Measuring Search Engine Bias in European Women's Image Results using Machine Learning Algorithms

Barbara Pisker [1], Kristian Dokic [1] and Marko Martinovic [2]

[1] *The Josip Juraj Strossmayer University of Osijek, Faculty of Tourism and Rural Development, Vukovarska 17, Pozega, Croatia*
[2] *University of Slavonski Brod, Technical department, Trg Ivane Brlic Mažuranic 2, Slavonski Brod, Croatia*

Abstract

This paper focuses on the issue of image search engine results, which many authors claim are the result of biases, thereby multiplying those same biases. The Google search engine was analysed, where images of women from nine countries of the European Union were searched, but using three different languages to generate queries. In this way, we tried to compare the prejudices of other language groups reflected in the results obtained using the search engine.
Two thousand seven hundred images of women were collected, and to quantify the results, an artificial intelligence algorithm was used to calculate the probability of nudity in the image. The hypothesis that there is no difference between the perception of women for a particular country by English, Chinese and Russian language users was generally rejected because there are statistically significant differences in 6 out of 9 countries.

Keywords

Nudity score, machine learning, gender bias, search engine

1. Introduction

Search engine bias is a significant concern in today's society. Various studies have been conducted on this topic, and scholars have contributed to understanding this bias's root causes, impacts and mitigation techniques. Search engine image gender bias has significant implications for how we perceive and understand societal inequalities, gender roles, stereotypes, and their perpetuation in a digital society. Over the past decade, researchers have investigated how search engines produce and reproduce gendered images and how these images are used to reinforce established social norms and power relations.

The widespread use of search engines in our daily lives has brought a new dimension to accessing information and images. However, search engine algorithms have been criticised for being biased towards displaying sexualised and objectifying images of women. Sociological authors have contributed to this debate, pointing out how such images can impact women's self-esteem and the perpetuation of gender stereotypes.

This paper will analyse images of women from nine EU countries obtained using the Google search engine but in three different languages. Considering that search engines index and tag different images depending on the language that is next to the image on a web page, it can be concluded that these same images indicate prejudice against women of certain nationalities by the population that uses one of those three languages. English, Chinese, and Russian languages will be used. The collected images will be analysed with an algorithm for detecting nudity, and a quantitative result will be obtained that indicates the differences in prejudices against women of the three mentioned populations.

Proceedings of RTA-CSIT 2023, April 26–27, 2023 Tirana, Albania
EMAIL: bpisker@ftrr.hr (A. 1); kdokic@ftrr.hr (A. 2); marko.martinovic@unisb.hr (A. 3)
ORCID: 0000-0001-9434-5541 (A. 1); 0000-0003-4358-9065 (A. 2); 0000-0003-1839-2471 (A. 3)

The explanation can be expressed as a research question: Is there a difference between the perception of women in a particular country by English, Chinese and Russian language users?

After the introduction, the literature is analysed in the second section. In the third section, the sample, methods, and results of the research are described, while in the fourth section, there is a discussion. The fifth section has an endnote, and below is a list of the literature used.

2. Literature review

One of the most important contributions to this field of research is Safiya Umoja Noble (2018) work exploring how search engine algorithms can perpetuate and reinforce societal biases, particularly concerning race and gender. Their research has proved that search engine algorithms perpetuate gender and racial biases. It has also revealed that search results for particular groups were often linked to negative or stereotypical content, which can reinforce further harmful societal stereotypes [1].

Nobles highlights the issue of the sexualisation of women in search engine image results. One of the key findings is that search engines tend to reinforce and perpetuate gender stereotypes in their image results. When searching for images of women, search engines often prioritise and display sexualised images, reinforcing the idea that women are primarily objects of male desire. This can seriously affect women's self-esteem and comprehension by other, different social groups.

Additionally, search engine algorithms tend to prioritise images shared and liked the most, rather than images most relevant to the search query. This can create a feedback loop where popular images continue to be prioritised, regardless of whether they perpetuate gender stereotypes or not. She argues that search engines must take a more proactive approach to identify and address these algorithms' biases.

