are expected to be successful and famous, while woman are not.

A growing number of social media platforms allows participants to collaboratively produce biographies. Given their potential influence, both in terms of a source of information for readers as well as a source of data for researchers, it is of growing importance to consider their quality and objectivity. In previous work, we analyzed English-language biographies of actors and actresses produced at the Internet Movie Database (IMDb).[18] Specifically, we considered gender- and race-based linguistic biases, and found that Caucasian men actors were more likely than other social groups to be described in an abstract, positive manner. The current work extends our research in a number of ways. First, we analyze biographies from Wikipedia, which as mentioned, is one of the most popular sites worldwide. In addition, we study biographies not only from the largest of the Wikipedias, English, but also from a smaller community, the Greek-language Wikipedia, in order to examine whether linguistic biases occur across languages, as predicted by theory.[19] To this end, we utilize a recently introduced tool, DidaxTo[20] that extracts opinion words that the authors use in a collection of documents. It operates in both languages: English and Greek.

As we will show, there are interesting cross-lingual differences in terms of the linguistic biases that manifest in biographies of famous intellectuals that are produced collaboratively at Wikipedia. In addition, our results suggest that it is not always that case that theories of linguistic bias, which have been developed by social psychologists in offline, experimental settings, can predict the types of biases observed in online, crowdsourced biographical texts. As will be explained, our findings underscore the need to continue to explore the phenomenon of linguistic bias in social media spaces where social actors (i.e., writers) are often anonymous. In the next section, we provide the theoretical background that underlies our work, before detailing our methodology.

2. Background

Our work is inspired by social psychology theories surrounding two types of linguistic biases: the Linguistic Expectancy Bias (LEB) and the Linguistic Intergroup Bias (LIB). Both are manifested through two stylistic characteristics of the language that a communicator uses to describe someone: the specificity of the description, as well as the use of subjective words (i.e., words that reveal sentiment). In order to provide adequate background on the LEB and the LIB, we must first start with an overview of the Linguistic Category Model (LCM).[21]

2.1. *Linguistic Category Model*

Semin and Fiedler's Linguistic Category Model is a tool for understanding language as a social behavior. More specifically, it proposes a shift in the methodological approach to analyzing language, from the individual to the social perceptive. According to a manual for analysts applying LCM, "to understand social behavior, one has to develop a handle on language as a tool that carries communication and makes social interaction possible,"[22] (page 4).

As shown in Figure 1, the LCM consists of four categories of predicates, which relate to the level of abstraction in a person description. The most concrete category comprises predicates involving a descriptive action verb, which describes an observed event with no additional interpretation of that event. At the other extreme, the most abstract predicates involve an adjective. Here, the respective description is general; it applies across events and scenarios. In between these two categories, we have the use of a state verb or the use of an interpretive action verb. A state verb describes an ongoing state of affairs, and is thus relatively abstract. An interpretive action verb denotes a description that can be attributed only to a specific event or action, and is thus, relatively less abstract.

		Description	Example
More Abstract	**Adjectives**	Describes a characteristic or feature of a person.	Albert Einstein was an amazing mind.
	State verb	Describes an enduring cognitive or emotional state with no clear beginning or end.	Albert Einstein still amazes students today.
	Interpretive action verb	Refers to various actions with clear beginning and end.	Albert Einstein was amazing as a professor at the Swiss Federal Polytechnic.
	Descriptive action verb	Refers to a single, specific action with a clear beginning and end.	Albert Einstein was a professor of theoretical physics at the Swiss Federal Polytechnic.

Figure 1. The Linguistic Category Model (LCM).

The two types of biases, the LEB and the LIB, can be detected based on the extent to which a person description uses abstract language. There are cognitive underpinnings to both, as familiar or expected (i.e., stereotypical) scenes are more easily processed.[23] But while their underpinnings are cognitive, the consequences of these biases are social in nature. As they are known to be pervasive in face-to-face interpersonal communication, it is important to understand the extent to which they are also pervasive in technology-mediated contexts, and especially in crowdsourcing platforms such as Wikipedia.

Abstract descriptions are more powerful than concrete ones. This is because they imply stability over time as well as generalizability. It has been confirmed that message recipients are impacted by the level of abstraction in the language used in person descriptions, with more abstract descriptions being interpreted as enduring characteristics of the target person, in contract to concrete descriptions, which are seen as being transient.[24] In this way, linguistic biases can contribute significantly to the maintenance and transmission of social stereotypes, as information encoded abstractly is more resistant to disconfirmation, as compared to very concrete information.

2.2. *Linguistic Expectancy Bias*

While known to be pervasive in interpersonal communication, the LEB has only been studied in laboratory settings, with few exceptions.[25] The LEB reflects the fact that it is easier for us to process information that is expected (e.g., persons who are stereotype-confirming). We tend to describe other people and situations that are consistent with expectations in a more interpretive and abstract manner. In turn, more abstract descriptions of the target person contain more information about their characteristics and traits, and less about a particular action taken by the person. Laboratory studies have demonstrated that when participants are asked to describe someone who violates their expectations, that they are likely to focus on particular details, providing tangible and concrete information.[19,24] On the other hand, when describing stereotype-congruent (i.e., expectation confirming) individuals, participants are more likely to provide abstract details, using language that references the perceived disposition and traits of the target person.

2.3. *Linguistic Intergroup Bias*

The LIB builds on the LEB; we expect our in-group members to have positive qualities and behaviors, while we may not hold such expectations for out-group members. The LIB predicts that we use language in such a way that it is difficult to disconfirm our preexisting ideas about social groups.[15] Therefore, we are more likely to describe the positive actions and attributes of fellow in-group members with abstract language, whereas any negative traits and actions are more likely described concretely. The converse is predicted for descriptions of out-group members.

2.4. *Detecting linguistic bias in crowdsourced texts*

Having reviewed the key theories of linguistic biases, we shall now propose a method for their detection in the collaboratively-produced Wikipedia biographies. Figure 2 summarizes the linguistic characteristics of our textual biographies, based on the relationship between Wikipedia authors and the target individuals being described.

