

# A tool to overcome technical barriers for bias assessment in human language technologies

Laura Alonso Alemany<sup>1,2</sup>, Luciana Benotti<sup>1,2</sup>, Lucía González<sup>1,2</sup>, Jorge Sánchez<sup>1</sup>,  
Beatriz Busaniche<sup>2</sup>, Alexia Halvorsen<sup>2</sup> and Matías Bordone<sup>1,2</sup>

<sup>1</sup> Sección de Computación, FAMAF, Universidad Nacional de Córdoba

<sup>2</sup> Fundación Via Libre  
Argentina

## Abstract

Automatic processing of language is becoming pervasive in our lives, often taking central roles in our decision making, like choosing the wording for our messages and mails, translating our readings, or even having full conversations with us. *Word embeddings* are a key component of modern natural language processing systems. They provide a representation of words that has boosted the performance of many applications, working as a semblance of meaning.

Word embeddings seem to capture a semblance of the meaning of words from raw text, but, at the same time, they also distill stereotypes and societal biases which are subsequently relayed to the final applications. Such biases can be discriminatory. It is very important to detect and mitigate those biases, to prevent discriminatory behaviors of automated processes, which can be much more harmful than in the case of humans because their of their scale.

There are currently many tools and techniques to detect and mitigate biases in word embeddings, but they present many barriers for the engagement of people without technical skills. As it happens, most of the experts in bias, either social scientists or people with deep knowledge of the context where bias is harmful, do not have such skills, and they cannot engage in the processes of bias detection because of the technical barriers.

We have studied the barriers in existing tools and have explored their possibilities and limitations with different kinds of users. With this exploration, we propose to develop a tool that is specially aimed to lower the technical barriers and provide the exploration power to address the requirements of experts, scientists and people in general who are willing to audit these technologies.

## 1 Introduction

Machine learning models and data-driven systems are increasingly being used to support decision-making processes. Such processes may affect fundamental rights, like the right to receive an education, or the right to non-discrimination. It is important that models can be assessed and audited to guarantee that such rights are not compromised. Ideally, a wider range of actors should be able to carry out those audits, specially those that are knowledgeable of the context where systems are deployed or those that would be affected.

Although data-driven systems can be audited, such audits often require technical skills that are beyond the capabilities of most of the relevant actors. The technical barrier has become a major hindrance to engage experts and communities in the assessment of the behavior of automated systems. Technicalities are not only a barrier to audit, but they also work as an obscurantism of sorts, making it very difficult for people from other areas to understand the capabilities and limitations of the tools. This makes it very difficult to plan for public policies that take into account the impact of these technologies. That is why are putting an effort to facilitate reduce the technical barriers to understand, inspect and modify some data-driven processes. In particular, we are focusing in a key component in the automatic treatment of language, namely, *word embeddings*.

In the last years the natural language processing (NLP) community has become increasingly worried over bias and stereotypes contained in models and how these biases can affect practical applications, such as personalized job advertisements. In particular, several studies found that word representations learned from corpora contain associations that produce harmful effects when brought into practice, like invisibilization, self-censorship or simply as deterrents. For a critical survey of 146

papers analyzing bias in NLP models see (Blodgett et al., 2020). To address these concerns, many techniques for measuring and mitigating the bias encoded in such word representations, namely word embeddings and language models, have been proposed (Bolukbasi et al., 2016; Caliskan et al., 2017).

Such approaches to dealing with biases tend to put the focus on complex technical questions. But the technical questions are often not the key issue to deploy fairness within Artificial Intelligence. Fairness requires an adequate understanding of complex sociological constructs, often involving phenomena that are not well understood, let alone systematized or formalized. Formalizing such constructs for computational treatment is challenging, and requires the involvement of experts: sociologists, linguists, physicians, psychologists, among others depending on the domain of application of the technology. However, in this paper we argue that fairness metrics and frameworks are based on nuanced technical instruments that hinder understanding and involvement of individuals without extensive technical education including programming skills.

In this project, we propose a methodology that facilitates the exploration of biases in word embeddings, keeping in mind the specific needs of the Latin American region. In Latin America, we need domain experts to be able to carry out these analyses with autonomy, not relying on an interdisciplinary team or on training, since both are usually not available.

This paper is organized as follows. Next section introduces word embeddings and argues that they are the simplest representation of word meanings that are widely used and that embed the biases present in the data on which other NLP technologies are developed. Section 3 presents relevant work in the area of bias diagnostics and mitigation in word embeddings. Section 4 describes two sets of case studies in which two groups of users with different profiles applied this methodology to carry out an exploration of biases, and the observations on usability and requirements that we obtained. The first group are data scientists with different expertise backgrounds but at least a 350 hours of training and education in data science (including the development and evaluation of machine learning models). The second group are social scientists with no previous training in pro-

gramming or technical aspects of machine learning. Section 5 explains our methodology in a worked out use case illustrating how it puts the power to diagnose biases in the hands of people with the knowledge and the societal roles to have an impact on current technologies. Finally, Section 6 includes a summary and an outline of the steps to follow in the prototype development and pilot.

## 2 Word embeddings and lexical bias

In this section we present the basic concepts of how lexical meaning is represented in NLP systems through word embeddings and then we discuss how biases arise in word embeddings.

### 2.1 Basic concepts

Word embeddings are widely used natural language processing artifacts that represent the meaning of words fully automatically, based on their usage in large amounts of text. This is why it is necessary to have large volumes of text to train word embeddings.

The gist of word embeddings consists in representing word meaning as similarity between words. Words are considered similar if they often occur in similar linguistic contexts, more concretely, if they share a high proportion of contexts of co-occurrence. Contexts of co-occurrence are usually represented as the  $n$  words that occur before (and after) the target word being represented. In some more sophisticated structures, contexts may include some measure of word order or syntactic structures. However, most improvements in current word representations have been obtained not by adding explicit syntactic information but by optimizing  $n$  for the NLP task (from a few words to a few dozen words) (Lison and Kutuzov, 2017).

Once words are represented by their contexts of occurrence (in a mathematical data structure called *vector*), the similarity between words can be captured and manipulated as a mathematical distance, so that words that share more contexts are closer, and words that share less contexts are farther apart, as seen in Figure 1. To measure distance, the cosine similarity is used (Singhal, 2001).

Word embeddings are a key component of applications such as text auto-completion or automatic translation, and have been shown to improve the performance of virtually any natural language

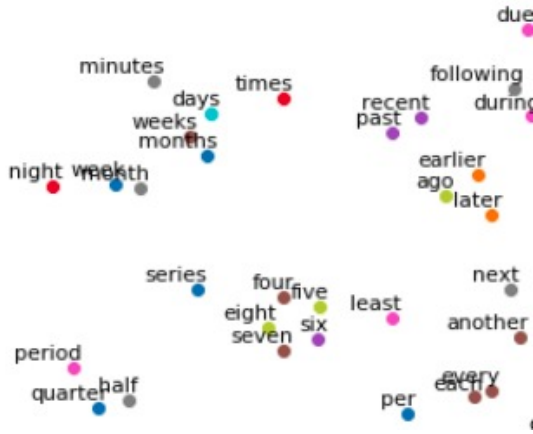


Figure 1: A representation of a word embedding in two dimensions, showing how words are closer in space according to the proportion of co-occurrences they share.

processing task they have been applied to. The problem is that, even if their impact in performance is overall positive, they systematically biased. Thus, even if word embeddings improve general performance, they may damage communities that are the object of those biases.