Nobles highlight the need for greater awareness and action around the issue of search engine image gender bias. By recognising how search engine algorithms can perpetuate gender stereotypes and sexualisation, we can work towards creating a more equitable and inclusive digital environment. She calls for greater transparency and accountability in the design and implementation of search engine algorithms, as well as increased awareness and education around the issue of search engine image gender bias [2], [1].

Another study that contributed to this field was Latanya Sweeney's research on gender and racial bias in search engines (2013). Sweeney found that search engine autocompletes suggestions for names associated with women were more likely to include negative terms and stereotypes than those associated with men. Additionally, her research found that job ads for male-dominated fields were more likely to display when search engine users entered terms related to the male gender than when they entered terms associated with the female gender [3].

Overall, Sweeney's research highlights the need for greater awareness and accountability regarding the potential for search engines to perpetuate biases and stereotypes related to race and gender. She emphasises the importance of transparent and inclusive algorithmic design and the need for ongoing evaluation of how search engines impact different social groups.

Similarly, Kate Crawford's study (2016) revealed that machine learning algorithms used to train facial recognition systems were often biased towards white people and males. This bias was due to the underrepresentation of women and people of colour in the training data. Crawford argued that these biases must be addressed by increasing diversity in the training data and algorithm development [4].

Kate Crawford has conducted several studies on search engine image gender bias, with some of the key findings as follows:

1. Stereotypical images: Her research has also found that search engines often return stereotypical images of women in specific fields, such as nursing or teaching. This reinforces traditional gender roles and biases.

2. Objectification of women: Crawford's research has revealed that search engines often display objectifying images of women, particularly concerning sexualised keywords. This can contribute to the objectification and sexualisation of women in society.

3. Intersectional biases: Crawford has also highlighted the intersectional nature of search engine bias, where women from marginalised communities, such as women of colour, are particularly likely

to be negatively impacted by search engine image bias.

4. The invisibility of specific groups: such as non-binary individuals and those who do not conform to traditional gender roles, is often rendered invisible in search engine image results, reinforcing societal biases and exclusion [5] [4].

Overall, Crawford's research highlights the need for greater awareness and accountability regarding search engine image gender bias and the importance of inclusive and diverse representation in search engine results.

One of the key themes in the literature on search engine image gender bias is the prevalence of stereotypical and objectifying, even sexualised, images of women. Researchers have found that search engines often prioritise images of women that conform to traditional gender roles, such as images of women in sexualised or domestic contexts. Different research studies revealed that the search results prioritised images of women in submissive, sexualised poses and that the images were often manipulated through editing software to enhance the sexualisation. This tendency to present women as passive and objectified reinforces patriarchal norms and contributes to women's marginalisation, objectification and oppression in society. In addition to perpetuating gender stereotypes and inequalities, search engine image gender bias can have tangible negative impacts on individuals and communities.

As the literature on search engine image gender bias grows, researchers explore various aspects of this issue and propose new approaches to address it. Otterbacher et al. (2018) conducted a study in which they presented participants with image search results for different keywords and asked them to rate the results for gender bias and sexism. They found that participants perceived images of women in sexualised or domestic contexts as more biased and sexist than images of men in similar contexts. This suggests that users are aware of and sensitive to gender bias in search engine image results and that biases may be reinforced by using stereotypical and objectifying images [6].

Another area of research in this field is the development of methods for detecting and diagnosing gender bias in image recognition systems. Schwemmer et al. (2020) developed a framework for analysing gender bias in image recognition systems and applied it to several publicly available methods, finding evidence of bias in all of them [7].

Furtheron, Fabris et al. (2020) explored the gender bias conveyed by ranking algorithms, finding that these algorithms often reinforce gender stereotypes by prioritising images and information that conform to traditional gender roles and excluding those that challenge or subvert those roles [8].

Additionally, Banet-Weiser, S. (2012) argues that search engines reinforce gender stereotypes by displaying sexualised images of women. Emphasis on sexualised images of women can harm women's self-esteem and perpetuate patriarchal attitudes [9]Banet-Weiser's argument is supported by several studies that have shown that exposure to sexualised images of women can lead to negative effects on women's self-esteem and body image [10], [11] Moreover, Banet-Weiser notes that the search engine algorithms are not neutral, but rather reflect the cultural biases and assumptions of the programmers who design them [9].

