

Inferring Gender from Names on the Web: A Comparative Evaluation of Gender Detection Methods

Fariba Karimi*, Claudia Wagner*, Florian Lemmerich*,
Mohsen Jadidi*, and Markus Strohmaier* †
* GESIS - Leibniz Institute for the Social Sciences
† University of Koblenz-Landau
{firstname.lastname}@gesis.org

ABSTRACT

Computational social scientists often harness the Web as a “societal observatory” where data about human social behavior is collected. This data enables novel investigations of psychological, anthropological and sociological research questions. However, in the absence of demographic information, such as gender, many relevant research questions cannot be addressed. To tackle this problem, researchers often rely on automated methods to infer gender from name information provided on the web. However, little is known about the accuracy of existing gender-detection methods and how biased they are against certain sub-populations. In this paper, we address this question by systematically comparing several gender detection methods on a random sample of scientists for whom we know their full name, their gender and the country of their workplace. We further suggest a novel method that employs web-based image retrieval and gender recognition in facial images in order to augment name-based approaches. Our findings show that the performance of name-based gender detection approaches can be biased towards countries of origin and such biases can be reduced by combining name-based and image-based gender detection methods.

Keywords: Gender detection, Discrimination on Web, Computational Social Science

1. INTRODUCTION

The Web enables studies of human social behavior on a very large scale. For many research questions, demographic information about individuals (such as age, gender or ethnic background) is highly beneficial but often particularly difficult to obtain.

This has led previous research to employ different methods for inferring the gender of individuals from names. For example, in [2] the authors determine the gender of individuals using the name repository from the US Social Security Administration and study the relationship between gender and job performance among brokerage firm. Mislove et al. [4] used the same name repository to infer the gender of Twitter users by mapping their self-reported names to the name database. In another study the authors aim to study gender disparities in science and infer the gender of scientists based on a similar approach [6]. Unfortunately, most previous work does not provide information on how accurate different gender detection methods are and/or how biased they are against certain sub-populations. Although crowd sourcing methods can be seen

as an alternative for automated gender detection methods, they do not scale well and are expensive. In the absence of a full name, more sophisticated methods such as supervised machine learning models are used to harness the users content for detecting the gender (see e.g. [5]). Yet, separate models are needed for gender detection methods in each language community [1].

In this paper we evaluate and compare frequently used name-based gender detection methods. We report overall accuracy and also bias, i.e., deviating accuracy for different demographic sub-populations (e.g. men, women and people living in different countries). Moreover, we propose novel methods that increase the accuracy of gender detection across heterogeneous sub-populations by augmenting traditional methods with face recognition techniques.

2. DATA AND METHOD

For our evaluation, we utilized ground truth data from a previous study on global gender disparities in science [3]. It consists of a manually labeled random sample of academics, their full names, institutions, countries, and their gender. The ground truth was created by inspecting CVs, pictures and institutional websites. After removing ambiguous and repetitive names, the final name list consist of 693 male names and 723 female names.

We then evaluate different name-based gender detection methods using the full names of our manually labeled scientists as input. Finally, we propose a new mixed method that combines name-based and image-based gender detection.

2.1 Gender detection methods

In the following, we review some prominent unsupervised approaches that only require a name or picture as an input. These approaches do not require training and are widely used in scientific research as mentioned in the introduction.

Social Security Administration’s baby names data. The US Social Security Administration (SSA) covers registered baby names in the United States since 1880. Many gender detection tools such as the “gender” package in *R*¹ or the OpenGenderTracking² rely on this database.

IPUMS Census data. Integrated Public Use Microdata Series (IPUMS) census data consists of samples of the American population drawn from fifteen federal censuses from the American Community Surveys between 1850 to 2000. This database is also used in the “gender” package in *R* and other web-based name extraction packages³.

