Inclusion Ethics in AI: Use Cases in African Fashion

Christelle Scharff[1], James Brusseau[2], Krishna Mohan Bathula[1], Kaleemunnisa Fnu[1], Samyak Rakesh Meshram[1], Om Gaikhe[1]

[1]Pace University, Seidenberg School of Computer Science and Information System, New York, USA
[2]Pace University, Department of Philosophy and Religious Studies, New York, USA
cscharff@pace.edu, jbrusseau@pace.edu, kbathula@pace.edu, klnu@pace.edu, smeshram@pace.edu, ogaikhe@pace.edu

Abstract

This paper addresses the ethics of inclusion in artificial intelligence in the context of African fashion. Despite the proliferation of fashion-related AI applications and datasets global diversity remains limited, and African fashion is significantly underrepresented. This paper documents two use-cases that enhance AI's inclusivity by incorporating sub-Saharan fashion elements. The first case details the creation of a Senegalese fashion dataset and a model for classifying traditional apparel using transfer learning. The second case investigates African wax textile patterns generated through generative adversarial networks (GANs), specifically StyleGAN architectures, and machine learning diffusion models. Alongside the practical, technological advances, theoretical ethical progress is made in two directions. First, the cases are used to elaborate and define the ethics of inclusion, while also contributing to current debates about how inclusion differs from ethical fairness. Second, the cases engage with the ethical debate on whether AI innovation should be slowed to prevent ethical imbalances or accelerated to solve them.

Introduction

There is a vested interest in applying Artificial Intelligence (AI) to fashion as indicated by the number of available fashion datasets and current range of fashion-related applications. The Fashion-MNIST dataset (Xiao 2017) was built by the Zalando German fashion and lifestyle e-commerce company to benchmark Machine Learning (ML) algorithms. It contains 70,000 grayscale images of fashion products from 10 categories (T-shirt/top, Trouser, Pullover, Dress, Coat, Sandal, Shirt, Sneaker, Bag, Ankle boot). DeepFashion (Liu, 2016) and Amazon Catalog images (for fashion) (Heilbron et al. 2019) are datasets that achieve state-of-the-art performance on fashion items recognition. DeepFashion contains 800,000 images and more than 50 categories. Amazon Catalog contains more more than 7 millions of images. We can however observe that these datasets focus mainly on Western fashion. Global fashion is missing from the existing datasets, and models would not perform illustratively on garments such as Boubou from Senegal, Sari from India or Sampot from Cambodia; boubous, saris and sampots would be most likely classified as “dresses”, ignoring the diversity of global fashion. More inclusive datasets and automated processes are required to categorize, organize, identify, advertise, recommend, and generate global fashion. This research describes two case studies that address diversity in AI through African fashion, which is a $31 billion industry, the largest economic sector of the continent after agriculture according to Euromonitor (CNN 2016).

The empirical work presented in this paper intersects two vigorous theoretical debates in contemporary AI ethics. One concerns the tension between inclusivity and fairness. The ethical principles overlap, but also differ, and the cases of African fashion will allow the principles to be distinguished while also clarifying the case that can be made for the inclusion of data and machine learning focusing on garments outside the scope of today’s mainstream fashion industry AI. The second debate is not about what inclusion is, or whether it should be promoted, but how? Fundamentally, this question concerns the status of technical innovation. On one side, there is the argument that innovation should wait for ethical alignment, while on the other there is the argument that innovation will create alignment. Both sides agree on the following: innovation in AI creates ethical harms as well as benefits. The disagreement is about how to respond effectively. Should innovation be slowed in order to limit the number and severity of the harms and ethical imbalances, or, should technical advance be accelerated to resolve them?

The paper is organized as follows. It begins by describing how data and machine learning techniques were expanded to address some African fashion classification and generative AI problems, and then addresses the ethics of fairness and inclusion, along with the debate between slow AI and acceleration.

