Outfit Cataloguer: A Digital Wardrobe Assistant
(Technical Paper)

Addressing Algorithmic Bias in Artificial Intelligence
(STS Paper)

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By

October 27, 2022

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Rosanne Vrugtman, Computer Science

Bryn E. Seabrook, PhD, Department of Engineering and Society

Introduction

Picture a scenario: you are getting dressed for the day and ask Siri or Alexa what the weather will be. You just used machine learning in three separate ways. Machine learning, or more broadly, artificial intelligence, has applications in every field and is present in many technologies. It helps save time for manufacturers and consumers by giving products the ability to make decisions without needed to be explicitly told what to do. However, sometimes the decisions it makes are biased towards majority groups. This research paper will detail different ways to identify those biases and remove them from machine learning algorithms.

In this scenario, machine learning helped to decide what type of clothes to wear, but there is also a technology that can help in selecting specifically what outfit to choose. Outfit Cataloguer gives users a place to upload and store all of their clothes so that they can build them into outfits. This website has features that suggests clothes that are rarely worn or forgotten about to the user, reducing their need to buy more clothes and the carbon footprint associated with clothing production. Implementing machine learning into these features will allow the product to intelligently compile clothes into outfits, further preventing user waste.

Technical

Clothing disposed of by consumers makes up 16 million tons of waste per year, or 6% of the total municipal waste (Porter, n.d.). While some of this is recycled or sent overseas, almost 90% of it is incinerated or put in landfills, increasing the already huge pollution and waste crisis. This also spills over to other environmental concerns. The processes involved in producing, dying, and finishing fabrics make up 3% of worldwide CO_2 and 20% of the world's water pollution (Igini, 2022). Even just one kilogram of cotton requires 20,000 liters of water to produce. When these effects are multiplied by fast fashion practices, the outcome is devastating.

Fast fashion is “cheaply produced and priced garments that copy the latest catwalk styles and get pumped quickly through stores in order to maximise on current trends” (University of Queensland, 2018). Consumers purchase items matching a quickly expiring trend, wear them a few times, and dispose of them. There is data that reports many garments are only worn 7-10 times before being thrown away (Igini, 2022). These clothes are being tossed before they reach their full lifespan.

These problems can be reduced by consumers being cognizant of what is in their closets and wearing clothing items more than a few times before disposing of them. “I feel like I have nothing to wear” is a sentiment that is echoed often by those with wardrobes full of clothing. If people had a way to conveniently view all of their clothes and experiment different outfits with them, it would break this writer’s block of the fashion world. This is the problem that Outfit Cataloguer seeks to solve. Outfit Cataloguer is a website where users can upload their articles of clothing and build outfits from them. Users can assign different attributes to these items such as brand, color, and material, which can be used later to filter out the item they are looking for. Having all their clothing digitized and catalogued in an organized way can aid in getting a fuller utilization of a person’s clothes, thereby reducing their need to purchase more. An added benefit of Outfit Cataloguer is it cuts down on time deciding what to wear. Having the ability to save outfits means users can filter on their outfits page for something that matches their desired occasion, reuse the same outfit, or switch out a component to give it a new life. Without this resource, people need to start from scratch every time they want to put something together.

Some features in development aim to further reduce the need to buy more clothes. For example, a page is being built that will suggest clothing items that are rarely used in outfits. This serves to remind users about articles they may have forgotten, increasing the perceived size of their closet and the satisfying the need to buy more. A goal of this website is to combine this feature with machine learning so that underused items can be assembled into outfits that are appeasing to the user. This would require large datasets to train the model and likely would need to be trained on outfits from many users, not just the one in question, to make sound decisions. Care will need to be taken in vetting the training data if many users are used to build the model. For example, data should be collected from a diverse set of backgrounds, factoring in age, ethnicity, sexual orientation, gender, and religious background. Analyzing outfit data from a non-diverse set will lead the algorithm to only suggest outfits that are typical of that set of people—marginalizing others who don't fall into that category.

STS Topic

Machine learning (ML) is a subcategory of artificial intelligence (AI), which seeks to create programs that are intelligent, or can mimic human processing (Brown, 2021). Intelligence in this case is defined as the ability to make decisions without explicitly being instructed what to do. The advantage of this is obvious but very powerful: programs can become automated and make fast, smart response to countless situations—even ones the programmers haven't considered. The cost for an intelligent machine like this involves incredibly large sets of data. The more examples a machine learning model is given to learn from, the better decisions it will make when presented with information (Brown, 2021). However, many companies with the resources find it worth the price. In 2021, 56% of companies used AI in some capacity (Chui, 2021). The leading uses in that year were service-operations optimization, AI-based enhancements of products, and contact-center automation, each with about a quarter of the respondents utilizing AI for that purpose. However, artificial intelligence can have applications in every field.