Rottenberg, C. (2014) argues that the emphasis on sexualised images of women in search engine results reflects a neoliberal feminist approach that values women's sexuality as a form of empowerment. Rottenberg notes that the commodification of women's sexuality has been a central feature of neoliberalism and argues that this is reflected in search engine algorithms. She argues that this approach can harm women, reducing them to objects of desire and reinforcing patriarchal attitudes towards women's bodies. Furthermore, search engines reinforce gender stereotypes by displaying sexualised images of women but also note that feminist activists are using the Internet to challenge these representations. Feminist activists use the Internet to create counterpublics that challenge dominant representations of women's bodies and sexuality [12].

The literature reviewed here shows that there is evidence to support the claim that search engines are biased toward displaying sexualised and objectifying images of women. These images can negatively affect women's self-esteem, perpetuate harmful stereotypes of women, and contribute to the objectification of women's bodies. Moreover, the literature reviewed here shows that the search engine algorithms are not neutral but reflect the cultural biases and assumptions of the programmers who design them. These biases can reinforce patriarchal attitudes towards women's bodies and sexuality and marginalise women of colour and other marginalised groups.

Efforts to address search engine image gender bias have primarily focused on two strategies: algorithmic interventions and community-led initiatives. Algorithmic interventions involve modifying the algorithms used by search engines to produce more diverse and representative image results. For example, a study by Fabrizzi S. et al. (2021 & 2022) and proposed a method for adjusting search engine algorithms to reduce gender and racial bias in image search results [23] Community-led initiatives involve engaging with communities affected by search engine image gender bias to raise awareness and develop strategies for challenging and subverting dominant stereotypes and biases.

Overall, the literature on search engine image gender bias highlights the complex and multifaceted nature of the issue and how search engines can perpetuate and reinforce gender stereotypes, biases, and inequalities. While efforts to address these issues are ongoing, it is clear that more work needs to be done to ensure that search engines reflect the diversity and complexity of human nature and experience, not perpetuating offensive, biased and harmful stereotypes and inequalities.

However, the literature reviewed here also shows that feminist activists use the Internet to challenge these dominant representations of women's bodies and sexuality. By creating counter-publics that challenge these representations, these activists can offer alternative images and narratives that celebrate women's diversity and challenge patriarchal attitudes towards women's bodies.

3. Research

3.1. Sample

The Google image search engine was used in the paper. The countries selected for analysis are the nine European Union countries with the most significant number of inhabitants. These are, in alphabetical order, Belgium, Czechia, France, Greece, Italy, the Netherlands, Poland, Romania and Spain. The image search engine Google indexes web pages on the Internet. It uses several methods and algorithms that are not publicly available to mark images and assign them tags with the most likely content. The first step before collecting the images was to determine whether the images that the Google search engine returns, as a result, depend on the language in which the query is made. Testing in different languages showed that the results depend on the language of the query and that we get different results for the same terms in other languages.

After that, three languages were chosen in which the queries were generated: English, Russian, and Chinese. In this way, we can compare the perception of the populations that use the mentioned languages.

The third step was the definition of the query itself, and two words were chosen, the first word being "woman", while the second word is the name of the country for which we are interested in pictures of women. For example, if we wanted to collect images of women from Spain, the query in English was: woman Spain. In the case of the Russian language, the query was: женщина Испания. The number of images collected was one hundred for each country and each language. The product of the number of images (100), the number of countries (9) and the number of languages (3) is a total of 2700 images. We collected images using a program written in the Python programming language, and the program is available at: https://github.com/kristian1971/RTA-CSIT-2023 .

3.2. Method

Authors mentioned in the literature often state that existing biases will likely be reinforced by transferring them to search engine systems, as it is a kind of feedback loop. Prejudice is initially part of a person's attitude, but publishing that prejudice on a web page on the Internet makes that same prejudice available to search engines. Based on several factors that are used to rank content on the Internet, prejudice becomes part of the results by which that same search engine influences the users' attitudes. People who are significant content creators have a powerful influence on this path, and this includes journalists and editors of portals with a large number of users. The literature also mentions the influence of the search engines themselves, the creators of algorithms for searching, indexing, tagging and ranking results. Figure 1 shows the path by which prejudice can spread from a prominent content creator to a broader population with the help of search engines.