Use Case 1: First Senegalese Fashion Apparels Classification

The lack of data is a problem for the development of inclusive AI solutions. Accessing data is difficult in the Global South because of non-existent open data policies and privately-owned data (Heng et al. 2022). This first use case is based on the creation of a small Senegalese fashion dataset and a model capable of classifying various Senegalese apparels, called Boubous and Taille Mames (Seck et all 2022), by using transfer learning (Raina 2007, Shao 2015) for image classification with MobileNetV2 as the base model (Dong 2020). Figure 1 illustrates the Senegalese dataset of 256 images that we created with permission from designers in Senegal. Transfer learning is specifically used to exploit the knowledge acquired from previous tasks (and large datasets) to improve generalization of other tasks with limited available data; it reuses a pre-trained model on new data. This use case raises the issue of the need for more inclusive datasets and proposes a preliminary reflection on addressing global fashion starting with datasets. Dataset creation comes with a high level of responsibility. The World Economic Forum recommends all entities to document the provenance, creation process, and possible uses of the datasets (World Economic Forum 2018). The authors of Datasheets for Datasets emphasize the importance of providing information about the motivation, composition, collection process (preprocessing, cleaning, labeling), recommended uses, distribution, and maintenance of the data (Gebru, 2021). Such requirements have the potential to ensure ML-dataset integrity and increase transparency and accountability within the ML community.

Figure 1: Senegalese Fashion Dataset Snapshot.

Use Case 2: AfricanFashionGAN/DM: Generating African Wax Textile Patterns

Generative AI has mostly focused on generating western-related artifacts, from text to music and images. This second use case explores the generation of African wax designs using generative adversarial networks (GAN) (Goodfellow 2014), StyleGAN2-ADA and StyleGAN3 architectures specifically (Karras 2018), as well as diffusion models (Ho 2020), from a DALL-E curated dataset of 5000 images. Wax are traditional colorful textiles composed of geometrical shapes and cultural symbols worn across Africa. This work studied different approaches of training the models, including hyperparameter tuning and data augmentation, in order to manipulate design aspects like color, pattern, and texture to obtain realistic wax. Figure 2 presents some wax patterns generated by AfricanFashionGAN/DM. This work involved Senegalese fashion professionals to validate the resulting designs and provide feedback for future work iterations. This use case provides wax pattern designers, mostly located in Europe, with a new approach for creating wax print designs, using generative AI but also using subject matter experts from Africa. It also raises the issue of the wide range of application of generative AI outside of the Western context.

Figure 2: Generated Wax Patterns (StyleGAN3, FID: 22.44, 188 hours on a HPC with 2 V100 GPU). Cover.

Inclusion Ethics in Fashion Datasets

The research described in this paper creates an opportunity to carefully define the value of inclusion, and to station it as

a specific ethical justification that is distinct from more common appeals to fairness. The distinction between the two has generated intense discussion in contemporary ethics, as Peters and co-authors (2020) have noted, as well as Jiang and Pardos (2021). The difficulty lies in the fact that inclusiveness is a subtle ethical principle that can overlap with fairness, while sometimes diverging. As a simple example to introduce the difference, we could consider the distribution of access to compute as divided between two researchers. In one case, the resource could be divided into smaller units and distributed sequentially, to one researcher and then the other. This would be fair as each one is treated equally, and it would also be inclusive since the distribution would never allow one researcher to get far ahead of the other. Conversely, the entire resource could be allotted to one or the other researcher by flipping a coin. This would also count as fair because both had an equal opportunity. But, it would not be inclusive since one researcher is manifestly excluded by being left with nothing.

African fashion research allows a more substantial exploration of inclusiveness and its specificity. To begin, the formal ethics of inclusiveness traces back to John Rawls' philosophy as articulated in his *Theory of Justice* (Rawls 1971). There, Rawls proposed a counter to the standard notion of fairness as defined by Aristotle. Rawls argued that an objective person placed in the role of distributing a societal good would, in fact, diverge from the Aristotelian notion of equal apportioning. Instead, for Rawls, a deliberative approach tilts the distribution toward those who have the least. Funding for technological exploration, for instance, may be tilted to favor subjects, schools, or regions with the historically poorest performance, as opposed to being distributed equally across possible recipients. Because Rawls was interested first in political philosophy, the societal goods and examples he envisioned align well with the kinds of benefits governments distribute, including educational opportunities and infrastructure. However, the specific content of his distributions are less important than the logic governing the distributing, and this logic is the core of formal inclusiveness, which has also been named by the terms "solidarity" and more recently as "equity" (Brusseau 2022).