Word embeddings are biased because they are obtained from large volumes of texts that have underlying societal biases and prejudices. Such biases are carried into the representation which are thus transferred to applications. But since these embeddings are complex, opaque artifacts, working at a subsymbolic level, it is very difficult for a person to inspect them and detect possible biases. This difficulty is even more acute for people without extensive skills in this kind of technologies. In spite of that opacity, readily available pre-trained embeddings are widely used in socio-productive solutions, like rapid technological solutions to scale and speed up the response to the COVID19 pandemic (Aigbe and Eick, 2021).

## 2.2 Assessing bias in word embeddings

Fortunately, given the importance that word embeddings have in language technologies, and the impact that biases may have, in the last years we have seen the emergence of a wide range of tools and techniques to assess bias in word embeddings and language models.

The core methodology to assess biases in word embeddings consists of three main parts, illustrated in Figure 2:

1. Defining a **bias space**, usually binary, delimited by two opposed extremes, as in *male* –

*female*, *young* – *old* or *high* – *low*. Each of the extremes of the bias space is characterized by a list of words, shown at the top of the diagrams in Figure 2.

2. Assessing the behaviour of **words of interest in this bias space**, finding how close they are to each of the extremes of the bias space. This assessment shows whether a given word is more strongly associated to any of the two extremes of bias, and how strong that association is. In Figure 2 it can be seen that the word "nurse" is more strongly associated to the "female" extreme of the bias space, while the word "leader" is more strongly associated with the "male" extreme.
3. Acting on the conclusions of the assessment. The actions to be taken vary enormously across approaches, as will be seen in the next section, but all of them are targeted to **mitigate** the strength of the detected bias in the word embedding.

However varied the approaches to assess bias, every one of them relies on **lists of words** to define the space of bias to be explored. These words have a crucial impact on how and which biases are detected and mitigated, but they are not central in the efforts devoted to this task, as argued in (Antoniak and Mimno, 2021). The methodologies for choosing the words to make these lists are varied: sometimes lists are crowd-sourced, sometimes hand-selected by researchers, and sometimes drawn from prior work in the social sciences. Most of them are developed in one specific context and then used in others without reflection on the domain or context shift. Most previous work uses word lists developed for English, or direct translations from English that do not take into account structural differences between languages (Garg et al., 2018). For example, in Spanish almost all nouns and adjectives are morphologically marked with gender, but this is not the case in English.

Figure 2 illustrates the differences in lexical biases measurements between translations of lists of words in English to Spanish over two different word embeddings in each language: the English embedding is described in (Bolukbasi et al., 2016) and the Spanish in (Cañete et al., 2020). From the 16 words analyzed, in English, 8 are more associated to the "feminine" extreme of the

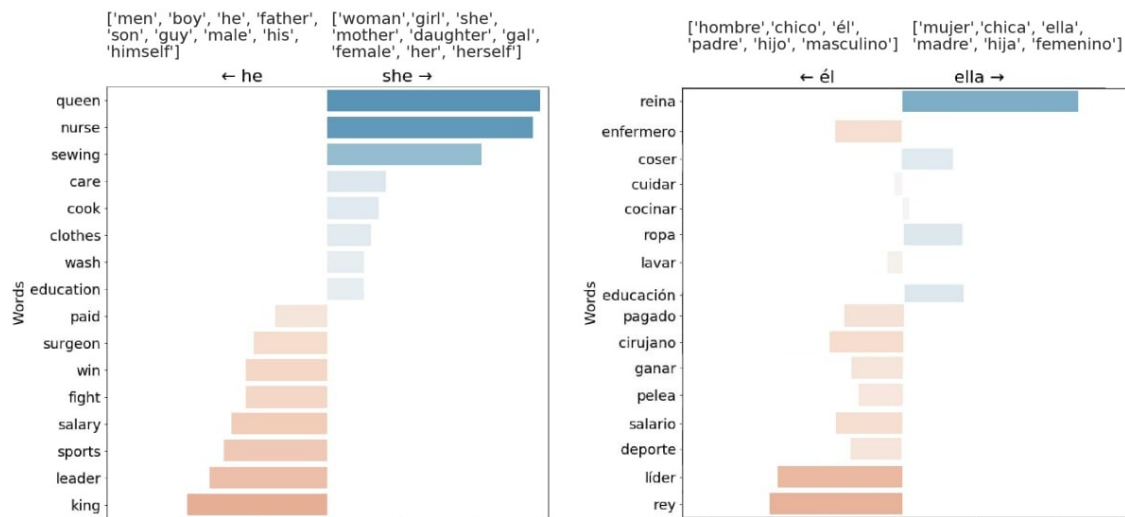


Figure 2: A list of 16 words in English (left) and a translation to Spanish (right) and the similarity of their word embeddings with respect to the list of words “*woman, girl, she, mother, daughter, feminine*” representing the concept “*feminine*”, the list “*man, boy, he, father, son, masculine*” representing “*masculine*”, and translations for both to Spanish. The English word embedding data and training is described in (Bolukbasi et al., 2016) and the Spanish in by (Cañete et al., 2020). From the 16 words of interest, in English, 8 are more associated to the concept of “*feminine*”, while in Spanish only 5 of them are. In particular, “*nurse*” in Spanish is morphologically marked with masculine gender in the word “*enfermero*” so, there is some degree of gender bias that needs to be taken into account to fully account for the behavior of the word. This figure illustrates that methodologies for bias detection developed for English are not directly applicable to other languages. Also, the figure illustrates that the observed biases depend completely on the list of words chosen.

bias space, while in Spanish only 5 of them are. The 3 words with different positions are “*nurse, care and wash*”. In particular, “*nurse*” in Spanish is morphologically marked with masculine gender in the word “*enfermero*”, so it is not gender neutral. This figure illustrates two things. First, the fact that methodologies for bias detection developed for English are not directly applicable to other languages. Second, that the list of words selected to analyze bias have a strong impact on the bias that is shown by the analysis.

### 2.3 Different languages, the same approach?

As illustrated by the previous example, linguistic differences have a big impact on the results obtained by the methodology to assess bias. Representing language idiosyncrasies is a crucial goal in our project, first, because we want to facilitate these technologies to a wider range of actors. Secondly, because to model bias in a given context or a given culture you need to do it in the language of that culture.

Different approaches have been proposed to capture specific linguistic phenomena. The paradigmatic example of linguistic variation are

languages with morphologically marked gender, which can get confused with semantic gender to some extent. Most of the proposals to model gender bias in languages with morphologically marked gender add some technical construct that captures the specific phenomena. That is the case of (Zhou et al., 2019), who add an additional space to represent morphological gender, independent of the dimension where they model semantic gender. This added complexity supposes an added difficulty for people without technical skills.

However, it is not strictly necessary to add technical complexity to capture these linguistic complexities. A knowledgeable exploitation of word lists can also adequately model linguistic particularities. In the work presented here, we adapted the approach to bias assessment proposed by Bolukbasi et al. (2016), resorting to its implementation in the *responsibly* toolkit. This approach is intended for English, and does not envisage morphologically marked gender or the different usage of pronouns. However, we could apply this approach to Spanish, with the following considerations for word lists:

- the extremes of bias cannot be defined by pro-




Figure 3: An assessment of how bias affects gendered words in Spanish. It can be seen that "female nurse", "enfermera" is much more strongly associated to the feminine extreme of bias than "male nurse", "enfermero" is associated to the masculine extreme.

nouns alone, because the pronouns do not occur as frequently or in the same functions in Spanish as in English. Therefore, the lists of words defining the extremes of the bias space need to be designed for the particularities of Spanish, not translated as is.