Copyright is held by the author/owner(s).
WWW’16 Companion, April 11–15, 2016, Montréal, Québec, Canada.
ACM 978-1-4503-4144-8/16/04.
<http://dx.doi.org/10.1145/2872518.2889385>.

¹<https://cran.r-project.org/web/packages/gender/>

²<http://opengendertracking.github.io/>

³<https://usa.ipums.org/usa-action/>

Table 1: Per-class and overall precision and recall of various gender detection methods. The mixed approach outperforms all other methods by at least 9%.

	SSA	IPUMS	Sexmachine	Genderize	Face++	Mixed1	Mixed2
female precision	0.96	0.96	0.97	0.95	0.92	0.91	0.93
female recall	0.79	0.69	0.77	0.86	0.81	0.95	0.94
female F_1	0.86	0.80	0.85	0.90	0.86	0.93	0.93
male precision	0.98	0.92	0.98	0.98	0.86	0.96	0.98
male recall	0.70	0.68	0.72	0.77	0.85	0.89	0.88
male F_1	0.82	0.78	0.83	0.86	0.85	0.92	0.93
accuracy	0.75	0.68	0.74	0.82	0.83	0.92	0.91

Sexmachine. The list of 40,000 names is primarily collected by Jörg Michael. Because of its availability, several libraries in various languages (see for example C’s (*gender.c*) or Python’s *Sexmachine* library⁴) use this database. Given a name, Sexmachine makes a guess whether the name is male, mostly male, female, mostly female or unclear. The advantage of this name list is that it provides detailed information about how popular a first name is in a country and how strongly it is associated with a given gender. Therefore, it enables the disambiguation of names based on the country of origin. The list also provides information for a variety of countries including China and India.

Genderize. Apart from publicly available name data bases there are numerous commercial applications that incorporate various databases from online resources to assess gender. The problem with commercial applications is the difficulty to determine how the data is gathered and processed. Among commercial detection methods are *Facebook graph API*, *Gender API*, *Namsor* which is based on *Gender API* and *Genderize*. In this work, we analyzed the latter method. *Genderize* utilizes big datasets of information, from user profiles across major social networks and exposes this data through its API. The response includes a confidence value⁵.

Face++. In addition to name-based gender detection methods, face recognition algorithms have become a popular tool for inferring the gender, e.g., for social media users. Among those, image-based application *Face++* seems to provide high performance [7]. This approach requires access to a picture of the person.

In order to derive the gender for a specific scientist, we propose to initially collect the first five *Google thumbnails* using the full name as search query term and then apply image-recognition on the search results. This approach does not necessarily require that the collected pictures depict the scientist we originally searched for, but the idea is that we collect a sample of pictures that depict people who are named like the person we searched for. The advantage of using the full name as input is that for first names that are ambiguous or unisex, the combination of first and last name is often a better indicator of the gender associated with the certain culture.

A Novel Mixed Approach. In addition, we propose mixed methods that combine name-based detection methods with an image-based face recognition approach. We test two variations of this method. In method *Mixed1*, the best name-based approach, namely *Genderize*, is used first. For the remaining unidentified names, the image-based method *Face++*, is used. In method *Mixed2*, *Genderize* and *Face++* have equal weight. For the weighting, we do not use a binary decision for each method, but also take the reported confidence as a numeric value into account. In doing so, this method can handle ambiguous names more efficiently. Note that method *Mixed1* does not require retrieving pictures for the whole population and is therefore more efficient than method *Mixed2*.

⁴<https://pypi.python.org/pypi/SexMachine/>
⁵<https://genderize.io/>

Table 2: Accuracy of various gender detection methods for people from different countries. For most countries mixed approaches perform best.