The specifics of the logic's development – which include discussions of Rawls' "original position" and his "veil of ignorance" – stand outside the scope of this paper, but the results distill into what Rawls (1971) called his "difference principle," which stipulates that inequalities in social groupings may be allowed only to the extent that they most benefit the worst-off. In an unequal world, this translates into the general rule that the maximum advantage should go to the least advantaged, and that sets the concept of inclusiveness apart from traditional fairness in the following fundamental way.

Inclusiveness is about outcomes, while fairness is about processes. This explains why a coinflip can fail the inclusion test while remaining perfectly fair: the outcome is totally unequal even while the process is perfectly even-handed. Translating into the example of African fashion and AI research, inclusiveness *begins* by conceiving a desirable outcome, which could well be equal access to AI tools in fashion across the world's regions. Then, that vision subsequently dictates a process. In this example, resources including the creation of, and access to datasets were concentrated on underserved areas of sub-Saharan Africa. It is important to note that the subject of inclusion is not African fashion, it is the underserved which happens to be African fashion today, though it may not remain that way tomorrow. In fact, this potential for any one of us to eventually find ourselves on the underserved side of the balance is a foundational human support for inclusion as an ethical principle.

Contrastingly, fairness means establishing an ethically justifiable process, which in turn justifies the outcome, whatever it may be. From this perspective, it could be argued that current and imbalanced distributions of AI resources across world regions is fair because it responds to an equal economic process: resources tend to cluster where they can be used to generate the most profit, and that clustering is intrinsically even-handed since profit may accumulate as easily in one place as another. Objections could be raised to this formulation, but the significance is the following: it is clear how the idea of inclusiveness diverges. Because of the focus on equality of outcomes, and because those outcomes are currently imbalanced, the ethics of inclusion forms an ethical imperative to channel AI resources toward Africa in ways consonant with the research in this paper.

The Tension Between Precaution and Acceleration in AI Ethics

AI tools in Africa – and elsewhere – have increased efficiency and supported creativity, but they have also aggravated inequalities and left significant populations behind (Gwagwa et al. 2020). Like technologies that came before, AI solves problems but also creates them, which opens a field of exploration dedicated to maximizing the former while minimizing the latter. At least two strategies are being developed (Phillips et al 2024). In sweeping terms, innovation can be slowed until ethical alignment with ideals like inclusion may be assured. Contrastingly, innovation can be accelerated to create alignment with ethical demands: innovation can be leveraged to solve innovation problems. This paper provides a case study to support the second alternative.

The divergence between braking and accelerating can be traced back several decades to a structurally similar dispute in biotechnological healthcare governance, one revolving around the idea of precaution. On one side of the precaution debate, the case was made that biomedical research should be slowed until potential humanist and ethical hazards had been surveyed, fully accounted for, and proactively managed (Martuzzi et al. 2004). The argument on the other side

was that slowing down is itself a hazard because of the benefits lost while innovation is constrained. The consequent recommendation was to forge ahead with biological research, to and respond to risks as they arose alongside progress (Sunstein 2003).

One practical complication is that it will never be possible to completely foresee and then perfectly calibrate the harms that any particular innovation will incidentally create. Additionally, and in the real world, no reasonable person will zoom to one or the other extreme on the precaution-innovation scale: it is untenable to insist on halting all innovation until every possible risk is captured and resolved, just like it would be ruinous to promote unbridled technological development without any consideration for downstream effects. So, the debate about precaution is more about attitudes toward technical innovation and ethics, than it is about determining strict analytical rules for deciding when to slow down and when to speed up. This does not mean the debate is unimportant.