	Expectancy-congruent (LEB) In-group (LIB)	**Expectancy-incongruent (LEB) Out-group (LIB)**
Familiar/ Desirable Actions and Traits	More abstract	More concrete
	More adjectives More subjective words	Fewer adjectives Fewer subjective words
Unfamiliar/ Undesirable Actions and Traits	More concrete	More abstract
	Fewer adjectives Fewer subjective words	More adjectives More subjective words

Figure 2. Linguistic features predicted by theory.

It is important to note that previous research on linguistic biases has involved manual annotation; in other words, researchers analyze texts with respect to the LCM guidelines. As mentioned in our previous work,[18] while the LCM is a rather precise and complicated model, it is possible to conduct an automated analysis, inspired by key elements of the LCM, such that one can analyze a large corpus of texts. Proponents of the LCM emphasize that the textual segments that one should annotate, and the particular manner in which we apply the LCM depend on the researcher's particular questions. As in previous work, we can note that adjectives play a key role in conveying abstract information about a target person. Furthermore, we can accurately distinguish between adjectives and other parts of speech.[26] Therefore, a person description in which there is a preference for more adjectives, is indicative of a relatively more abstract description.

Secondly, subjective words also play a key role in the construction of more abstract descriptions. Such words inject authors' sentiment into the description, as well as their inferences about the target person. In fact, in psychological studies of social stereotypes, which involved the "trait adjective method," researchers often ask participants to associate subjective adjectives with target social groups.[27]

Given the above observations, we consider four linguistic characteristics of the biographies in our corpus:

(1) The extent to which the text comprises adjectives.
(2) The extent to which the text comprises subjective, positive words.
(3) The extent to which the text comprises subjective, negative words.
(4) The ratio of subjective positive to negative words in a text.

2.5. *Domain-independent sentiment analysis*

For the extraction of subjective words we utilize our recently introduced unsupervised tool for domain-specific opinion word discovery, DidaxTo.[28,29] The novelty of the the tool is that it enables the extraction of subjective terms that are used in each domain independently. Therefore our analysis does not rely on a pre-defined list of opinion words like the one offered in Ref.[30] This is an important feature for several reasons. For instance, some opinion words might be used only in a sub-set of circumstances while others might change meaning and even polarity depending on the respective domain. In order to achieve this goal, DidaxTo utilizes opinion modifiers, sentiment consistency theories, polarity assignment graphs and pattern similarity metrics.

In previous work, DidaxTo was used to learn subjective lexicons in a number of different domains and the resulting lexicons were compared against those obtained via other state-of-the-art approaches. In an explicit evaluation (i.e., where human judges evaluated the accuracy of the learned lexicons), DidaxTo outperformed other sentiment classifiers. Likewise, in the implicit evaluation, where human judgments were available only for the overall sentiment of a text and not individual words, DidaxTo outperformed other methods in the majority of domains tested. Details of the algorithm, as well as the evaluations, can be found in Agathangelou et al.[28]

3. Data and Method

In this section, we first lay out three specific research questions motivated by the theory presented in Section 2, which extends our previous research on linguistic biases in crowdsourced biographies.[18] We then detail the construction of our corpus as well as the processing of the English- and Greek-language biographies.

3.1. *Research questions*

The present study will examine the following three research questions:

- RQ1: Is there evidence of linguistic biases based on the gender of the target person?
- RQ2: Is there evidence of linguistic biases based on the ethnicity of the target person?
- RQ3: Do we observe linguistic biases more frequently in one language as compared to the other?

We shall answer questions one and two within each language edition of Wikipedia, before then comparing the linguistic style used in biographies between the two languages.

3.2. *Corpus*

We built a corpus of biographies of famous intellectuals, including scientists and engineers, inventors and writers. To be included, the target individual must have a biography entry in both the English- and Greek-language Wikipedias. This criterion significantly limited the number of biographies available for inclusion in our corpus; the Greek Wikipedia is a small resource with just over 133,000 entries at the time of writing, compared to the nearly 5.5 million entries at the English-language site[b]. In the end, we have collected and processed 197 biographies of men and 187 biographies of women, in both the Greek and English languages. In other words, the corpus consists of 768 carefully chosen biographies.

Next, biographies were coded for ethnicity of the target persons. To have an objective means to do this, we used citizenship. We distinguished individuals whose Wikipedia biography indicates that they are/were a citizen of an anglophone country (the United States, Canada, England, Ireland, Australia and New Zealand citizens appeared in our corpus) from those who were/are a citizen of other countries. It should be noted that in classifying immigrants and dual citizens (e.g., Albert Einstein, Zaha Hadid) we based the classification on the country in which they died or currently reside. The corpus contains biographies of 154 anglophone intellectuals (61 men, 93 women) and 230 intellectuals from other regions (136 men, 94 women). Given the small size of the Greek Wikipedia, it was not feasible to create groups based on more specific ethnic background criteria.

Table 1 provides summary statistics concerning the length of biographies (number of words). Since the distributions are skewed to the right, we used non-parametric tests to compare lengths across languages and gender.

[b]`https://en.wikipedia.org/wiki/List_of_Wikipedias`

Table 1. Median/mean length of biographies.

	Men	Women
English	4,798/3,070	3,557/2,515
Greek	1,313/601	900/543

The non-parametric Wilcoxon Signed-Ranked Test[31] (an alternative to the paired t-test) revealed that English-language biographies tend to be longer than Greek-language biographies of both women ($V = 17080$, $p < .001$) and men ($V = 19124$, $p < .001$). The two-group Mann-Whitney U test was used to compare across genders within a given language. Results indicated that at the English-language Wikipedia, men's biographies tend to be longer than those of women ($W = 20592$, $p < .05$). However, there was no significant gender-based difference for the Greek-language biographies.

Below, we provide two example biographies of famous men, which appear in our corpus. Below each Greek-language entry, we have provided an English translation. In particular, we provide the opening sentence for each biography, which sets the tone for the text and is likely read by anyone visiting the entry at Wikipedia. Likewise, the first few sentences of a biography are important as they appear in the snippet provided by a search engine in response to a query on the respective person's name. As can be seen, there are some subtle differences between these opening sentences. In the García Márquez biographies, while slightly different information is detailed in the English versus Greek entries, both contain a bit of subjectivity, with the use of the words "affectionately" in the English entry and "important" in the Greek entry. In contrast, in the opening sentences describing Thomas Edison, the Greek entry is notably more objective, as compared to English entry, which refers to Edison as the country's "greatest inventor."