- with respect to the lists of words of interest to be placed in the bias space, Bolukbasi et al. (2016)'s approach is strongly based on gender neutral words. However, in Spanish most nouns and adjectives are morphologically marked for their semantic gender (as in "enfermera", "female nurse", vs. "enfermero", "male nurse"). To address this difference, we constructed gender neutral words resorting to patterns like: verbs, adverbs, abstract nouns, collective nouns, and adjective suffixes that are gender neutral, like "-ble".
- a proper assessment of bias for Spanish cannot be made with gender-neutral words only, because most nouns and adjectives morphologically marked for their semantic gender, or are morphologically gendered even if they have no semantics for gender (as in "mesa", "table", which is morphologically feminine). To assess bias also in that wide range of words, we constructed word lists containing both words, and compared how far they were positioned with respect to the corresponding extreme of bias. In Figure 3 it can be seen that "female nurse", "enfermera" is much more strongly associated to the feminine extreme of bias than "male nurse", "enfermero" is associated to the masculine extreme.

In fact, putting the focus in a careful, language-aware construction of word lists has the side effect of putting in the spotlight not only linguistic differences, but also other cultural factors like stereotypes, cultural prejudices, or the interaction be-

tween different factors. Thus, in the construction of word lists, different factors need to be taken into account, not only the primary object of research that is the bias.

### 3 Relevant related work

In the last years the study of biases in language technology has been gaining growing relevance, with a variety of approaches accompanied by insightful critiques (Nissim et al., 2020) and successive elaborations that build upon the experience of early proposals (Bolukbasi et al., 2016; Gonen and Goldberg, 2019).

In this section we review relevant related work in four parts. First, we describe that most previous work has focused on a rather narrow set of biases and languages. Then, we discuss the limitations of previous work which focuses on developing algorithms for measuring and mitigating biases automatically. Later we discuss the role that training data play in the process and review work that focuses on the data instead of focusing on the algorithms. Finally, we present existing tools for biases explorations and situate our proposal with respect to them.

#### 3.1 Race and gender, the most studied biases

Most of the published work on biases exploration and mitigation has been produced by computer scientists based on the northern hemisphere, in big labs which have access to large amounts of founding, computing power and data. Unsurprisingly, most of the work has been carried out the English language and for gender and race biases (Garg et al., 2018; Blodgett et al., 2020; Field et al., 2021). Meanwhile there are other biases that deeply affect the global south such as nationality, power and social status. Also aligned with the rest of the NLP area, work has been focused on the technical nuances instead of the more impacting qualitative aspects, like who develops the word list used for bias measurement and evaluation techniques (Antoniak and Mimno, 2021). Since gender-related biases are one of the most studied ones, previous work has shown that the different bias metrics that exist for contextualized and context independent word embeddings only correlate with each other for benchmarks built to evaluate gender-related biases in English (Badilla et al., 2020).

English is a language where morphological

marking of grammatical gender is residual, observable in the form of very few words, mostly personal pronouns and some lexicalizations like “actress - actor”. Some of the assumptions underlying this approach seemed inadequate to model languages where a big number of words have a morphological mark of grammatical gender, like Spanish or German, where most nouns or adjectives are required to express a morphological mark for different grammatical genders. The proposal from the computer science community working on bias was a more complex geometric approach, with a dimension modelling semantic gender and another dimension modelling morphological gender (Zhou et al., 2019). Such approach is more difficult to understand for people without a computer science background, which is usually the case for social domain experts that could provide insight on the underlying causes of the observable phenomena.

In this work we have explored the effects of putting the complexity of the task in constructs that are intuitive for domain experts to pour their knowledge, formulate their hypotheses and understand the empirical data. Conversely, we try to keep the technical complexity of the methodology to a bare minimum.

Lauscher and Glavaš (2019) make a comparison on biases across different languages, embedding techniques, and texts. Zhou et al. (2019) and (Gonen et al., 2019) develop 2 different detection and mitigation techniques for languages with grammatical gender that are applied as a post processing technique. These approaches add many technical barriers that require extensive machine learning knowledge from the person that applies these techniques. Therefore they fail to engage interactively with relevant expertise outside the field of computer science, and with domain experts from particular NLP applications.

### 3.2 Automatically measuring and mitigating

There is a consensus (Field et al., 2021) that what we call bias are the observable (if subtly) phenomena from underlying causes deeply rooted in social, cultural, economic dynamics. Such complexity falls well beyond the social science capabilities of most of the computer scientists currently working on bias in artificial intelligence artifacts. Most of the effort of ongoing research and innovation with respect to biases is concerned with technical

issues. In truth, these technical lines of work are aimed to develop and consolidate tools and techniques more adequate to deal with the complex questions than to build a solid, reliable basis for them. However, such developments have typically resulted in more and more technical complexity, which hinders the engagement of domain experts. Such experts could provide precisely the understanding of the underlying causes that computer scientists lack, and which could help in a more adequate model of the relevant issues.

Antoniak and Mimno (2021) argues that the most important variable when exploring biases in word embeddings are not the automatizable parts of the problem but the manual part, that is the word lists used for modelling the type of bias to be explored and the list of words that should be neutral. They conclude that *word lists are probably unavoidable, but that no technical tool can absolve researchers from the duty to choose seeds carefully and intentionally.*

There are many approaches to "the bias problem" that aim to automatize every step from bias diagnosis to mitigation. Some of these approaches argue that when subjective, difficult decisions on how to model certain biases are involved, automating the process via an algorithmic approach is the solution. (Guo and Caliskan, 2021; Guo et al., 2022; An et al., 2022; Kaneko and Bollegala, 2021) introduce diverse methods to automatically identify (and even mitigate) biases in word embeddings and language models.

On the contrary, the methodology we propose in this paper hides the technical complexity. We develop on the insights of Antoniak and Mimno (2021) by facilitating access to these technologies to domain experts with no technical expertise, so that they can provide well-founded word lists, by pouring their knowledge into those lists. We argue that evaluation should be carried out by people aware of the impact that bias might have on downstream applications. Our methodology focuses on delivering a technique that can be used by such people to evaluate the bias present in text data as we explain in the next section.

### 3.3 A closer look at the training data

(Brunet et al., 2019) trace the origin of word embedding bias back to the training data, and show that perturbing the training corpus would affect the resulting embedding bias. Unfortunately, as

argued in (Bender et al., 2021), most pre-trained word embeddings that are widely used in NLP products do not describe those texts. Interestingly, (Dinan et al., 2020) show that training data can be selected so that biases caused by unbalanced data are mitigated. Also, Kaneko and Bollegala (2021) show that better curated data provides less biased models.

(Brunet et al., 2019) show that debiasing techniques have a are more effective when applied to the texts wherefrom embeddings are induced, rather than applying them directly in the already induced word embeddings. Prost et al. (2019) show that overly simplistic mitigation strategies actually worsen fairness metrics in downstream tasks. More insightful mitigation strategies are required to actually debias the whole embedding and not only those words used to diagnose bias. However, debiasing input texts works best. Curating texts can be done automatically (Gonen et al., 2019) but this has yet to prove that it does not make matters worse. It is better that domain experts devise curation strategies for each particular case. Our proposal is to offer a way in which word embeddings created on different corpora can be compared.

### 3.4 Existing tools for bias exploration

There are tools like WordBias (Ghai et al., 2021) or the Language Interpretability Toolkit (Tenney et al., 2020) that aim to lower the technical barrier that needs to be climbed to use bias detection and mitigation techniques. These kinds of tools provide implementations of many NLP techniques, and graphical interfaces to avoid having to write code.

(Badilla et al., 2020) is an open source python library called WEFÉ which is similar to WordBias in that it allows for the exploration of biases different to race and gender and in different languages. One of the focuses of WEFÉ is the comparison of different automatic metrics for biases measurement and mitigation, however, in order to use this library python programming skills are needed as it doesn't provide a graphical interface.