What is certain is that today, the older tension in biological ethics is spilling into artificial intelligence (Fuller et al. 2023). The pole occupied by precaution is converting into something similar to the idea of "Slow AI" (Strickland 2022), and the anti-precautionary stance is transforming toward the idea of acceleration AI ethics (Brusseau 2023b).

Abundant examples of the tension can be found in the area of generative AI image-creation, especially involving deepfake applications (Grossman et al. 2023) and copyright violations (Abbott and Rothman 2022). On the precaution side, proposals have tended toward constricting the technology by limiting or delaying releases to the wider public (Shoaib et al. 2023), and many casual users of platforms including Dall-E have written prompts that elicited a refusal to produce an image, and then a suggestion that the platform's community guidelines may have been violated (Strickland 2022a). The broader strategy is to contain the ethical risk and harm that generative AI can do by limiting what generative AI is allowed to do.

On the acceleration side, the same humanist and ethical challenges are acknowledged, but the solution is not restriction so much as expansion. Instead of constraining generative AI to dimmish the proliferation of harmful deepfake images and copyright violations, the acceleration response is to reformat the technological innovation that produced the problems to provide solutions (Brusseau 2023a). The same deep learning explorations and techniques that initially enabled generative AI deepfakes can be harnessed to detect them automatically (Raza et al. 2022). And, with watermarking and fingerprinting techniques established to preserve the value of intellectual property, copyright protections can be fortified even while image-generation proliferates (Regazzoni et al. 2021). While this is not the forum to fully explore the techniques of acceleration, or to circumscribe the debate dividing precaution against acceleration, what has been made clear is the terms of the disagreement. It is not about whether technology creates ethical challenges, but about how the challenges should be addressed. The cases of AI tools applied to African fashion developed in this paper exemplify the acceleration option.

To begin, acceleration is a response to ethical challenges generated by technological innovation, and the mainstream Fashion-MNIST dataset – along with AI performance surrounding African garments – provides a clear challenge as the test of inclusion fails. Certain geographic regions and practices are not just underrepresented, they are barely represented at all. The question is, how should the imbalance be addressed? For proponents of precaution including Vince Madai (Allahabadi et al. 2022), innovation and AI application should be slowed to limit the number and severity of these kinds of imbalances. Botes (2023) articulates the stance this way: "[Precaution] emphasizes caution, pausing, and review before leaping into new innovations." As applied to AI fashion technology, precaution would have delayed work everywhere in the world once it became clear that optimizing for some garments (mainly belonging to Western cultures) was effectively leaving others out. Then, and only after inclusionary effort were made, would the larger research efforts recommence.

The reticence is reasonable and defensible, but what happened in this paper's use cases illustrates the acceleration alternative: *the same* potential that created the social exclusion problem surges ahead to resolve it. The answer is not to slow innovation but push it ahead in a slightly different direction. For that reason, it is significant that the data techniques and algorithmic tools employed to include sub-Saharan garments in this paper are not radically new or different kinds of AI created to resolve the inclusiveness shortcomings. They are not modifications of AI applications produced *at the cost* of research advances in previously established directions. Instead, the tools are fundamentally identical to those already employed, they are just re-oriented with new data gathered from previously unexplored regions. Consequently, the logic of precaution that potentially sets innovation against ethics is overridden by confidence that innovation may serve ethics. This is the attitude underlying this paper's advances, it is captured by innovations where the same architectures (StyleGAN2-ADA, StyleGAN3) and image classification models (MobileNetV2) that currently leave some global regions behind are amplified to include those garments and traditions. The same innovative machine that caused the inclusion problem, begins to resolve it.

Of course, there is no guarantee of fortuitous outcomes. For example, advances in diagnostic dermatological AI for users with lighter skin have proven stubbornly resistant to darker skin tones (Melarkode 2023): the solution to this insufficient inclusion may require a more precautionary strategy. Regardless, part of what this paper provides is a real case that researchers may use going forward to justify confidence that technical innovation can solve technical innovation's ethical problems. Slowing down for ethics is not

necessary. In some cases at least, more ethics means more innovation.