Gabriel José de la Concordia García Márquez
English
Gabriel José de la Concordia García Márquez (6 March 1927 – 17 April 2014) was a Colombian novelist, short-story writer, screenwriter and journalist, known affectionately as Gabo or Gabito throughout Latin America.

Greek
Ο Γκαμπριέλ Γκαρσία Μάρκες (ισπ. Gabriel José García Márquez, 6 Μαρτίου 1927 – 17 Απριλίου 2014) ήταν σπουδαίος Κολομβιανός συγγραφέας, βραβευμένος με Βραβείο Νόμπελ Λογοτεχνίας.
Gabriel García Márquez (Spanish: Gabriel José García Márquez, 6 March 1927 – 17 April 2014) was an important Colombian author, awarded with the Nobel Prize in literature.

Thomas Alva Edison
English
Thomas Alva Edison (February 11, 1847 – October 18, 1931) was an American inventor and businessman, who has been described as America's greatest inventor.

Greek
Ο Τόμας Έντισον (Thomas Alva Edison, 11 Φεβρουαρίου 1847 – 18 Οκτωβρίου 1931) ήταν Αμερικανός εφευρέτης και επιχειρηματίας.
Thomas Edison (Thomas Alva Edison, 11 February 1847 – 18 October 1931) was an American inventor and businessman.

In addition, we provide examples of the opening sentences of two famous women. Between languages, we can again observe some differences. For both women, the Greek biographies open with more subjective sentences. Austen is described as "popular" and "widely-read", while Yourcenar is a "top literary figure."

Jane Austen
English
Jane Austen (16 December 1775 – 18 July 1817) was an English novelist known primarily for her six major novels, which interpret, critique and comment upon the British landed gentry at the end of the 18th century.

Greek
Η Τζέιν Όστεν (Jane Austen, 16 Δεκεμβρίου 1775 - 18 Ιουλίου 1817) είναι μία από τις πιο δημοφιλείς και πολυδιαβασμένες μυθιστοριογράφους της αγγλικής λογοτεχνίας.
Jane Austen (16 December 1775 – 18 July 1817) is one of the most popular and widely-read novelists of English literature.

Marguerite Yourcenar
English
Marguerite Yourcenar (8 June 1903 – 17 December 1987) was a French novelist and essayist born in Brussels, Belgium, who became a US citizen in 1947.

Greek
Η Μαργκερίτ Γιουρσενάρ (γαλλ. Marguerite Yourcenar) (8 Ιουνίου 1903 – 17 Δεκεμβρίου 1987) ήταν Γαλλίδα συγγραφέας και ποιήτρια, μια από τις κορυφαίες λογοτεχνικές μορφές της Γαλλίας του εικοστού αιώνα.
Marguerite Yourcenar (8 June 1903 – 17 December 1987) was a French author and poet, one of France's top literary figures of the twentieth century.

3.3. *Text processing*

DidaxTo was used to learn a lexicon of subjective words of each polarity (negative and positive) for each domain (i.e., by language in which the biography was written and gender of the target individual). Using the extracted lexicon, we obtained the number of positive and negative domain words used in each textual biography, as well as the number of adjectives used. For the POS tagging process we used the implementation of the Stanford parser included in the NLTK Python library[c].

[c]http://www.nltk.org/

3.3.1. *Learned dictionaries*

Table 2 details the sizes of the eight dictionaries that were learned using DidaxTo. One can immediately note a striking difference across genders in the sizes of the dictionaries learned from the English-language entries. In other words, Wikipedians appear to have a much larger vocabulary of subjective words for describing women versus men. It remains to be seen if these words are used at a greater frequency over all (i.e., if, in general, subjective words are used more often in biographies of women versus men), or if more unique words are used in women's biographies (i.e., while the dictionaries of subjective positive and negative words are larger, the words are not necessarily used more often in the entries). It can also be seen that across languages but within a given gender, the dictionary sizes differ. This can be partially explained by the fact that, as examined in Table 1, Greek biographies are significantly shorter than the respective English versions. However, it again remains to be seen from the analysis if the frequency of use of subjective words exhibits a systematic difference between languages.

Table 2. Sizes of dictionaries.

	Positive		Negative	
	Men	Women	Men	Women
GR	553	424	260	185
EN	887	2,175	979	2,086

In order to explore how the learned sentiment dictionaries vary by the target persons' gender, we first translated the Greek-language dictionary entries to English, so that we could make comparisons. Specifically, the extracted Greek words were submitted to Google Translate[d] and then manually corrected where necessary by a native speaker. Next, the words in all four dictionaries were stemmed via Porter's stemming algorithm,[32] such that we could find all unique lemmas.

For each gender and polarity (i.e., negative and positive), we found the intersection of the Greek- and English-language dictionaries. We then identified which words were uniquely used to describe one gender but not the other. Table 3 summarizes this analysis and provides example words that are used to describe persons of each gender. For instance, there are 111 negative words that are used to describe men at both the Greek and English Wikipedias, and 52 of these are used to describe men and not women.

[d]`https://translate.google.com/`

A qualitative observation one can make from Table 3 is that the dictionaries learned by DidaxTo largely reflect prevalent gender-based stereotypes. Researchers of person perception have found that there are two universal dimensions upon which we judge other people - how *warm* (i.e., non-threatening) someone is and how *competent* or *agentic* she is.[33] Traditional gender stereotypes include expectations that men are (and should be) high in competence / low in warmth, and vice versa for women.[34] Similarly, in the dictionaries, we find that many positive sentiment-bearing words used to describe women refer to warmth (e.g., affection, cheer, family, nice) while those used to describe men are more often related to competence (e.g., best, glory, inspire, rich). While a critique of the social and ethical implications of the dictionaries is beyond the scope of the current work, we can conclude that DidaxTo's results are quite logical, given the documented underlying gender biases in Wikipedia biographical texts[35,36] as well as the nature of prevalent gender stereotypes.