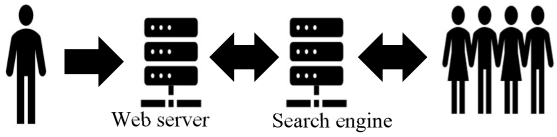

Figure 1: Information flow (authors)

Two thousand seven hundred images of women from 9 nations were collected, and queries were made in three languages. As previously stated, the collected images indicate the attitude of the population in the manner described. By looking at the pictures, you can see the differences between the nationalities of the women, and it is evident that the pictures show women of different ages. In addition, some nations are represented by younger and more freely dressed women than others. To quantify these differences, an algorithm was used to calculate the nudity score of each image. The nudity score results from an algorithm based on a deep neural network, which can range from 0 to 1. The value 0 represents a probability of 0 that a person is without clothes in the image. In contrast, a value of 1 means the same probability that there is a person without clothes in the image. The analysis of the obtained values should reveal the prejudices of the mentioned three groups (English, Chinese and Russian speakers) towards women who belong to the mentioned nine EU countries or ethnicities.

3.3. Nudity score algorithms

There are many artificial intelligence algorithms for obtaining a nudity score. Ananthram et al. analysed ten of the most famous ones for which an Application Programming Interface (API) is available on the Internet. In the title of Ananthram's paper, the author mentioned the abbreviation NSFW for the algorithms. It is a set of algorithms often used by companies to detect content unsuitable for use during working hours, and the acronym comes from "**N**ot **S**afe **F**or **W**ork". The authors stated that NSFW algorithms detect five categories of content, namely:

- Explicit Nudity
- Suggestive Nudity
- Porn/sexual act
- Simulated/Animated porn
- Gore/Violence [13]

For this paper, some listed categories are unimportant but do not affect the result.

There are different approaches to detecting human beings without clothes, and the main goal was to automate this process. Garcia et al. proposed a model based on human skin colour and achieved a precision of 90.33% and an accuracy of 80.23%. They transferred the images to YCbCr space and classified them depending on whether a pixel was skin-coloured [14]. Moreira et al. proposed a new dataset of 376,000 images categorised into pornography and non-pornography. The authors used convolutional neural networks, namely Densenet-121, with a batch size of 128 trained with an SGD optimiser and a learning rate of 2^{-8}. The authors state that they achieved an overall accuracy of 97.1% on the combined datasets, and they consider convolutional neural networks to be the best choice for detecting pornography in images [15]. At the end of the second decade of this century, a whole series of authors started using convolutional neural networks to detect nudity, and this approach is currently considered the best [16] [17] [18] [19] [20].

The paper uses the Nudity Detection algorithm, which is available on the website of DeepAI [21]. The algorithm mentioned above was created by adapting the Open NSFW model that Yahoo presented is based on convolutional neural networks [22]. An API is available for using the algorithm, which can easily automate the calculation process for many images. A simple program that performed this is available at: https://github.com/kristian1971/RTA-CSIT-2023

3.4. Results

Tables 1, 2, and 3 provide descriptive statistical data on the obtained values for the analysed countries and languages. In Table 1 are the mean values of the nudity score, shown in Figure 2.

Table 1

Average nudity score

	Chinese	English	Russian
Belgium	0,0387	0,0573	0,0291
Czechia	0,0250	0,0322	0,0401
France	0,0528	0,0414	0,0281
Greece	0,0811	0,0252	0,0339
Italy	0,0434	0,0405	0,0410
Netherland	0,0368	0,0504	0,0283
Poland	0,0253	0,0287	0,0328
Romania	0,0467	0,0287	0,0104
Spain	0,1216	0,1112	0,0801