None of these frameworks were designed with the goal of being usable by social scientists or people without technical and programming skills in general. WordBias is designed for exploring many different kind of biases, however, it has a complex graphical interface not suited for iterative defini-

tion of word lists. The Language Interpretability Tool includes (among many other capabilities) interactive visualizations, integrated fairness and explicability metrics, counterfactual analysis, etc. These tools require extensive Natural Language Processing understanding from the user. In general, most of the toolkits, frameworks and libraries providing functionalities to assess model behaviour with respect to biases are confuse and opaque even for developers with extensive technical knowledge (Richardson et al., 2021).

In this project we will build on the Responsibly (Hod, 2018–) toolkit as the library providing basic functionalities for our tool. We selected Responsibly because it focuses on exploration with the most simple metric of bias (direction projection) rather than on the comparison of different automatic biases measurements as WEFÉ does. Responsibly is based on previous work by Bolukbasi et al. (2016) which establishes an approach to the problem that is intuitive and effective, but at the same time establishes some of the oversimplifications that have been carried on in most of the subsequent work. This technique is presented for the English language and gender and racial biases but it can be applied to any bias that can be modeled as binary in languages that do not have grammatical gender.

However, in order to be able to use the Responsibly toolkit a person should have python programming skills, understanding on word embeddings, natural language processing and bias in word embeddings. Moreover, this person, unless they are working in an interdisciplinary team, should also have expertise on which biases present in word embeddings are relevant to the problem and could affect downstream applications. In our project, we propose to integrate the functionalities provided by Responsibly in a tool that facilitates the engagement of people without technical skills in the process of bias assessment.

## 4 Two case studies

The main goal of our project is to facilitate access to the tools for bias exploration in word embeddings for people without technical skills. To do that, we explored the usability of our adaptation to Spanish of the Responsibly toolkit<sup>1</sup>. In Section 3.4 we explain the reasons for using Responsibly as the basis for our work.

<sup>1</sup><https://docs.responsibly.ai/>

In order to assess where the available tools are lacking and barriers for their use, we conducted two usability studies with different profiles of users: junior data scientists most of them coming from a non-technical background but with a 350 hour instruction in machine learning, and social scientists without technical skills but a 2 hour introduction on word embeddings and bias in language technology. Our objective was to teach these two groups of people how to explore biases in word embeddings, while at the same time gathering information on how they understood and used the technique proposed by Bolukbasi et al. (2016) to model bias spaces. We focused on difficulties to understand how to model the bias space, shortcomings to capture the phenomena of interest and the possibilities the tools offered.

Our design goals are the following. First, reduction of the technical barrier to a bare minimum. Second, a focus on exploration and characterization of bias (instead of focus on a compact, opaque, metric-based diagnostic). Third, an interface that shows word lists in a dynamic, interactive way that elicits, shapes, and expresses contextualized domain knowledge (instead of taking lists as given by other papers, even if these are papers from social scientists). Fourth, guidance about linguistic and cultural aspects that may bias word lists (instead of just translating word lists from another language or taking professions as neutral).

The first group we studied were students at the end of a 1 year nano-degree on Data Science, totalling 180 hours and a project. All of them had some degree of technical skills, the nano-degree providing extensive practice with machine learning, but most came from a non-computer science background and had not training on natural language processing as part of the course. The second group were journalists, linguists and social scientists without technical skills.

Both of the groups were given a 2 hour tutorial, based on a Jupyter Notebook we created for Spanish<sup>2</sup> as explained in Section 2 by adapting the implementation of the Responsibly toolkit (Bolukbasi et al., 2016) done by Shlomi<sup>3</sup>. The groups conducted the analysis on Spanish FastText vectors, trained on the Spanish Billion Word Corpus (Cardellino, 2019) using a 100 thousand word vo-

<sup>2</sup>Our notebook is available here <https://tinyurl.com/ycxz8d9e>

<sup>3</sup>The original notebook for English is available here <https://learn.responsibly.ai/word-embedding/>

cabulary.

#### 4.1 Case Study 1: Junior data scientists

The group composed by junior data scientists were given a 1-hour explanation on how the tools were designed and an example of how it could be applied to explore and mitigate gender bias, the prototypical example of application. This was part of an 8 hour course on Practical Approaches to Ethics in Data Science. As part of the explanation, mitigation strategies built upon the same methodology were also provided, together with the assessment that performance metrics did not decrease in a couple of downstream applications with mitigated embeddings. The presentation of the tools was made for English and Spanish explaining the analogies and differences between both languages to the students, who were mostly bilingual. Then, as a take-home activity, they had to work in teams to explore a 2 dimensional bias space of their choice, different from gender.

During the presentation of the tools, students did not request clarifications or extensive explanations into the nuances of word embeddings, biases represented as lists of words, or the linguistic differences between English and Spanish and the adaptation of the tool. We suspect this was due to the fact that they were conditioned by the methodology of the nano-degree, which was based on classes explaining a methodology and showing how it was applied, followed by practical sessions when they actually applied the methodology to other cases. Thus, they were not trying to be critical, but to reproduce the methodology.

The teams successfully applied the methodology to characterize biases other than gender, namely economic bias (wealthy vs poor), geographical bias (latin american vs north american), ageism (old vs young). Figure 4 illustrates one of their analyses of the bias space defined by *rich* vs *poor*, exploring how negative and positive words were positioned in that space. The figure shows the list of words they used to define the two extremes of the bias space, the concepts of poor and rich. It also shows how words such as *gorgeous* and *violence* are closer to the rich or poor concepts. Students concluded that the concept of poor is more associated with words with a negative sentiment and rich more with a positive sentiment.

They did not report major difficulties or frustrations, and in general reported that they were satis-



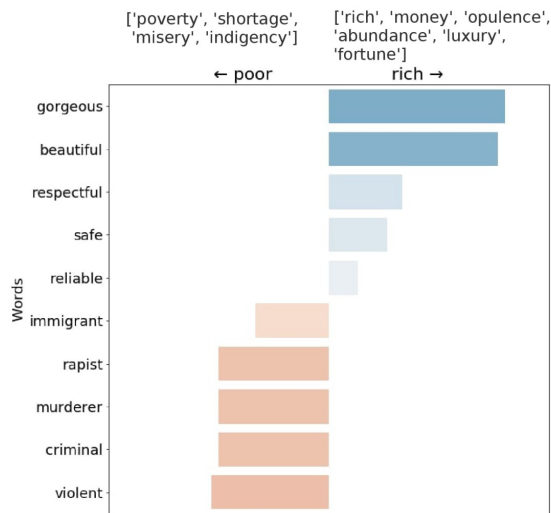


Figure 4: Exploration of the *rich vs poor* bias space carried out by data scientists, showing the words used to define the two extremes of the bias space and how words of interest, like "gorgeous" or "violence", are positioned with respect to each of the extremes. The original exploration was carried out in Spanish and has been translated into English for readability.

fied with their findings by applying the tool. They also applied mitigation strategies available with the Responsibly toolkit, but they made no analysis of its impact.

As it was not required by the assessment, they did not systematically report their exploration process. To gain further insights on the exploration process, we included an observation of the process in the experience with social scientists.

Overall, their application of the tool was satisfactory but rather uncritical. This is not to suggest that the participants were uncritical themselves, but rather that the way the methodology was presented, aligned with a consistent approach to applying methodologies learnt throughout the course, inhibited a more critical, nuanced exploitation of the tool.

#### 4.2 Case Study 2: Social Scientists

The social scientists were presented the tool as part of a twelve-hour workshop on tools to inspect NLP technologies. A critical view was fostered, and we explicitly asked participants for feedback to improve the tool.