Table 3. Subjective words by polarity and associated gender.

		Common words in both GR/EN	Unique for gender	Example words unique for gender
Negative	M	111	52	awkward, barbarian, careless, cheap, foolish, poorly, stuck, stupid, weak, wrong
Negative	W	107	48	fail, greed, guilt, miser, nightmare, sad, shock spoil, weird, wreck
Positive	M	212	86	affirm, best, charm, competitive, convince, fair, glory, inspire, passion, rich
Positive	W	222	96	affection, cheer, clever, creation, colleague, family, friend, host, nice, pioneer sexual, stylish, young

4. Analysis

We analyze the extent to which the target person's gender and ethnicity, as well as the interaction between them, explain the variance in four independent variables, all of which correlate to an increased level of abstraction in the language used, per the Linguistic Category Model. In particular, we examine the variables described in Table 4

Table 4. Linguistic characteristics analyzed.

Characteristic	Continuous variable	Binary variable
Adjectives	$\frac{Adjectives}{Total\_words}$	$\geq$ one adjective used
Positive words	$\frac{Positive\_words}{Total\_words}$	$\geq$ one positive word used
Negative words	$\frac{Negative\_words}{Total\_words}$	$\geq$ one negative word used
Ratio of positive-to-negative	$\frac{Positive\_words+1}{Negative\_words+1}$	Ratio ≥ 1

We examine the four variables in each of the three textual segments of interest: the full-text biography, the first five sentences (i.e., textual "snippet") and the opening or first sentence of the biography. For the full-text biographies and snippets we examine the continuous variables. In contrast, for opening sentences, we examine the binary variables, as the length of opening sentences can often be brief. As we used Analysis of Variance (ANOVA) to analyze the continuous variables, we transformed those that did not meet the normality assumption as described in Table 5. Entries of "N/A" in the table indicate cases where the variables met the normality assumption.

Table 5. Variable transformations used.

Characteristic	Full-text	Snippet
Adjectives	GR: N/A	N/A
	EN: N/A	N/A
Positive words	GR: N/A	sqrt
	EN: N/A	N/A
Negative words	GR: log	GR: log
	EN: N/A	EN: log
Ratio of positive-to-negative	GR: log	log
	EN: log	log

For each of the four independent variables, we fit models to examine the extent to which the target person's gender and ethnicity, and their interaction, explain a significant portion of variance in the variable, and in each case (full-text biography, snippet of five sentences, opening sentence of the biography). In other words, all models test main effects for gender and ethnicity, and an interaction term. In the case of continuous variables, we fit a two-way ANOVA. In the event of statistically significant effects, we report effect sizes using partial η^2 and use Cohen's conventions for interpreting their magnitude.[37]

Briefly, partial η^2 aids in the interpretation of effect sizes between studies. It expresses the ratio of the sum of squares of the effect in question (e.g., gender) to the sum of squares of the effect and the sum of squares of the error associated with the effect. An η^2 ranging from 0.01 to 0.05 is interpreted as indicating a small effect size, while an η^2 ranging from 0.06 to 0.13 indicates a medium effect size. In addition, we follow up with Tukey's Honestly Significant Differences (HSD) test,[38] which compares all possible pairs of means, in order to determine which pairwise differences are meaningful.

For the binary variables, we fit a logit regression model in which we predict the likelihood of a text exhibiting the respective linguistic characteristic. In the event of statistically significant effects, we report the odds ratios as a measure of effect size following Ref.[39] As the odds ratio indicates how many times bigger the odds of one outcome is for a given value of the independent variable as compared to the other (e.g., for women intellectuals versus men, or for anglophone intellectuals versus others), it can be interpreted as an unstandardized effect size.

4.1. *Full-text biographies*

Table 6 details the ANOVA results with respect to the use of adjectives in the biographies, in each language. The right-most column provides the pairwise differences that are significant, per the HSD test, where the p-value is less than 0.05. As can be seen, in both the Greek- and English-language biographies, there is a tendency for Wikipedians to use more adjectives when describing men as compared to women (i.e., there is a significant main effect on gender, albeit with a small effect size, in both models). In addition, target persons from non-anglophone countries tend to be described with significantly more adjectives as compared to citizens of anglophone countries (i.e., main effect on ethnicity).

The ANOVAs on the proportion of words that are positive and negative in full-text biographies of both languages are detailed in Tables 7 and 8, respectively. In addition, Table 9 details the ANOVA on the ratio of positive to negative words. From Tables 7 and 8, we observe that in Greek biographies, positive words are used more often when describing men as compared to women although negative words are used just as often to de-

Table 6. ANOVA on proportion of words that are adjectives (full-text): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	13.8***	5.30*	2.49	M>W
	(.03)	(.01)	n.s.	Other>Ang
EN	17.3***	6.50*	0.48	M>W
	(.03)	(.02)	n.s.	Other>Ang

[a] $***p < .001, **p < .01, *p < .05$

scribe both genders; however, there are no ethnicity-based differences. In contrast, in English-language biographies, women are more often described with subjective words of both polarities (positive and negative) as compared to men. In addition, famous persons from anglophone countries are more likely to be described with positive words than are other individuals. When it comes to the ratio of positive to negative words used in biographies, there are no significant effects on either gender or ethnicity, for either language version, in Table 9.

Table 7. ANOVA on proportion of words that are subjective and positive (full-text): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	6.65*	0.000	0.168	M>W
	(.02)	n.s.	n.s.	
EN	41.0***	19.7***	1.31	W>M
	(.07)	(.04)	n.s.	Ang>Other

[a] $***p < .001, **p < .01, *p < .05$

Table 8. ANOVA on proportion of words that are subjective and negative (full-text): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	2.23	0.408	0.539	n.s.
	n.s.	n.s.	n.s.	
EN	10.4**	0.639	3.40	W>M
	(.02)	n.s.	n.s.	

[a] $***p < .001, **p < .01, *p < .05$

Table 9. ANOVA on ratio of positive to negative words (full-text): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	0.526 n.s.	0.728 n.s.	0.094 n.s.	n.s.
EN	0.043 n.s.	1.646 n.s.	3.407 n.s.	n.s.