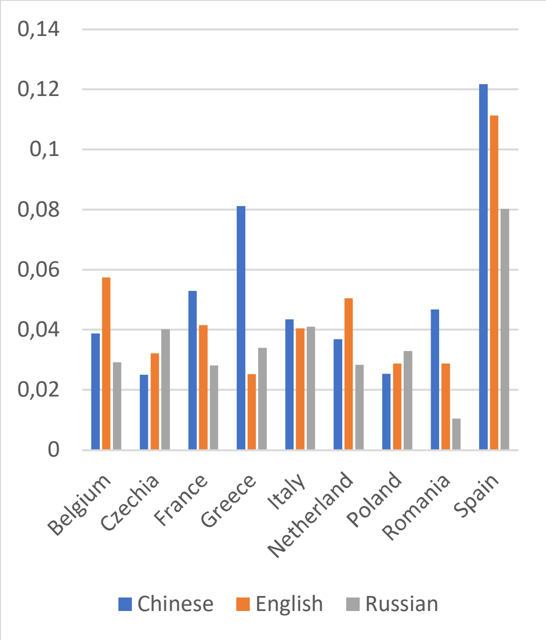

Figure 2: Average nudity score

In Table 2 are the median values of the nudity score, shown in the graph in Figure 3.

Table 2

Median nudity score

	Chinese	English	Russian
Belgium	0,011	0,0084	0,0021
Czechia	0,0038	0,0044	0,0067
France	0,0059	0,006	0,0039
Greece	0,0216	0,0051	0,0082
Italy	0,0106	0,0071	0,0128
Netherland	0,0067	0,0068	0,0017
Poland	0,0057	0,0025	0,006
Romania	0,004	0,0041	0,0018
Spain	0,0378	0,0144	0,0139

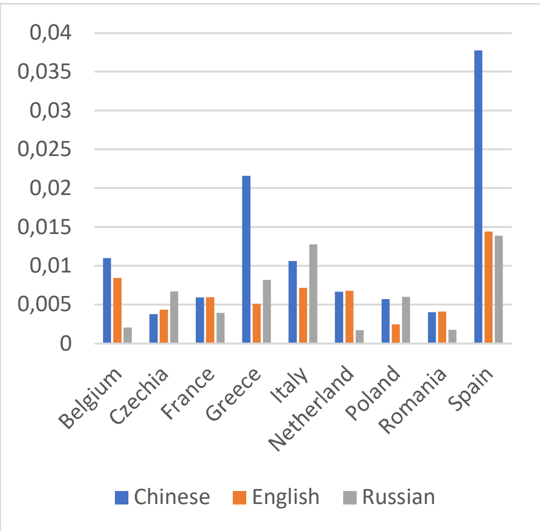

Figure 3: Median nudity score

In Table 3 are the maximum values of the nudity score, which are the same as shown on the graph in Figure 4.

Table 3

Maximum nudity score

	Chinese	English	Russian
Belgium	0,3638	0,4669	0,4566
Czechia	0,38	0,4327	0,6414
France	0,3078	0,8365	0,5745
Greece	0,7974	0,3416	0,603
Italy	0,3232	0,3065	0,308
Netherland	0,4433	0,7559	0,7548
Poland	0,453	0,65	0,4085
Romania	0,461	0,3683	0,2513
Spain	0,803	0,838	0,9376

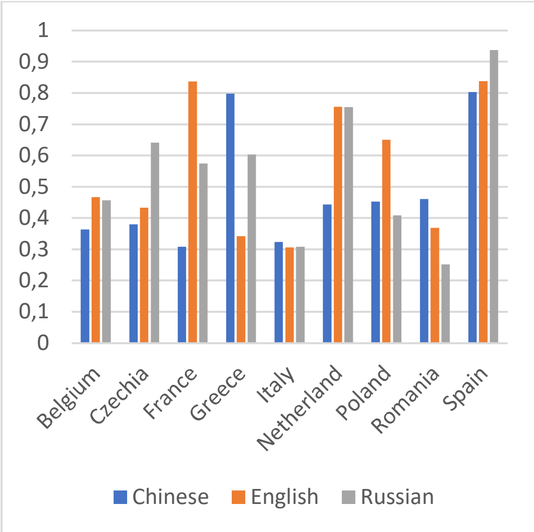

Figure 4: Maximum nudity score

The description of the algorithm states that scores <0.2 indicate that the image is likely to be safe. Scores > 0.8 suggest that the image is highly probable to be NSFW. Scores between 0.2 and 0.8 may be binned for different NSFW levels. Table 4 shows the number of pictures whose nudity score exceeds 0,2, and the same values are shown graphically in Figure 5.