There was extensive time within the workshop devoted to carry out the exploration of bias. We observed and sometimes elicited their processes, from the guided selection of biases to be explored,

based on personal background and experience, to the actual tinkering with the available tools. Also, explicit connections were made between the word embedding exploration tools available via Responsibly and an interactive platform for NLP model understanding, the Language Interpretability Tool (LIT) (Tenney et al., 2020), which was also inspiring for participants as to what other information they could obtain that could enrich their analysis in exploration.

Since there was no requirement for a formal report, bias exploration was not described systematically. Different biases were explored, in different depths and lengths. Besides gender and age, also granularities of origin (cuban - north-american) and the intersection between age and technology were explored. Participants were creative and worked collaboratively to find satisfactory words that represented the phenomena that they were trying to assess. They were also insightful in their analysis of the results: they were able to discuss different hypotheses as to why a given word might be further in one of the extremes of the space than another.

Figure 5 illustrates one of their analysis of the bias space defined by the concepts of *young* and *old*, using the position of verbs in this space to explore bias. In this analysis they did not arrive to any definite conclusions, but found that they required more insights on the textual data wherefrom the embedding had been inferred. For example, they wanted to see actual contexts of occurrence of "sleep" or "argument" with "old" and related words, to account for the fact that they are closer to the extreme of bias representing the "old" concept. Analyzing these results, the group also realized there were various senses associated to the words representing the "old" concept, some of them positive and some negative. They also realized that the concept itself may convey different biases, for example, respect in some cultures or disregard in others. Such findings were beyond the scope of simple analysis of the embedding, requiring more contextual data to be properly analyzed and subsequently addressed.

Participants were active and creative while requiring complementary information about the texts from which word embeddings had been inferred. Many of these requirements will be included in the prototype of the tool we are devising, such as the following: frequencies of the words

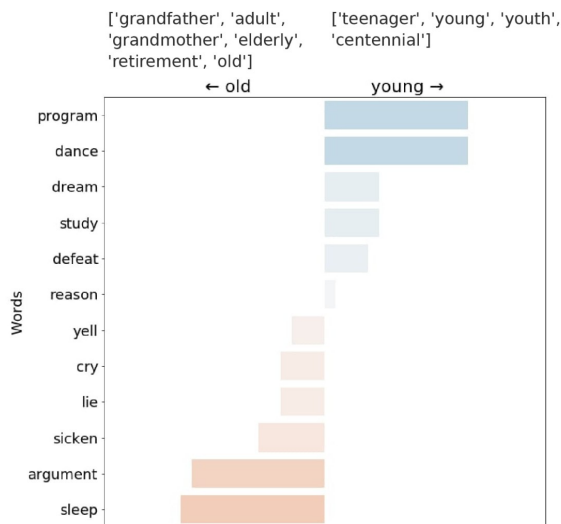


Figure 5: Exploration of the *old* vs *young* bias space carried out by social scientists, showing the words used to define the two extremes of the bias space and how words of interest, like "dance" or "sleep", are positioned with respect to each of the extremes. The original exploration was carried out in Spanish and has been translated into English for readability.

being explored concordances of the words being explored, that is, being able to examine the actual textual productions when building word lists, especially for the extremes that define the bias space, suggestions of similar words or words that are close in the embedding space would be useful, as it is often difficult to come up with those and the space is better represented if lengthier lists are used to describe the extremes functionalities and a user interface that facilitates the comparison between different delimitations (word lists) of the bias space, or the same delimitations in different embeddings.

They also stated that it would be valuable that the tool allowed them to explore different embeddings, from different time spans, geographical origin, publications, genres or domains. For that purpose, our prototype includes the possibility to upload a corpus and have embeddings inferred from that corpus, which can then be explored.

In general, we noted that social scientists asked for more context to draw conclusions on the exploration. They were critical of the whole analysis process, with declarations like "I feel like I am torturing the data". Social scientists study data in context, not data in itself, as is more common in the practice of data scientists. Also, the con-

text where the tool was presented also fostered a more critical approach. We asked participants to formulate the directions of their explorations in terms of hypotheses. Such requirement made it clear that more information about the training data was needed in order to formulate hypotheses more clearly.

### 4.3 Wishlist for the prototype

In these two case studies we could test the adequacy of some aspects of our approach, describe the shortcomings of the tools that we are currently using at this stage of development, and also reassess some design decisions with respect to the proposed tool.

Summing up, below we highlight those limitations that we have identified during this case studies that we will address in the prototype. We then discuss intrinsic limitations of the current capabilities of the tools for bias exploration.

Working with **multiword expressions** as a linguistic unit.

Being able to retrieve the actual **contexts of occurrence** of words in the corpus wherefrom the embedding has been obtained.

Being able to compare how the same set of words is related to **extremes of bias defined with different lists of words**.

Being able to explore and **compare different embeddings** and how the same bias is represented there.

Integrating **different bias spaces into a single visualization**, so that we can see how a word is associated to different meanings at the same time.

#### 4.3.1 Limitations of the approach

The most evident, frustrating limitation is the fact that only binary biases can be represented with this approach, because of the way the space of bias is mathematically defined, with its two extremes. However, this limitation was somewhat overcome with intersectionality expressed in the word lists themselves, that is, building word lists that were intersections, for example, an word list that is an intersection between gender and age would contain the words "grandmother", "granny", "old lady" in one extreme and "girl", "lass", "young lady" in the other.

One of the very interesting shortcomings to arise was that ambiguities in words could not be

accounted for: does a word have different senses, different morphosyntactic categories? These are all merged in a single representation of the word, which can then introduce noise in the representation of an extreme of bias. For example, the word "gusano" used for "Cuban" has the primary meaning "worm", so when it is used to define an extreme of bias for Cubans, it takes with it many animal-related contexts. We will test different ways to address this problem in usability studies.

Also, we found that in some cases one of the extremes of the bias direction was lexicalized, but not the other, as is the case of "football". The methodology is strongly binary, and cannot account for such cases, frequent as they may be. It also falls short to account for intersectional bias, although some approximations can be made by using words in the extremes that include a combination of meanings, thus representing intersections.

#### 4.4 Conclusions of the case studies

With these study cases, we show that reducing the technical complexity of the tool and explanations to the minimum fosters engagement by domain experts. Providing intuitive tools like word lists, instead of barriers like vectors, allows them to formalize their understanding of the problem, casting relevant concepts and patterns into those tools, formulating hypotheses in those terms and interpreting the data. Such engagement is useful in different moments in the software lifecycle: error analysis, framing of the problem, curation of the dataset and the artifacts obtained.

Our conclusion is that the inspection of biases in word embeddings can be understood without most of the underlying technical detail. However, the Responsibly toolkit adapted to Spanish needs the improvements we discuss in this section and develop as an applied research plan in Section 6.

### 5 User story for our prototype

Up to this point we have motivated the need for bias assessment in language technologies and in word embeddings in particular, we have explained our differences in the framing of the solution with respect to existing tools, we have discussed the artificial technical barriers of existing approaches, that hinder the engagement of actual experts in the exploration process, and we have put together a wishlist from social scientists describing their ideal tool for exploration of bias in word embed-

dings. In this Section we are describing a user story that presents the intended functionalities of the proposed tool, and we finish in the following Section with a detail of the steps to develop it, if the next phase of the project is granted.

We would like to note that this user story was originally thought as situated in Argentina, the local context of this project. However, in order to make understanding easier for non-Spanish speaking readers, we adapted the case to work with English, and consequently situated the use case in the United States. The original use case in Spanish will be published later on as a media piece.

**The users.** Marilina is a data scientist working on a project to develop a software that helps the public administration to classify citizens' requests and route them to the most adequate department. Tomás is a social worker within the non-discrimination office, and wants to assess the possible discriminatory behaviours of such software.