[a] $***p < .001, **p < .01, *p < .05$

4.2. *First paragraph*

As previously explained, we also analyzed the first five sentences of biographies. This is meant to approximate the textual "snippet" of a Wikipedia entry that appears in search engine results.

As shown in Table 10, in the first five sentences, just as in the full-text English-language texts, there is a tendency for Wikipedians to use more adjectives when describing men versus women. However, there is also a significant interaction between gender and ethnicity, such that women from anglophone regions are described with fewer adjectives as compared to other women.

Table 10. ANOVA on proportion of words that are adjectives (snippet): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	0.192 n.s.	0.759 n.s.	1.443 n.s.	n.s.
EN	4.74* (.01)	.542 n.s.	6.02* (.01)	M>W Other-W>Ang-W

[a] $***p < .001, **p < .01, *p < .05$

Tables 11 and 12 detail the analysis of the use of positive and negative words in the first five sentences (i.e., the "snippet") of a biography, respectively. Table 13 analyzes the ratio of positive to negative words used. As can be seen, the trends across language diverge. In the Greek-language biographies, more abstract language (negative words) is used in biographies of men and individuals from non-anglophone countries. In English-language biographies, more abstract language (positive and negative words) is observed in biographies of women and those from anglophone countries.

Table 11. ANOVA on proportion of words that are subjective and positive (snippet): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	3.218 n.s.	1.372 n.s.	0.015 n.s.	n.s.
EN	17.5*** (.05)	11.1*** (.03)	0.971 n.s.	W>M Ang>Other

[a] $* * * p < .001, * * p < .01, *p < .05$

Table 12. ANOVA on proportion of words that are subjective and negative (snippet): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	5.129* (.01)	4.134* (.01)	0.169 n.s.	M>W Other>Ang
EN	13.04*** (.02)	5.306* (.01)	0.708 n.s.	W>M Ang>Other

[a] $* * * p < .001, * * p < .01, *p < .05$

Table 13. ANOVA on ratio of positive to negative words (snippet): F-statistic[a] and effect size (η^2)

	Gender	Ethnicity	Gender∗Ethnicity	Sig. Diff.
GR	0.031 n.s.	8.552** (.02)	0.026 n.s.	Ang>Other
EN	0.411 n.s.	0.065 n.s.	0.171 n.s.	n.s.

[a] $* * * p < .001, * * p < .01, *p < .05$

4.3. *Opening sentence*

Finally, we analyzed the opening sentences of biographies, as they are arguably the most-read unit of text in a biography. In addition, the first sentence (i.e., the topic sentence) sets the overall tone of the text and are often used as a summary, which indicates the article's content and tone.[40]

In Table 14, it can be seen that in the English-language biographies, there is evidence that women of any ethnicity tend to be described with more abstract language (i.e., more adjectives) as compared to men. This is an interesting finding as it is the opposite of what the LEB would predict.

Table 14. Logit regression to predict the use of $\geq$ one adjective in opening sentences: coefficients[a] and odds ratios

	Intercept	Gender	Ethnicity	Gender∗Ethnicity
GR	2.088∗∗∗	-0.2582	-0.5736	0.7529
	(8.07)	n.s.	n.s.	n.s.
EN	0.7376∗∗∗	1.948∗∗∗	0.3830	-0.9523
	(2.09)	(7.01)	n.s.	n.s.

[a] $***p < .001, **p < .01, *p < .05$

Tables 15 and 16 present the logit regression models for the prediction of there being at least one positive and negative word in the opening sentence of a given biography. Similarly, Table 17 presents the model for predicting the event that, in the first sentence, there are more positive than negative words (i.e., the ratio of positive to negative words is greater than one). We observe no gender- or ethnicity-based differences in the use of abstract language in Greek-language biographies. However, in the English-language texts, women are described with more abstract language (i.e., using words of both positive and negative polarity) as compared to men. It is also interesting to note that women's biographies appear to be more positive than those of men (i.e., there is a significant effect on gender in predicting the ratio of positive to negative words).

Table 15. Logit regression to predict the use of $\geq$ one positive word in opening sentences: coefficients[a] and odds ratios

	Intercept	Gender	Ethnicity	Gender∗Ethnicity
GR	-0.9474∗∗∗	-0.2382	0.07656	0.3175
	(0.39)	n.s.	n.s.	n.s.
EN	0.3264	1.047∗∗∗	0.5444	-1.0237∗
	n.s.	(2.85)	n.s.	(0.36)

[a] $***p < .001, **p < .01, *p < .05$

5. Discussion

We now interpret our results, answering each of the three research questions put forward. As previously explained, two theories, Linguistic Expectancy Bias (LEB) and Linguistic Intergroup Bias (LIB), have been developed in the context of "offline," laboratory experiments. These theories make

Table 16. Logit regression to predict the use of $\geq$ one negative word in opening sentences: coefficients[a] and odds ratios

	Intercept	Gender	Ethnicity	Gender∗Ethnicity
GR	-3.076*** (0.05)	0.3902 n.s.	0.1139 n.s.	-0.5307 n.s.
EN	-1.540*** (0.21)	0.6316* (1.88)	0.9679** (2.63)	-0.5186 n.s.

[a] $***p < .001, **p < .01, *p < .05$

Table 17. Logit regression to predict the ratio of positive to negative words being $\geq$ one in opening sentences: coefficients[a] and odds ratios

	Intercept	Gender	Ethnicity	Gender∗Ethnicity
GR	-1.021*** (0.36)	-0.4187 n.s.	0.0707 n.s.	0.5781 n.s.
EN	0.02941 n.s.	0.9858*** (2.68)	0.4035 n.s.	-0.8208 n.s.

[a] $***p < .001, **p < .01, *p < .05$

particular predictions regarding the characteristics of language that we can expect to find in descriptions of people. Figure 2 details the features that we would expect to find in Wikipedia biographies of famous intellectuals, under the LEB and LIB, as relates to the use of abstract, subjective language. Tables 18 and 19 provide a summary of the observations from our analyses presented in Section 4, to aid interpretation.