Table 4

Nudity scores below 0,2

	Belgium	Czechia	France	Greece	Italy	Netherland	Poland	Romania	Spain
CHI	4	2	10	9	6	4	4	8	20
ENG	12	6	5	4	3	9	4	6	16
RUS	5	4	3	5	8	5	5	1	10

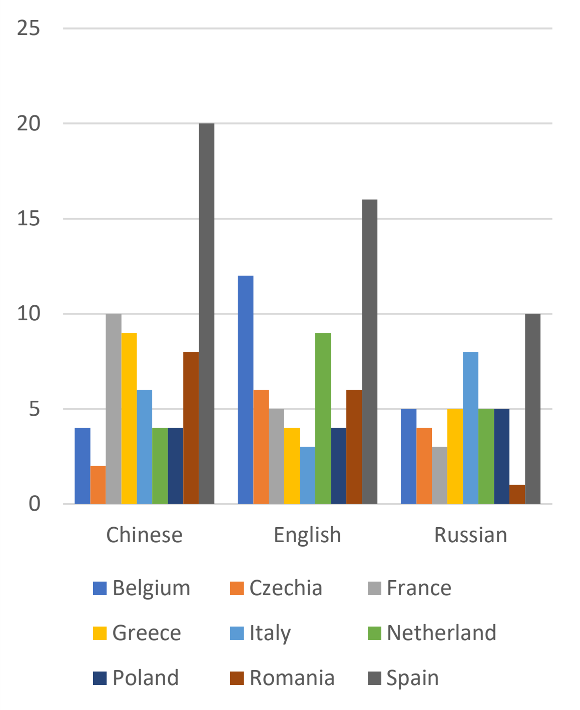

Figure 5: Nudity scores below 0,2

The research question deals with a difference between the perception of women in a particular country by English, Chinese and Russian language users, so the hypothesis is: there is no difference between the perception of women in a specific country by English, Chinese and Russian language users.

Figure 6 shows the distribution of 2700 results obtained by the algorithm, which shows that this is not normally distributed data. For this reason, the non-parametric Kruskal-Wallis test was chosen for testing the hypothesis.

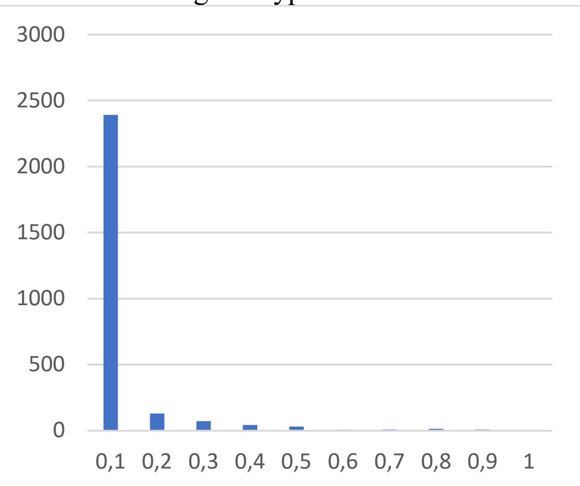

Figure 6: Nudity scores distribution

The SPSS 20 tool was used for statistical analysis. Considering that data was collected for nine countries, nine hypothesis tests were conducted, for each country separately. The test results are visible in Table 5.

Table 5

Kruskal-Wallis test results

	p	Rank CHI	Rank ENG	Rank RUS
Belgium	**0,002**	**168**	**157**	**126**
Czechia	0,810	147	148	155
France	0,234	156	156	138
Greece	**<0.001**	**184**	**124**	**142**
Italy	0,757	151	145	154
Netherland	**0,024**	**155**	**163**	**131**
Poland	**0,004**	**161**	**127**	**163**
Romania	**0,003**	**158**	**166**	**126**
Spain	**0,033**	**168**	**142**	**139**

In the second column of Table 5, there are p-values, while in the third, fourth, and fifth columns are the ranks of the results for each country depending on the language in which the query was made. Rows for which the p-value is less than 0.05 are in bold. Suppose the rank value for a particular language is higher. In that case, it indicates that the nudity score values were also higher, which further suggests that the probability that the pictures show nudity is higher. As a rule, the search engine did not return images with nudity, but the algorithm is sensitive to the very signs of nudity, that is, scantily clad women.