**The context.** Marilina addresses the project as a supervised text classification problem. To classify new texts from citizens, they are compared to documents that have been manually classified in the past. New texts are assigned the same label as the document that is most similar. Calculating similarity is a key point in this process, and can be done in many ways: programming rules explicitly, via machine learning with manual feature engineering or by deep learning, where a key component is word embeddings. Marilina observes that the latter approach has the least classification errors on the past data she separated for evaluation (the so called test set). Moreover, deep learning seems to be the preferred solution these days, it is often presented as a breakthrough for many natural language processing tasks. So Marilina decides to pursue that option.

An important component of the deep learning approach are word embeddings. Marilina decides to try a well-known pre-trained word embedding, pre-trained on Wikipedia content. When she integrates it in the pipeline, there is a boost in the performance of the system: more texts are classified correctly in her test set.

**Looking for bias.** Marilina decides to look at the classification results beyond the figures. Being a descendant of Latin American immigrants, she looks at documents related to this societal group. She finds that applications for small busi-

ness grants presented by Latin American immigrants or citizens of Latin American descent are sometimes misclassified as immigration issues and routed to the wrong department. These errors imply a longer process to address these requests in average, and sometimes misclassified requests get lost. In some cases, this mishap makes the applicant drop the process.

**Finding systematicities in errors.** Intrigued by this behaviour of the automatic pipeline, she makes a more thorough research into how requests by immigrants are classified, in comparison with requests by non-immigrants. As she did for Latin American requests, she finds that documents presented by other immigrants have a higher misclassification rate than the non immigrants requests. She suspects that other societal groups may suffer from higher misclassification rates, but she focuses on Latin American immigrants because she has a better understanding of the idiosyncrasy of that group, and it can help her establish a basis for further inquiry. She finds some patterns in the misclassifications. In particular, she finds that some particular business, like hairdressing or bakery, accumulate more errors than others.

**Finding the component responsible for bias.** She traces the detail of how such documents are processed by the pipeline and finds that they are considered most similar to other documents that are not related to professional activities, but to immigration. The word embedding is the pipeline component that determines similarities, so she looks into the embedding. She finds that there is the `responsibly.ai` library to inspect bias word embeddings, and uses some of its utilities to assess bias: the projection of neutral words in the direction of bias or the metric to measure bias. She defines a bias space with "*Latin American*" in one extreme and "*North American*" in the other, and checks the relative position of some professions with respect to those two extremes, as can be seen in Figure 6, left. She finds that, as she suspected, some of the words related to the professional field are more strongly related to words related to Latin American than to words related to North American, that is, words like "*hairdresser*" are closer to Latin American. However, the words more strongly associated to North American do not correspond to her intuitions. She is at a loss as to how to proceed with this inspection beyond

the anecdotal findings, and how to take action with respect to the findings. That is when she calls for help to the non-discrimination office.

**Assessing harm.** The non-discrimination office appoints Tomás for the task of assessing the discriminatory behavior of the software. Briefed by Marilina about her findings, he finds that misclassifications do involve some harm to the affected people that is typified among the discriminatory practices that the office tries to prevent. Misclassification implies that the processes are made longer than for other people, because they need to be reclassified manually before they can actually be taken care of. Sometimes, they are simply dismissed by the wrong civil servant, resulting in unequal denial of benefits. In many cases, the mistake itself has a negative effect on the self-perception of the issuer, making them feel less deserving and discouraging the pursuit of the grant or even the business initiative. Tomás can look at the output of the system, but he cannot see a rationale for the system's (mis)classifications, since he doesn't know the technical details of the processing.

**Understanding word embeddings without unnecessary technicalities.** Tomás understands that there is an underlying component of the software that is impacting in the behaviour of classification. Marilina explains to him that it is a pre-trained word embedding, and that a word embedding is a projection of words from a sparse space where each context of occurrence is a dimension into a dense space where there are less dimensions, obtained with a neural network, and each word is a vector with values in each of those dimensions. Tomás feels that understanding the embedding is beyond his capabilities. Then Marilina realizes and explains to him that words are represented as a summary of their contexts of occurrence in a corpus of texts, but this cannot be directly seen, but explored using similarity between words, so that more similar words are closer.

**Finding an intuitive tool for bias exploration.** She shows him some of the tools available to assess bias in the `responsibly` toolkit, but Tomás cannot program and feels overwhelmed by the code. Then she looks for a tool that does not require this kind of expertise and finds the prototype develop by our project. This is a tool accessible for Tomás, that presents the key concepts in

an intuitive way, and that he can manipulate autonomously. Then Tomás feels empowered and goes on with the exploration.

**Get to know the corpus underlying the embedding.** To begin with, Tomás wants to explore the words that are deemed similar to "*Latin American*", because he wants to see which words may be strongly associated to the concept, besides what Marilina already observed. He finds that the embedding has been trained with texts from newspapers. Most of the news containing the word Latin American deal with catastrophes, troubles and other negative news from Latin American countries, or else portray stereotyped Latin Americans, referring to the typical customs of their countries of origin rather than to their facets as citizens in the United States. With respect to business and professions, Latin Americans tend to be depicted in accordance with the prevailing stereotypes and historic occupations of that societal group in the States, like construction workers, waiters, farm hands, etc.

He concludes that this corpus, and, as a consequence, the word embedding obtained from it, contains many stereotypes about Latin Americans which are then relayed to the behaviour of the classification software, associating certain professional activities and demographic groups more strongly with immigration than with business. Possibly they will have to find a better word embedding, that does not have such biases or so marked, but he wants to characterize the biases first so that he can make a quicker assessment in other word embeddings.

**Understanding bias exploration in word embeddings without unnecessary technicalities.** Tomás needs to focus his exploration of the word embedding in the bias of interest, in this case, in the Latin American versus North American. Marilina resorts to the available materials for our tool to explain bias definition and exploration easily to Tomás. He quickly grasps the concepts of bias space, definition of the space by lists of words, assessment by observing how words are positioned within that space, and exploration by modifying lists of words, both defining the space and positioned in the space. He gets more insights on the possibilities of the techniques and on possible misunderstandings by reading examples and watching short tutorials that can be found with

the tool. He then understands that word ambiguity may obscure the phenomena that one wants to study, that word frequency has a big impact, and that language-specific phenomena, like grammatical gender or levels of formality, need to be carefully taken into account.

**Formalize a starting point for bias exploration.** Now Tomás is able to systematize bias exploration, with the final objective to make a report and take principled, informed action to prevent and redress any discriminatory behaviour that the automated process may have deployed. First, he builds the sets of words that will be representing each of the extremes of the bias space. He realizes that Marilina's approach with only one word in each extreme is not quite robust, because it may be heavily influenced by properties of that single word. That is why he defines each of the extremes of the bias space with longer word lists, and experiments with different lists and how they determine the relative position of his words of interest. Words of interest are the words being positioned in the bias space, words that Tomás wants to characterize with respect to this bias because he suspects that their characterization is one of the causes for the discriminatory behavior of the classification software.

To find words to include in the word lists for the extremes, Tomás resorts to the functionality of finding the closest words in the embedding. Using "*Latin American*" as a starting point, he finds other similar words like "*latino*", and also nationalities of Latin America.

He also explores the contexts of his words of interest. Doing this, he finds that "*shop*" occurs in many more contexts than he had originally imagined, many with different meanings, for example, short for Photoshop. This makes him think that this word is probably not a very good indicator of the kind of behavior in words that he is trying to characterize. He also finds that some professions that were initially interesting for him, like "*capoeira trainer*" are very infrequent and their characterization does not have a correspondence with his intuition about the meaning and use of the word, so he discards them.