Table 18. Summary of observations (Full-text & Snippet).

		Full-text		Snippet	
		GR	EN	GR	EN
ADJ	Gender	M>W	M>W	n.s.	M>W
	Ethnicity	Other>Ang	Other>Ang	n.s.	Other>Ang
Positive	Gender	M>W	W>M	n.s.	W>M
	Ethnicity	n.s.	Ang>Other	n.s.	Ang>Other
Negative	Gender	n.s.	W>M	M>W	W>M
	Ethnicity	n.s.	n.s.	Other>Ang	Ang>Other
Pos/neg	Gender	n.s.	n.s.	n.s.	n.s.
	Ethnicity	n.s.	n.s.	Ang>Other	n.s.

Table 19. Summary of observations (Opening sentence).

		Opening sentence	
		GR	EN
ADJ	Gender	n.s.	W>M
	Ethnicity	n.s.	n.s.
Positive	Gender	n.s.	W>M
	Ethnicity	n.s.	Other-W>Ang-W
Negative	Gender	n.s.	W>M
	Ethnicity	n.s.	Ang>Other
Pos/neg	Gender	n.s.	W>M
	Ethnicity	n.s.	n.s.

5.1. *Gender-based differences*

Our first research question (RQ1), asked whether there are systematic differences in terms of the frequency of markers of abstract language in the biographies of women versus men intellectuals. Both the LEB and the LIB predict that we should observe more abstract language in the biographies of men as compared to women. As mentioned, the prevailing gender stereotypes in Western societies tend to cast women as less agentic or competent as compared to men.[34] Therefore, men can be said to be more stereotype-congruent, and therefore more expected, as famous intellectuals. In addition, it is well known that there is a substantial gender gap at Wikipedia, with men participants greatly outnumbering women.[41,42] While the gender distribution of participants who collaborate in the writing of a given entry certainly varies, the LIB predicts that men will describe other men using more abstract language, and men are the majority in general.

In the Greek-language Wikipedia, we observe several cases where men are indeed described more frequently using markers of abstract, subjective language, as compared to women. Specifically, in full-text biographies, men are described using more adjectives and positive words than women. In snippets, men are described with more negative words than are women. However, in English-language Wikipedia biographies, the results are less consistent. That said, one salient finding concerns the opening sentences of English-language biographies, in which women's biographies contain significantly more markers of abstract language (adjectives, positive and negative words) as compared to men.

Previous studies have revealed that while famous women and men are covered in Wikipedia equally well (i.e., there is little to no gender-based coverage bias),[35] the topics covered in men and women's biographies vary. Specifically, women's biographies are more likely to emphasize her family and social relationships while men's biographies do not.[43] As mentioned in Section 3, our lexicons of subjective words certainly reflect this difference and this explains why, in the English data, women are systematically described more positively and abstractly than are men. Still another possible explanation is that women who make it into Wikipedia are systematically more notable than men[43] and thus, may be more familiar figures to Wikipedians as compared to men in our corpus.

5.2. *Ethnicity-based differences*

RQ2 asked whether there are significant differences in the use of abstract, subjective language, in the biographies of famous intellectuals from anglophone countries versus others. Even though Wikipedia is open to participants worldwide and enjoys an international user base, the fact remains that the project is based in North America and that the English-language version draws the most participation. It is also well known that a good deal of translation from one version to another takes place.[44] Therefore, in light of the LEB, anglophone individuals might be seen as more "expected" at Wikipedia and therefore, more likely to be described in a more abstract manner. Likewise, the LIB would predict that Wikipedians describe intellectuals from their own ethnic in-group in a more positive, abstract manner. Therefore, we might expect that at the English-language site, that intellectuals hailing from anglophone countries would be described in a more abstract manner as compared to those from other regions.

In the Greek-language biographies, there are only three cases where we observe ethnicity-based differences, and these do not reveal a consistent trend, as shown in Table 18. In the English-language texts, while there are more cases where there are significant differences in the use of abstract, subjective language, again the findings are rather inconsistent. It should be noted that since we studied only biographies of people that appeared at both the Greek and English Wikipedias, the individuals described are internationally revered. Perhaps this explains why we do not observe a salient ethnicity-based linguistic bias (i.e., Wikipedians may not differentiate between "ours" and "theirs" when describing individuals that are internationally famous and thus, familiar to all).

It also may be the case that, because our corpus has a large number of individuals who have immigrated to anglophone countries, our classification technique did not capture ethnic in- and out-groups precisely enough. Another reason for a lack of ethnicity-based linguistic bias is that participation at the largest of the Wikipedias, the English-language community, is international. For instance, our own previous research found that Greek Wikipedians very often contribute to the English-language site in addition to (or even instead of) the Greek-language site.[45]

5.3. *Between-language differences*

Our third research question asked whether linguistic biases are more likely to occur in one language over the other. As illustrated in Tables 18 and 19, there were many more significant differences in the use of markers of abstract language discovered in the English-language biographies as compared to the Greek-language texts. In addition, as mentioned above, there were also more consistent differences in the English-language texts.

It is believed that linguistic biases occur regardless of the language in which the target person(s) are described, as some of the key studies that developed the LEB and LIB were conducted in non-English, European languages.[19] However, we are not aware of previous work that makes cross-lingual comparisons of the extent to which linguistic biases occur. As mentioned, the cross-lingual study of linguistic biases is challenged by the need for appropriate sentiment or subjectivity lexicons for each domain/language. Ours is the first study to attempt such a cross-lingual comparison and the results suggest intriguing differences in the types of linguistic biases that occur.

There are reasons to expect varying degrees of linguistic bias across language editions of Wikipedia. For example, in their comparison of six Wikipedia language communities, Wagner and colleagues[35] analyzed gender biases with respect to structural features (i.e., to which other pages a biography links) and the use of gender-related words. They found significant differences between English- and Russian-language biographies, with more gender-based biases in the Russian-language data. They suggested that there may be a correlation between offline measures of gender equality (e.g., the Gender Inequality Index of the World Economic Forum) and inequality at Wikipedia. Similarly, it is likely the case that different content about a given individual is covered at different versions of Wikipedia. For instance, in their bilingual analysis of Polish- and English-language biographies, Callahan and Herring[44] noted a tendency for the English-language edition to highlight a target person's social life and personal preferences as compared to the Polish edition. Therefore, future work on cross-linguistic comparisons of linguistic bias should take into account these factors.