4. Discussion

Analysis of the descriptive statistics of the nudity score for each country and language provides much information. First of all, the results obtained for women from Spain have the highest average nudity score values (0.1216; 0.1112; 0.0801), and the same is true for the median (0.0378; 0.0144; 0.0139) and maximum values (0.803; 0.838; 0.9376). The values that are also visually significantly higher, but only for the Chinese language, are the values of the nudity score for Greece, i.e. images of Greek women. The nudity score anomaly for France (English language) and the Netherlands (English and Russian language) is somewhat noticeable. Whether the nudity score value for the Netherlands is higher because of the liberal attitude towards prostitution is a hypothesis that should be further investigated by analysing the content of the pages from which the images were obtained.

When analysing images for all nine countries, the average nudity score values for Chinese, English, and Russian are 0.0524, 0.0462, and 0.0360.

In the description of the algorithm, it is stated that nudity score values less than 0.2 are probably safe, so Table 4 and Figure 5 shows a total of 100 images per group whose score is greater than 0.2. Considerably higher values for Spain are also visible in that analysis.

Finally, we analyse the hypothesis that there is no difference between the perception of women in a particular country by English, Chinese and Russian language users. Nine tests were conducted, and it is evident that for only three countries, there is no statistically significant difference in the ranks of nudity score values for all three languages. Those countries are Czechia, France and Italy. There is a statistically significant difference between the nudity score ranks for all other countries (Belgium, Greece, Netherlands, Poland, Romania, and Spain). Interestingly, we get the three highest values for the Chinese language (Belgium, Greece, Spain), while for the English language, we get two (Netherlands, Romania), and for the Russian language, only one (Poland).

5. Conclusion

Search engine image gender bias is a complex issue that significantly impacts a contemporary digitalised society. Recent research has contributed substantially to understanding this bias's root causes and impacts. Scholars have found that search engine algorithms perpetuate gender and racial biases, reinforce harmful stereotypes, and limit opportunities for different societal groups. Addressing these biases will require greater awareness and regulation of search engine algorithms and, perhaps even more important, a system of automatic regulation of results that eliminates user-generated biases. One way is to use the algorithm that was used in the paper.

The hypothesised equality between the perception of women for a particular country by English, Chinese and Russian language users was generally rejected, but with an indication that it was rejected for six analysed countries. In comparison, it was not rejected for the three countries. From the descriptive statistics, it can be concluded that the nudity score values are significantly higher for Spain compared to the other eight analysed countries and that queries in the Chinese language usually return images with a higher nudity score. The average value of the nudity score for the Chinese language is the highest when analysing images for all nine countries.

Further research can be extended to all other EU countries and some other countries outside the EU and Europe. In addition, reviewing the images shows that the images with a higher nudity score are generally images of younger women. Further research could use algorithms to analyse the age of women in images, which would give a different perspective on practically the same problem.

6. References

[1] S. U. Noble, »Algorithms of oppression,« u *Algorithms of oppression*, New York University Press, 2018.

[2] S. U. Noble, »Google search: Hyper-visibility as a means of rendering black women and girls invisible,« 2013.

[3] L. Sweeney, »Discrimination in online ad delivery: Google ads, black names and white names, racial discrimination, and click advertising,« *Queue,* svez. 11, p. 10–29, 2013.

[4] K. Crawford, »Can an algorithm be agonistic? Ten scenes from life in calculated publics,« *Science, Technology, & Human Values,* svez. 41, p. 77–92, 2016.

[5] K. Crawford, The atlas of AI: Power, politics, and the planetary costs of artificial intelligence, Yale University Press, 2021.

[6] J. Otterbacher, A. Checco, G. Demartini i P. Clough, »Investigating user perception of gender bias in image search: the role of sexism,« u *The 41st International ACM SIGIR conference on research & development in information retrieval*, 2018.