Finally, he is satisfied with the definition provided by the word lists that can be seen in Figure 6, right. With that list of words, the characterization of the words of interest shows tendencies that have a correspondence with the misclassification

tions of the final system: applications from hairdressers, bakers, dressmakers of latino origin or descent are misclassified more often than applications for other kinds of businesses.

**Report biases and propose a strategy for mitigation** . With this characterization, Tomás can make a detailed report of the discriminatory behavior of the classification system. From the beginning, he could describe the systematicities that can be found in errors, which affect more often people of Latin American descent applying for a subsidies for a certain kind of business. However, his understanding of the underlying word embedding allows him to describe a pattern of behavior, going beyond the cases that he has actually been able to see as misclassified by the system, and predicting other cases.

Moreover, understanding the pattern of behavior allows him to describe properties of the underlying corpus that will make for a better word embedding, with less biases. He can propose strategies like editing the sentences containing hairdressers, designers and bakers to show a more balanced mix of nationalities and ethnicities in them.

**Wider and deeper.** If time allows, he will also be able to explore intersectionality and compare this with other embeddings, all in a visual interface with intuitive concepts

## 6 Applied research proposal

In this section we first describe the functionalities that the prototype will include. Then we list the activities planned to implement such functionalities. In the last subsection we explain the crucial role that the pilot experiences will have in the iterative development of our prototype. We also present our regional partners for the pilot experiences.

### 6.1 Functionalities of the prototype

The design principles and functionalities of the prototype that we will prioritize are the following:

- Focus on **exploration** of bias, not on metric-based diagnosis or mitigation within word embeddings, as is the case of existing approaches.
- Facilitate the **comparison of word lists** used to define the bias space, as a mean to assess the effects of different words and their combinations to define bias space.

- Allow the **comparison across different collections of texts**: different times, different regions, different authors, different media, etc.
- Facilitate the **inspection** of the contexts of occurrence of the words in the lists, their frequency, and any other information that allows to identify reasons for unsuspected behaviours or biases, like infrequent words being strongly associated to other words merely by chance occurrences.
- Characterize **multi-word expressions** as a lexical unit.
- Reduce **technical complexity** to the minimum necessary, and substituting highly technical concepts by more intuitive concepts whenever possible.

### 6.2 Planned activities

In this paper we have presented a methodology for the kind of involvement that can enrich approaches to bias exploration of NLP artifacts with the necessary domain knowledge to adequately model the problems of interest. Up to this point in our research, we have relied on the tools provided by the **responsibly** toolkit and our adaptation of it to Spanish. Given our fieldwork with different types of users, we are now in the point where we can begin the development of a standalone tool that integrates both the technical capabilities of existing approaches and the design requirements valued by our approach and the users we studied.

In this section we detail the activities that we will carry out to develop such prototype.

The development of the prototype will be iterative, following an agile methodology that will be validated at various points during the development process with usability studies that we describe in the next subsection, followed by a more complete pilot. Table 1 organizes the following activities in 7 months from July 2022 to January 2023.

1. Provide a graphical interface inspired by some of the HCI ideas in the tool WordBias that we present in Section 3.
2. Provide comparative visualizations that record the history of interactions with the prototype, allowing to compare:
  - variations in the extremes of the bias space




Figure 6: Different characterizations of the space of bias "Latin American" vs "North American", with different word lists created by a data scientist (left) and a social scientist (right), and the different effect to define the bias space as reflected in the position of the words of interest (column in the left).

- variations across different embeddings, diachronically (e.g., newspapers)
  - combinations of different spaces (i.e., intersectionality)
- Host our prototype in [huggingface](https://huggingface.co/)<sup>4</sup>, to import pre-trained embeddings and to offer our tool to the NLP community of practitioners that huggingface has.
  - Possibility to train word embeddings, given a corpus, provide metrics of reliability of the embedding word by word.
  - Show the following additional information about the words
    - frequency with respect to corpus size
    - concordances, context of occurrences
    - n most similar words
    - average similarity with n most similar words
  - Define metrics assessing quality of word lists, based on their statistical properties.
  - Extend embeddings with n-grams as explained in above with multi-word expressions.
  - Suggest mitigation strategies that involve comparing different corpora or modifying the

- original corpora. For example, a corpus in Spanish could be made gender neutral before training word embeddings by using the neutral gender 'e'.
- Assess whether our methodology for exploring biases could be applied to contextual embedding methods used in large language models (Zhao et al., 2019; Sedoc and Ungar, 2019).
- Usability studies for agile development that we describe below.
- Integration with public policy practice.

Months	J	A	S	O	N	D	J
1. graphical interface	X		X				
2. visualizations	X	X					
3. huggingface			X				
4. train embeddings			X				
5. word information				X			
6. word list quality				X			
7. n-gram embeddings				X	X		
8. corpus mitigation			X			X	
9. usability studies			X	X			
10. public policy					X	X	X
11. prototype delivery							X

Table 1: Planification of the activities month by month from July 2022 to January 2023.

<sup>4</sup><https://huggingface.co/>

### 6.3 Usability studies and pilot

The general goal of our project is to develop usable tools for anyone without a technical profile or specialization in AI and data processing natural language, but, at the same time, we seek to promote the usage of this tool within social science communities, to integrate it with their established practices. That is why we are working with FLACSO (University of Latin American Social Sciences) as a strategic partner to carry out usability studies and a pilot.

FLACSO is an institution with a long history in the region, with offices in Argentina, Costa Rica, Ecuador and Mexico, and students and researchers from all over Latin America, working in areas of social sciences linked to gender studies, migrations, native peoples, culture and communication, economy, climate change and emergent issues from a public policy perspective. At the Argentine headquarters, FLACSO has an area of bioethics, intellectual property and public goods. Some members of our team and in the extended Fundación Via Libre team are graduates, professors and FLACSO researchers, so we have well established alliances and a long history of joint work between Fundación Vía Libre and FLACSO Argentina.

We envision two types of interaction with the FLACSO community. First, usability studies with small focus groups to obtain immediate qualitative feedback on the features that we will be implementing during the development of the prototype, and that allow us to make rapid adjustments to the development. These are planned as "usability studies" within the applied research plan, and are aimed to choose more adequate options for the graphical user interface, visualizations, usability of the functionalities to load and train embeddings, and to inspect the underlying corpora and contexts of occurrence of words.

Then, if the following phase of the project is granted, we plan to deploy a study on bigger groups, to assess difficulties and potentialities of the tool in use in a wider population of users. Then, we will use these findings to develop accompanying materials to facilitate and promote the use of this tool in different contexts: tutorials, videos, media pieces, documentation and online help pages.

We plan to put a special effort to facilitate the integration of this tool and similar approaches as

an integral part of the processes to monitor, assess and mitigate discriminatory behaviours of language technologies. We will focus our efforts to integrate the tool within the workflow of agents of public policy, either regulatory agencies, or government agencies that need to make data-informed decisions. We are thinking of use cases within the consumer protection area, as in inspection of biases in language technologies of state providers or products with large audiences, or assessment of bias after users' complaints.

### 6.4 Integration with public policy

As stated in the introduction, one of our goals in lowering the technical barriers to access tools for bias assessment is precisely to make it possible for decision makers to understand the capabilities and limitations of the current language processing applications, and to plan for public policies that take into account the impact of these technologies.

Within the scope of our work, we will be researching how this tool can be integrated with efforts to assess whether language technologies are compliant with regulations and legal standards. A tool for bias assessment seems necessary to carry out compliance assessments, because the declared objective of transparency from the technologies does not seem to be enough to solve the problems derived from the potential associated problems.