6. Conclusion

Linguistic biases are known to be extremely common in interpersonal communication, and are believed to be an important means by which social stereotypes are maintained in society.[15] Increasingly, scholars are noting the need to examine the presence of linguistic biases in media, including social media. Collaborative knowledge sources such as Wikipedia are of particular interest in this respect, not only because of their popularity with the public but also because of their open, participatory nature.

Our previous monolingual work motivated the need to develop methods to detect linguistic biases in crowdsourced descriptions of people.[26] Specifically, we relied on the Subjectivity Lexicon[46] to assess the extent to which English-language biographies of Hollywood actors used abstract, subjective language in describing various social groups. In order to extend our work to a more general, cross-lingual case of Greek- and English-language biographies of intellectuals from various backgrounds, we used DidaxTo, a novel unsupervised method, to train sentiment lexicons on biographies for each language (Greek and English) and gender.

As previously explained, DidaxTo offers a significant advantage over lexicon-based approaches, in that we can discover words that are used in a particular domain to convey sentiment. For instance, in the Subjectivity Lexicon, words are labeled as being either strongly subjective (i.e., always sentiment-bearing, regardless of context of use) and weakly subjective (i.e., could be sentiment-bearing, depending on the specific context of use). Indeed, as discussed in Section 3, even the extracted lexicons reveal differences in terms of how Wikipedians describe women versus men, and which aspects of famous women and men are seen as positive and negative.

In future work, we plan to continue to explore cross-lingual differences in linguistic biases exhibited in crowdsourced descriptions of people, using DidaxTo, as well as comparing to other approaches. One direction will be the study of other types of documents like historical events (e.g. wars, inventions, empires, important constructions or art) or research papers. In addition, we plan to conduct longitudinal studies of how linguistic biases evolve over time as a function of the pool of participants collaborating on a given article, taking into consideration their demographic and behavioral characteristics exhibited through their digital traces at Wikipedia. In conclusion, this work will help to ensure that Wikipedia serves as an objective knowledge source for the public. In addition, it will also allow us to learn more about the linguistic behaviors that correlate to the maintenance of social stereotypes, by enabling new methodological approaches beyond laboratory experiments.

Acknowledgements

J. Otterbacher is supported by the European Union's Horizon 2020 Research and Innovation Programme under grant agreements No. 739578 and No. 810105.

References

1. S. Downes, New technology supporting informal learning, *Journal of Emerging Technologies in Web Intelligence*. **2**(1), 27–33 (2010).
2. A. Forte and A. Bruckman. From wikipedia to the classroom: Exploring online publication and learning. In *Proceedings of the 7th international conference on Learning sciences*, pp. 182–188 (2006).
3. M. Strube and S. P. Ponzetto. Wikirelate! computing semantic relatedness using wikipedia. In *AAAI*, vol. 6, pp. 1419–1424 (2006).
4. G. Giannakopoulos, M. El-Haj, B. Favre, M. Litvak, J. Steinberger, and V. Varma, Tac 2011 multiling pilot overview (2011).
5. H.-F. Yu, P. Jain, P. Kar, and I. Dhillon. Large-scale multi-label learning with missing labels. In *International Conference on Machine Learning*, pp. 593–601 (2014).
6. M. Kimura, K. Saito, and R. Nakano. Extracting influential nodes for information diffusion on a social network. In *AAAI*, vol. 7, pp. 1371–1376 (2007).
7. A. Capocci, V. D. Servedio, F. Colaiori, L. S. Buriol, D. Donato, S. Leonardi, and G. Caldarelli, Preferential attachment in the growth of social networks: The internet encyclopedia wikipedia, *Physical review E*. **74**(3), 036116 (2006).
8. M. Hu, E.-P. Lim, A. Sun, H. W. Lauw, and B.-Q. Vuong. Measuring article quality in wikipedia: models and evaluation. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pp. 243–252 (2007).
9. T. Wöhner and R. Peters. Assessing the quality of wikipedia articles with lifecycle based metrics. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration*, p. 16 (2009).
10. D. Hasan Dalip, M. André Gonçalves, M. Cristo, and P. Calado. Automatic quality assessment of content created collaboratively by web communities:

a case study of wikipedia. In *Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries*, pp. 295–304 (2009).
11. J. E. Blumenstock. Size matters: word count as a measure of quality on wikipedia. In *Proceedings of the 17th international conference on World Wide Web*, pp. 1095–1096 (2008).
12. A. Kittur and R. E. Kraut. Harnessing the wisdom of crowds in wikipedia: Quality through coordination. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, CSCW '08, pp. 37–46, ACM, New York, NY, USA (2008). ISBN 978-1-60558-007-4. doi: 10.1145/1460563.1460572. URL `http://doi.acm.org/10.1145/1460563.1460572`.
13. C. Pentzold, Fixing the Floating Gap: The Online Encyclopedia Wikipedia as a Global Memory Place, *Memory Studies.* **2**(2), 255–272 (2009).
14. L. Flekova, O. Ferschke, and I. Gurevych. What makes a good biography?: Multidimensional quality analysis based on wikipedia article feedback data. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, pp. 855–866, ACM, New York, NY, USA (2014). ISBN 978-1-4503-2744-2. doi: 10.1145/2566486.2567972. URL `http://doi.acm.org/10.1145/2566486.2567972`.
15. A. Maass. Linguistic Intergroup Bias: Stereotype Perpetuation through Language. In ed. M. Zanna, *Advanced in Experimental Social Psychology*, pp. 79–121. Academic Press, San Diego, CA (1999).
16. W. von Hippel, D. Sekaquaptewa, and P. Vargas, The Linguistic Intergroup Bias as an Implicit Indicator of Prejudice, *Journal of Experimental Social Psychology.* **33**, 490–509 (1997).
17. C. Beukeboom. Mechanisms of Linguistic Bias: How Words Reflect and Maintain Stereotypic Expectations. In eds. J. Laszlo, J. Forgas, and O. Vincze, *Social Cognition and Communication*, pp. 313–330. Psychology Press, New York, NY (2013).
18. J. Otterbacher. Linguistic bias in collaboratively produced biographies: crowdsourcing social stereotypes? In *ICWSM*, pp. 298–307 (2015).
19. A. Maass, D. Salvi, L. Arcuri, and G. Semin, Language Use in Intergroup Context: The Linguistic Intergroup Bias, *Journal of Personality and Social Psychology.* **57**(6), 981–993 (1989).
20. P. Agathangelou, I. Katakis, I. Koutoulakis, F. Kokkoras, and D. Gunopulos, Learning patterns for discovering domain-oriented opinion words, *Knowledge and Information Systems.* pp. 1–33 (2017).
21. G. Semin and K. Fiedler, The Cognitive Functions of Linguistic Categories in Describing Persons: Social Cognition and Language, *Journal of Personality and Social Psychology.* **54**, 558–568 (1988).
22. L. Coenen, L. Hedebouw, and G. Semin. Measuring Language Abstraction: The Linguistic Category Model Manual. Technical report, Free University Amsterdam, Amsterdam, The Netherlands (June, 2006). URL `http://www.cratylus.org/resources/uploadedFiles/1151434261594-8567.pdf`.
23. P. Winkielman, J. Halberstadt, T. Fazendeiro, and S. Catty, Prototypes are Attractive because they are Easy on the Mind, *Psychological Science.* **17**(9), 799–806 (2006).