Conclusion

This research contributes to inclusive AI both technologically and ethically. Technologically, datasets and applications are produced and tested for sub-saharan fashion. Ethically, inclusion is defined and elaborated as humanist support for the technical research. Acceleration ethics is also elaborated and deployed to support the strategy for inclusion implemented in this research.

References

Abbott, R., and Rothman, E. 2022. Disrupting Creativity: Copyright Law in the Age of Generative Artificial Intelligence." Florida Law Review.

Allahabadi, H. et al. 2022. Assessing Trustworthy AI in Times of COVID-19: Deep Learning for Predicting a Multiregional Score Conveying the Degree of Lung Compromise in COVID-19 Patients. In IEEE Transactions on Technology and Society, vol. 3, no. 4, pp. 272-289, Dec. 2022, doi: 10.1109/TTS.2022.3195114.

Botes, M. 2023. Regulating Scientific and Technological Uncertainty: The Precautionary Principle in the Context of Human Genomics and AI. South African Journal of Science, 119(5-6), 1-6. https://dx.doi.org/10.17159/sajs.2023/15037.

Brusseau, J. 2023a "Acceleration AI Ethics and the Debate between Stability AI's Diffusion and OpenAI's Dall-E." In International Conference on Computer Ethics, vol. 1, no. 1. https://soremo.library.iit.edu/index.php/CEPE2023/article/download/293/301.

Brusseau, J. 2023b. How to Accelerate Ethics for Innovation and Against Precaution in Generative AI. International Association of Computing and Philosophy Conference (IACAP), Prague, CZ. https://doi.org/10.5281/zenodo.8024360.

Brusseau, J. 2022. Using Edge Cases to Disentangle Fairness and Solidarity in AI Ethics. AI and Ethics, 2(3), 441-447. https://doi.org/10.1007/s43681-021-00090-z.

CNN. Gbadamosi, N. 2016. The rising stars of a $31 billion industry. Accessed 30 January 2024, https://www.cnn.com/style/article/fashion-cities-africa/index.htm.

Dong, K., et al. 2020. MobileNetV2 Model for Image Classification. In Proceedings of the 2nd International Conference on Information Technology and Computer Application (ITCA).

Fuller, S., Nordmann, A., Coenen, C., Dickel, S., and, M. and Von Schomberg, R. 2023. Proactionary or Precautionary – A Debate1." Hermeneutics, History, and Technology: The Call of the Future.

Gebru, T., et al. 2021. Datasheets for Datasets. Communications of the ACM, vol. 64, no. 12, 2021, pp. 86–92.

Goodfellow, I. J., Mirza, M., Xu, B., Ozair, S., Courville, A., and Bengio, Y. 2014. Generative Adversarial Networks. ArXiv. https://doi.org/10.48550/arXiv.1406.2661.

Grossman, M. R., Grimm, P. W, Brown, D., and Xu, M. 2023. The GPTJudge: Justice in a Generative AI World. Duke Law & Technology Review 23, no. 1.

Gwagwa, A., Kraemer-Mbula, E., Rizk, N., Rutenberg, I., and De Beer, J. (2020). Artificial intelligence (AI) Deployments in Africa: Benefits, Challenges and Policy Dimensions. Afr. J. Inf. Commun. 26, 1–28. doi:10.23962/10539/30361.

Heilbron, F. C. et al. 2019. Clothing Recognition in the Wild using the Amazon Catalog. In IEEE/CVF International Conference on Computer Vision Workshop (ICCVW) (2019): 3145-3148

Heng, S., et al.: Understanding AI Ecosystems in the Global South: The Cases of Senegal and Cambodia. International Journal of Information Management, vol. 64, 2022, p. 102454.

Ho, J., Jain, A., and Abbeel, P. 2020. Denoising diffusion probabilistic models. In Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS'20). Article 574, 6840–6851.