We will also be researching how this tool can be integrated in a more proactive way to prevent discriminatory behavior before automations are effectively deployed. We are thinking about a collaborative benchmark to detect different kinds of biases that tools can be run through before being deployed, as a quality standard for software with social impact.

Various international organizations, including UNESCO, have initiated processes for the creation of ethical frameworks for the development, adoption and implementation of artificial intelligence technologies in public life. One of the guidelines in such frameworks states that all technology should have, before it is put into operation, various metrics and tools for an appropriate impact assessment. We will be working on this tool with this framework of reference.

#### 6.4.1 A note on the legal framework: Human Rights in Latin America

The Universal Declaration of Human Rights approved in 1948 establishes the general normative



framework saying that "All human beings are born free and equal in dignity and rights and, endowed as they are with reason and conscience, they must behave fraternally towards one another" and continues "Everyone is entitled to the rights and freedoms set forth in this Declaration, without distinction of any kind as to race, color, sex, language, religion, political or other opinion, national or social origin, property, birth or other status."

However, human society is far from fulfilling this vision of those who promoted and promote equality and justice as fundamental rights. Discrimination based on ethnicity, color, gender, language, religion, opinion, economic position and diversity of abilities continue to be the order of the day despite the legal frameworks established for its eradication. In recent years, language technologies have been a major agent of discrimination, not least because of their massive scale, where a single program can affect millions of people throughout millions of devices working 24 hours a day. Discriminatory behaviors from automated systems can be subjected to anti-discrimination laws.

In Argentina, since 1994, the "American Declaration of the Rights and Duties of Man; the Universal Declaration of Human Rights; the American Convention on Human Rights; the International Covenant on Economic, Social, and Cultural Rights; the International Covenant on Human Rights have constitutional status." Civil and Political Rights and its Optional Protocol, the Convention on the Prevention and Punishment of the Crime of Genocide, the International Convention on the Elimination of All Forms of Racial Discrimination, the Convention on the Elimination of All Forms of Discrimination Against Women; the Convention Against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment; the Convention on the Rights of the Child; in the conditions of their validity, they have constitutional hierarchy, do not repeal any article of the first part of this Constitution and must be understood as complementary of the rights and guarantees recognized by it. They can only be denounced, where appropriate, by the National Executive Power, prior approval of two thirds of all the members of each Chamber.

That same constitutional framework is shared by the countries of all of Latin America, the region from which we carried out this investigation. It is this framework of Human Rights that estab-

lishes the normative base that later derives in specific regulations on each of the fields addressed in this field. However, we understand that specific legislation is not necessary for the uses and applications of automated decision-making systems, but rather the possibility of applying the current regulatory and legal frameworks, with the consequent fulfillment of guarantees and rights established with the highest legal rigor.

## References

- Steve Aibuedefe Aigbe and Christoph Eick. 2021. [Learning domain-specific word embeddings from covid-19 tweets](#). In *2021 IEEE International Conference on Big Data (Big Data)*, pages 4307–4312.
- Haozhe An, Xiaojiang Liu, and Donald Zhang. 2022. [Learning bias-reduced word embeddings using dictionary definitions](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1139–1152, Dublin, Ireland. Association for Computational Linguistics.
- Maria Antoniak and David Mimno. 2021. [Bad seeds: Evaluating lexical methods for bias measurement](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1889–1904, Online. Association for Computational Linguistics.
- Pablo Badilla, Felipe Bravo-Marquez, and Jorge Pérez. 2020. [WEFE: the word embeddings fairness evaluation framework](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 430–436.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in](#)

- NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, page 4356–4364, Red Hook, NY, USA. Curran Associates Inc.
- Marc-Etienne Brunet, Colleen Alkalay-Houlihan, Ashton Anderson, and Richard Zemel. 2019. [Understanding the origins of bias in word embeddings](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 803–811. PMLR.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Cristian Cardellino. 2019. [Spanish Billion Words Corpus and Embeddings](#).
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish pre-trained bert model and evaluation data. In *Workshop Practical Machine Learning for Developing Countries: learning under limited resource scenarios at International Conference on Learning Representations (ICLR)*.
- Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. 2020. [Queens are powerful too: Mitigating gender bias in dialogue generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8173–8188, Online. Association for Computational Linguistics.
- Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. [A survey of race, racism, and anti-racism in NLP](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1905–1925, Online. Association for Computational Linguistics.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. [Word embeddings quantify 100 years of gender and ethnic stereotypes](#). *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Bhavya Ghai, Md Naimul Hoque, and Klaus Mueller. 2021. [Wordbias: An interactive visual tool for discovering intersectional biases encoded in word embeddings](#). In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, CHI EA '21*, New York, NY, USA. Association for Computing Machinery.
- Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In *NAACL-HLT*.
- Hila Gonen, Yova Kementchedjhieva, and Yoav Goldberg. 2019. [How does grammatical gender affect noun representations in gender-marking languages?](#) In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 463–471, Hong Kong, China. Association for Computational Linguistics.
- Wei Guo and Aylin Caliskan. 2021. [Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, page 122–133, New York, NY, USA. Association for Computing Machinery.
- Yue Guo, Yi Yang, and Ahmed Abbasi. 2022. [Auto-debias: Debiasing masked language models with automated biased prompts](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1012–1023, Dublin, Ireland. Association for Computational Linguistics.
- Shlomi Hod. 2018–. [Responsibly: Toolkit for auditing and mitigating bias and fairness of machine learning systems](#). [Online; accessed <today>].

- Masahiro Kaneko and Danushka Bollegala. 2021. [Dictionary-based debiasing of pre-trained word embeddings](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 212–223, Online. Association for Computational Linguistics.
- Anne Lauscher and Goran Glavaš. 2019. [Are we consistently biased? multidimensional analysis of biases in distributional word vectors](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\*SEM 2019)*, pages 85–91, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pierre Lison and Andrey Kutuzov. 2017. [Redefining context windows for word embedding models: An experimental study](#). In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 284–288, Gothenburg, Sweden. Association for Computational Linguistics.
- Malvina Nissim, Rik van Noord, and Rob van der Goot. 2020. [Fair is better than sensational: Man is to doctor as woman is to doctor](#). *Computational Linguistics*, 46(2):487–497.
- Flavien Prost, Nithum Thain, and Tolga Bolukbasi. 2019. [Debiasing embeddings for reduced gender bias in text classification](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 69–75, Florence, Italy. Association for Computational Linguistics.
- Brianna Richardson, Jean Garcia-Gathright, Samuel F. Way, Jennifer Thom, and Henriette Cramer. 2021. [Towards fairness in practice: A practitioner-oriented rubric for evaluating fair ml toolkits](#). In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA. Association for Computing Machinery.
- João Sedoc and Lyle Ungar. 2019. [The role of protected class word lists in bias identification of contextualized word representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 55–61, Florence, Italy. Association for Computational Linguistics.
- Amit Singhal. 2001. Modern information retrieval: A brief overview. *IEEE Data Engineering Bulletin*, 24(4):35–43.
- Ian Tenney, James Wexler, Jasmijn Bastings, Tolga Bolukbasi, Andy Coenen, Sebastian Gehrmann, Ellen Jiang, Mahima Pushkarna, Carey Radebaugh, Emily Reif, and Ann Yuan. 2020. [The language interpretability tool: Extensible, interactive visualizations and analysis for NLP models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 107–118, Online. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. [Gender bias in contextualized word embeddings](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 629–634, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pei Zhou, Weijia Shi, Jieyu Zhao, Kuan-Hao Huang, Muhao Chen, Ryan Cotterell, and Kai-Wei Chang. 2019. [Examining gender bias in languages with grammatical gender](#). 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 5276–5284, Hong Kong, China. Association for Computational Linguistics.