	# instances	SSA	IPUMS	Sexmachine	Genderize	Face++	Mixed1	Mixed2
United States	419	0.82	0.76	0.84	0.83	0.91	0.91	0.90
China	113	0.20	0.11	0.67	0.28	0.65	0.50	0.56
United Kingdom	96	0.94	0.92	0.92	0.94	0.81	0.98	0.94
Germany	82	0.87	0.88	0.96	0.94	0.87	0.96	0.93
Italy	75	0.93	0.92	0.94	0.98	0.79	0.99	1
Canada	60	0.87	0.77	0.86	0.91	0.90	0.96	0.93
France	58	0.93	0.92	0.80	0.96	0.81	0.97	1
Japan	56	0.79	0.70	1	0.90	0.62	0.91	0.94
Brazil	44	0.29	0.29	0.15	0.44	0.81	0.90	0.93
Spain	39	0.96	0.92	0.92	1	0.92	1	1
Australia	31	0.89	0.89	0.90	0.86	0.86	0.94	0.93
India	29	0.67	0.17	0.71	0.78	0.83	0.83	0.93
South Korea	27	0.04	0.00	0.58	0.11	0.74	0.37	0.66
Switzerland	25	0.78	0.70	0.56	0.83	0.88	0.90	0.92
Turkey	21	0.43	0.14	0.79	0.81	0.86	1	1

3. RESULTS AND DISCUSSION

The results displayed in Table 1 show that among individual methods, image-based *Face++* and *Genderize* perform relatively better than others. However, the overall best results are achieved by the mixed approaches, which outperform all others by at least 8% accuracy. Although all evaluated methods achieve high overall precision, recall rates vary. All gender detection methods show comparable results for both classes (male and female) and therefore no systematic gender-bias can be asserted.

By contrast, Table 2 indicates that the error rates strongly depend on the country of residence of an individual. While name-based approaches work quite well for western industrialized countries, their performance deteriorates for emerging nations such as China, South Korea or Brazil. Clearly, popular names of these countries are not covered sufficiently in the databases at this point in time. For these countries, an image-based approach leads to substantially better results (e.g., for South Korea the accuracy of image-based approaches is at least 16% better than the best name-based method). The *Genderize* method that also harnesses social media performs poorly for China, presumably due to accessibility to the Chinese social networking websites. Our proposed mixed approaches outperform the existing methods for the majority of the countries.

Conclusions: Our results suggest that the performance of name-based gender detection approaches varies according to the country of origin and that performance for emerging nations is particularly weak. Significant enhancements can be achieved by combining name-based with image-based gender detection methods. In the future, our findings could be combined with machine learning approaches to develop better methods for assessing demographic attributes of users on the Web.

References

- [1] M. Ciot, M. Sonderegger, and D. Ruths. Gender inference of twitter users in non-english contexts. In *EMNLP*, pages 1136–1145, 2013.
- [2] C. Green, N. Jegadeesh, and Y. Tang. Gender and job performance: Evidence from wall street. *Financial Analysts Journal*, 65(6):65–78, 2009.
- [3] V. Larivière, C. Ni, Y. Gingras, B. Cronin, and C. R. Sugimoto. Bibliometrics: Global gender disparities in science. *Nature*, 504(7479):211–213, Dec. 2013.
- [4] A. Mislove, S. Lehmann, Y.-Y. Ahn, J.-P. Onnela, and J. N. Rosenquist. Understanding the demographics of twitter users. *ICWSM*, 11:5th, 2011.
- [5] D. Rao, D. Yarowsky, A. Shreevats, and M. Gupta. Classifying latent user attributes in twitter. In *Proceedings of the 2nd International Workshop on Search and Mining User-generated Contents*, SMUC ’10, pages 37–44, New York, NY, USA, 2010. ACM.
- [6] J. D. West, J. Jacquet, M. M. King, S. J. Correll, and C. T. Bergstrom. The role of gender in scholarly authorship. *PLoS ONE*, 8(7):e66212, 07 2013.
- [7] E. Zhou, Z. Cao, and Q. Yin. Naive-deep face recognition: Touching the limit of LFW benchmark or not? *arXiv preprint arXiv:1501.04690*, 2015.