Jiang, W., and Pardos, Z. A. 2021. Towards Equity and Algorithmic Fairness in Student Grade Prediction. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, pp. 608-617. 2021.

Karras, T., Laine, S. and Aila, T. 2018. A Style-Based Generator Architecture for Generative Adversarial Networks. ArXiv. https://doi.org/10.48550/arXiv.1812.04948.

Liu, Z., et al. 2016. Deepfashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR),N.p., 2016.1096-1104.

Martuzzi, M., Joel A. Tickner, and World Health Organization. 2004. The Precautionary Principle: Protecting Public Health, the Environment and the Future of Our Children. World Health Organization. Regional Office for Europe, 2004.

Melarkode, N., Srinivasan, K., Qaisar, S. N., and Plawiak, P. 2023. AI-Powered Diagnosis of Skin Cancer: A Contemporary Review, Open Challenges and Future Research Directions. Cancers 15, no. 4 (2023): 1183.

Peters, D., K. Vold, D. Robinson and R. A. Calvo, 2020, Responsible AI - Two Frameworks for Ethical Design Practice. In IEEE Transactions on Technology and Society, vol. 1, no. 1, pp. 34-47, March 2020, doi: 10.1109/TTS.2020.2974991.

Phillips, C. Jiao, J. and Clubb, E.. 2024. Testing the Capability of AI Art Tools for Urban Design. In IEEE Computer Graphics and Applications.

Raina, R., et al. 2007. Self-Taught Learning: Transfer Learning from Unlabeled Data. Proceedings of the 24th International Conference on Machine Learning - ICML '07.

Rawls, J. 1971. A Theory of Justice. Cambridge: Harvard University Press

Raza, A., Munir, K., and Almutairi, M. 2022. A Novel Deep Learning Approach for Deepfake Image Detection. Applied Sciences 12, no. 19 (2022): 9820.

Regazzoni, F., Palmieri, P., Smailbegovic, F., Cammarota, R., and Polian, I. 2021. Protecting Artificial Intelligence IPs: A Survey of Watermarking and Fingerprinting for Machine Learning. In CAAI Transactions on Intelligence Technology 6, no. 2 (2021): 180-191.

Seck, A.C., Kaleemunnisa, Bathula, K. M., Scharff, C. 2022. Senegalese Fashion Apparels Classification System Using Deep Learning. In: Visvizi, A., Troisi, O., Grimaldi, M. (eds) Research and Innovation Forum 2022. RIIFORUM 2022. Springer Proceedings in Complexity. Springer, Cham. https://doi.org/10.1007/978-3-031-19560-0_60.

Shao, L., et al. 2015. Transfer Learning for Visual Categorization: A Survey. IEEE Transactions on Neural Networks and Learning Systems, vol. 26, no. 5, 2015, pp. 1019–1034.

Shoaib, M. R., Wang, Z., Ahvanooey, M. T., and Zhao, J. 2023. Deepfakes, Misinformation, and Disinformation in the Era of Frontier AI, Generative AI, and Large AI Models. In International Conference on Computer and Applications (ICCA), pp. 1-7.

Strickland, E. 2022. Timnit Gebru Is Building a Slow AI Movement. IEEE Spectrum, 31 Mar. Accessed 30 January 2024, https://spectrum.ieee.org/timnit-gebru-dair-ai-ethics.

Strickland, E. 2022. DALL-E 2's Failures Are the Most Interesting Thing About It. IEEE Spectrum. 14 Jul. https://spectrum.ieee.org/openai-dall-e-2.

Sunstein, C. R. 2003. Beyond the Precautionary Principle. University of Pennsylvania Law Review 151, no. 3 (2003): 1003–58. https://doi.org/10.2307/3312884.

World Economic Forum. How to Prevent Discriminatory Outcomes in Machine Learning. Accessed 30 January 2024: https://www3.weforum.org/docs/WEF_40065_White_Paper_How_to_Prevent _Dis- criminatory_Outcomes_in_Machine_Learning.pdf.

Xiao, H., et al. 2017. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. ArXiv.org, 15 Sept. 2017.30