Unfortunately, a large percentage of those companies integrating machine learning into their products do so unaware of the discriminatory biases it may have. If ML models are trained on data that is not representative of the user base, or if it possesses inappropriate content, that can be reflected in the output of the algorithm. These flaws are very hard to detect when the dataset is large, but some companies don't consider looking at all. Depending on the field and application that the algorithm is used in, the consequences of discrimination in ML can be monumental. For example, Equivant's Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), that uses the results from a questionnaire to predict if a criminal defendant will commit another crime (Stewart, 2020). After several controversial predictions were made by the algorithm, ProPublica conducted a study on its risk scores. They found that COMPAS was very inaccurate, correctly predicting recidivism only 20% of the time, and that it was racially biased against black defendants. COMPAS incorrectly predicted black defendants to commit another crime twice as often as it did white defendants, and it labeled white defendants as low risk more than it did black defendants. While this is an extreme example of how discriminatory bias in machine learning can affect people, there are plenty of smaller examples of this in everyday technologies. Facial tracking software that only functions when the face is white ("The Coded Gaze", 2020). Racial bias in determining health care for patients (Baker & Hawn, 2021). Gender bias in translation algorithms (Feldman & Peake, 2021). With such a widespread problem in our midst, it begs the question, can algorithmic bias in AI be eliminated, and if so, how?

To assist in answering this question, it will be analyzed through the lens of Pinch and Bijker's Social Construction of Technology (SCOT). The premise of SCOT is that society and its circumstances determine how technologies are formed (Klein and Kleinman, 2002). SCOT attributes four concepts to this social construction: interpretive flexibility, which states that technology can take different forms depending on the circumstances of its development; relevant social groups, which states that social groups have different interpretations of what a technology should look like; closure and stabilization, which states that design of a technology continues until all relevant social groups agree; and wider context, which states that the small scale social differences between technology designer social groups also plays a role in constructing the technology.

The most important of these for this analysis is closure and stabilization. Algorithmic bias in AI occurs often because the algorithm was verified by social groups that are not representative of the user base for the algorithm. Majority social groups will find nothing wrong with a technology and agree upon its construction, only to find later that it is biased against minority groups. Herein lies a weakness of SCOT; it assumes that all relevant social groups have equal power and are present in the design of the technology. For technologies that service diverse sets of people or the world, this is a difficult feat, but this is one of the ways that discriminatory bias is allowed to arise. The main principle of SCOT of how society shapes technology and its flaw of assuming relevant social groups presence in construction will be used in tandem to analyze how bias can be reduced in artificial intelligence.

Methodology

To further answer the question of how discrimination in AI can be prevented, a combination of documentary research and discourse analysis will be used. The types of documentary resources that will be used include research papers on the how to detect bias in ML algorithms, where algorithmic bias currently exists, and proposals of different methods for preventing or removing bias. The information from these papers will be organized chronologically in the order of implementing a product that integrates AI: first considerations during the design of the product/algorithm, selecting/vetting training data and training the model, verifying the results of the model, deploying the product, and continuing to monitor its behavior. In addition, discourse analysis will be used to incorporate anecdotal evidence of discriminatory algorithmic bias as part of the problem statement. Part of the reason that these biases arise is because of lack minority representation, and this analysis should give those voices a platform to speak and bring attention to those harms.

Conclusion

Outfit Cataloguer is a website and resource that aims to reduce the waste and pollution caused by fast fashion and the clothing industry. Its features make finding something to wear quicker and easier, while helping users to incorporate more of their wardrobe. Machine learning can strengthen the benefits of this tool, but data must be selected carefully to avoid erasing minority groups.

Bias can be produced from technologies implementing AI when the designers of the technology are not representative of the user base. It is this group’s responsibility to find and remove discriminatory tendencies from their product. The research conducted will seek to compile the measures that can be taken to identify and eliminate these biases from AI technologies at every step of the building process, and document the impacts experienced by users of such discriminatory algorithms.

References

Baker, R. S., & Hawn, A. (2021). Algorithmic Bias in Education. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-021-00285-9

Brown, S. (2022, April 21). *Machine learning, explained*. MIT Sloan. https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained

Chui, M. (2021, December 8). *Global AI Survey 2021 – Desktop*. http://ceros.mckinsey.com/global-ai-survey-2020-a-desktop-3-1

Fast fashion quick to cause environmental havoc. (2018, March 21). https://sustainability.uq.edu.au/projects/recycling-and-waste-minimisation/fast-fashion-quick-cause-environmental-havoc

Feldman, T., & Peake, A. (2021). *End-To-End Bias Mitigation: Removing Gender Bias in Deep Learning* (arXiv:2104.02532). arXiv. http://arxiv.org/abs/2104.02532

Igini, M. (2022, August 2). *10 Stunning Fast Fashion Waste Statistics*. Earth.Org. https://earth.org/statistics-about-fast-fashion-waste/

Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values*, *27*(1), 28–52. https://doi.org/10.1177/016224390202700102

Porter, B. (n.d.). *What Really Happens to Unwanted Clothes? | Green America*. Retrieved October 30, 2022, from https://www.greenamerica.org/unraveling-fashion-industry/what-really-happens-unwanted-clothes

Stewart, M. (2020, July 29). *Handling Discriminatory Biases in Data for Machine Learning*. Medium. https://towardsdatascience.com/machine-learning-and-discrimination-2ed1a8b01038

The Coded Gaze: Unpacking Biases in Algorithms That Perpetuate Inequity. (2020, December 16). *The Rockefeller Foundation*. https://www.rockefellerfoundation.org/case-study/unpacking-biases-in-algorithms-that-perpetuate-inequity/