24. D. Wigboldus, R. Spears, and G. Semin, When do We Communicate Stereotypes? Influence of the Social Context on the Linguistic Expectancy Bias, *Group Processes & Intergroup Relations*. **8**(3), 215–230 (2005).
25. A. Hunt. *The Linguistic Expectancy Bias and the American Mass Media*. PhD thesis, Temple University, Philadelphia, PA (2011).
26. J. Otterbacher. Crowdsourcing stereotypes: Linguistic bias in metadata generated via gwap. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 1955–1964 (2015).
27. P. Devine and A. Elliot, Are Racial Stereotypes Really Fading? The Princeton Trilogy Revisited, *Personality and Social Psychology Bulletin*. **21**(11), 1139–1150 (1995).
28. P. Agathangelou, I. Katakis, I. Koutoulakis, F. Kokkoras, and D. Gunopulos, Learning patterns for discovering domain oriented opinion words, *Knowledge and Information Systems* (2017).
29. P. Agathangelou, I. Katakis, F. Kokkoras, and K. Ntonas, *Mining Domain-Specific Dictionaries of Opinion Words*, In eds. B. Benatallah, A. Bestavros, Y. Manolopoulos, A. Vakali, and Y. Zhang, *Web Information Systems Engineering – WISE 2014: 15th International Conference, Thessaloniki, Greece, October 12-14, 2014, Proceedings, Part I*, pp. 47–62. Springer International Publishing (2014).
30. M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, pp. 168–177, ACM, New York, NY, USA (2004). ISBN 1-58113-888-1. doi: 10.1145/1014052.1014073. URL `http://doi.acm.org/10.1145/1014052.1014073`.
31. J. D. Gibbons and S. Chakraborti. Nonparametric statistical inference. In *International encyclopedia of statistical science*, pp. 977–979. Springer (2011).
32. P. Willett, The porter stemming algorithm: then and now, *Program*. **40**(3), 219–223 (2006).
33. S. T. Fiske, A. J. Cuddy, and P. Glick, Universal dimensions of social cognition: Warmth and competence, *Trends in cognitive sciences*. **11**(2), 77–83 (2007).
34. S. T. Fiske, A. J. Cuddy, P. Glick, and J. Xu, A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition., *Journal of personality and social psychology*. **82**(6), 878 (2002).
35. C. Wagner, D. Garcia, M. Jadidi, and M. Strohmaier. It's a man's wikipedia? assessing gender inequality in an online encyclopedia. In *Proceedings of the AAAI International Conference on Web and Social Media (ICWSM)*, pp. 454–463 (2015).
36. E. Graells-Garrido, M. Lalmas, and F. Menczer. First women, second sex: gender bias in wikipedia. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, pp. 165–174 (2015).
37. J. Cohen, P. Cohen, S. G. West, and L. S. Aiken, *Applied multiple regression/correlation analysis for the behavioral sciences*. Routledge (2013).
38. H. Abdi and L. J. Williams, Tukey's honestly significant difference (hsd) test,

Encyclopedia of Research Design. Thousand Oaks, CA: Sage. pp. 1–5 (2010).
39. J. M. Hilbe. Logistic regression. In *International Encyclopedia of Statistical Science*, pp. 755–758. Springer (2011).
40. D. Radev, J. Otterbacher, A. Winkel, and S. Blair-Goldensohn, Newsinessence: summarizing online news topics, *Communications of the ACM*. **48**(10), 95–98 (2005).
41. J. Antin, R. Yee, C. Cheshire, and O. Nov. Gender differences in wikipedia editing. In *Proceedings of the 7th international symposium on Wikis and open collaboration*, pp. 11–14 (2011).
42. B. Collier and J. Bear. Conflict, criticism, or confidence: an empirical examination of the gender gap in wikipedia contributions. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*, pp. 383–392 (2012).
43. C. Wagner, E. Graells-Garrido, D. Garcia, and F. Menczer, Women through the glass ceiling: gender asymmetries in wikipedia, *EPJ Data Science*. **5**(1), 5 (2016).
44. E. S. Callahan and S. C. Herring, Cultural bias in wikipedia content on famous persons, *Journal of the Association for Information Science and Technology*. **62**(10), 1899–1915 (2011).
45. I. Protonotarios, V. Sarimpei, and J. Otterbacher. Similar gaps, different origins? women readers and editors at greek wikipedia. In *Wiki@ ICWSM* (2016).
46. T. Wilson, J. Wiebe, and P. Hoffman. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In *Proceedings of the ACL HLT / EMNLP* (2005).