[7] C. Schwemmer, C. Knight, E. D. Bello-Pardo, S. Oklobdzija, M. Schoonvelde i J. W. Lockhart, »Diagnosing gender bias in image recognition systems,« *Socius,* svez. 6, p. 2378023120967171, 2020.

[8] A. Fabris, A. Purpura, G. Silvello i G. A. Susto, »Gender stereotype reinforcement: Measuring the gender bias conveyed by ranking algorithms,« *Information*

Processing & Management, svez. 57, p. 102377, 2020.

[9] C. Wilkinson, *Banet-Weiser, S.(2012). Authentic™ the Politics of Ambivalence in a Brand Culture,* Sage Publications Sage UK: London, England, 2014.

[10] B. L. Fredrickson i T.-A. Roberts, »Objectification theory: Toward understanding women's lived experiences and mental health risks,« *Psychology of women quarterly,* svez. 21, p. 173–206, 1997.

[11] L. M. Groesz, M. P. Levine i S. K. Murnen, »The effect of experimental presentation of thin media images on body satisfaction: A meta-analytic review,« *International Journal of eating disorders,* svez. 31, p. 1–16, 2002.

[12] C. Rottenberg, »The rise of neoliberal feminism,« *Cultural studies,* svez. 28, p. 418–437, 2014.

[13] A. Ananthram, »Comparison of the best NSFW Image Moderation APIs 2018,« Towards Data Science, 22 November 2018. [Mrežno]. Available: https://towardsdatascience.com/comparison-of-the-best-nsfw-image-moderation-apis-2018-84be8da65303. [Pokušaj pristupa 7 January 2023].

[14] M. B. Garcia, T. F. Revano, B. G. M. Habal, J. O. Contreras i J. B. R. Enriquez, »A pornographic image and video filtering application using optimised nudity recognition and detection algorithm,« u *2018 IEEE 10th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM)*, 2018.

[15] D. C. Moreira, E. T. Pereira i M. Alvarez, »PEDA 376K: A Novel Dataset for Deep-Learning Based Porn-Detectors,« u *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020.

[16] Y. Huang i A. W. K. Kong, »Using a CNN ensemble for detecting pornographic and upskirt images,« u *2016 IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, 2016.

[17] K. Li, J. Xing, B. Li i W. Hu, »Bootstrapping deep feature hierarchy for pornographic image recognition,« u *2016 IEEE International Conference on Image Processing (ICIP)*, 2016.

[18] X. Ou, H. Ling, H. Yu, P. Li, F. Zou i S. Liu, »Adult image and video recognition by a deep multicontext network and fine-to-coarse strategy,« *ACM Transactions on Intelligent Systems and Technology (TIST),* svez. 8, p. 1–25, 2017.

[19] O. Surinta i T. Khamket, » Recognising pornographic images using deep convolutional neural networks,« u *2019 Joint International Conference on Digital Arts, Media and Technology with ECTI Northern Section Conference on Electrical, Electronics, Computer and Telecommunications Engineering (ECTI DAMT-NCON)*, 2019.

[20] K. Zhou, L. Zhuo, Z. Geng, J. Zhang i X. G. Li, »Convolutional neural networks based pornographic image classification,« u *2016 IEEE Second International Conference on Multimedia Big Data (BigMM)*, 2016.

[21] R. Kumar Thakur, »Detect Nudes Using Python Programming and Deep AI,« Medium, 10 April 2022. [Mrežno]. Available: https://medium.com/geekculture/detect-nudes-using-python-programming-and-deep-ai-a9be69b2e9af. [Pokušaj pristupa 4 January 2023].

[22] A. Woodie, »Yahoo Shares Algorithm for Identifying 'NSFW' Images,« datanami, 3 Octobar 2016. [Mrežno]. Available: https://www.datanami.com/2016/10/03/yahoo-shares-algorithm-identifying-nsfw-images/. [Pokušaj pristupa 7 January 2023].

[23] S. Fabbrizzi, S. Papadopoulos, E. Ntoutsi i I. Kompatsiaris, »A survey on bias in visual datasets,« *Computer Vision and Image Understanding,* svez. 223, p. 103552, 2022.