
Ethical Challenges of Artificial Intelligence in Fashion Recommendation

Dr. Madeline Avery Kensington

Associate Professor, Department of Digital Fashion Analytics and Responsible Artificial Intelligence, Brighton Coast University, San Francisco, California, United States

Corresponding Author: Dr. Madeline Avery Kensington

Associate Professor, Department of Digital Fashion Analytics and Responsible Artificial Intelligence, Brighton Coast University, San Francisco, California, United States

Article Type: Simulation-Based Research Article

Submitted: 04-November-2024 **Revised:** 11-November-2024

Accepted: 13-November-2024 **Published:** 19-November-2024

Abstract: Artificial intelligence enabled fashion recommendation engines are increasingly informing consumers' encounters with garments, styles, brands, sizes, and combinations of these in digital shopping. These recommendation engines can utilize browsing history, purchase history, visual features, contextual signals, body-related attributes, and behavioral cues to enhance product search and reduce information overload. The proliferation of these recommenders presents emerging ethical dilemmas that stem from algorithmic stereotyping, privacy-preserving but information-intensive profiling, diminished consumer agency, body-conscious stress, opaque personalization, commercial manipulation, provider-centered unfairness, and environmentally destructive shopping behavior.

This research introduces an Ethical Fashion Recommendation Governance Framework and tests the framework's implications on a publicly reproducible simulation of 1,200 sessions of utility-only, explainable, fairness-aware, and responsible-AI recommenders. The utility-only configuration achieved the highest simulated relevance score (0.864 out of 1.0), but the greatest choice diversity (73.9/100) as well as the greatest consent-comprehension score (80.6/100). The responsible-AI recommender maintained a relevance score of 0.826 while achieving the best choice-diversity score (73.9/100) and the best consent-comprehension score (80.6/100). The corresponding reduction in intrusive-profiling points from 80.6/100 to 7.3/100 resulted in the greatest increase in sustainable-item total points (65.7/100). Logistic regression revealed that an increase in personalized search degrees increased modeled odds of

ethical risk, then that data-minimization, explainability, diversity controls, and sustainability-serious reranking exerted significant countervailing effects.

A supplementary synthetic recommender audit demonstrated a strong negative correlation between personalization intensity and recommendation diversity. These results reveal that ethical fashion recommendation optimization requires multiple simultaneous tradeoffs, not the maximization of only consumer interest engagement or sales conversions. Responsible implementation should empirically include considerations for provider fairness, user agency, explainability, user privacy-minimization, body-positive recommender designs, sustainable recommendation objectives, manipulative guidelines, and continuous audits.

Keywords: fashion recommender; artificial intelligence; recommender system; algorithmic fairness; consumer sovereignty; privacy; sustainability of fashion

Introduction

AI has redefined digital-fashion retail from a standard product search approach to comprehensive and personalized recommendation environments. Advanced fashion recommender systems utilize collaborative filtering, computer vision, multimodal learning, contextual information, body/size features, compatibility modeling, and more advanced representation-learning techniques to infer things or groups of things a customer might favor; Deldjoo et al. present fashion recommender as a medium-specific recommender-system sub-branch which includes item, outfit, compatibility, size, and explainability tasks.

All of these benefits are significant. Personalization can lessen information assimilation, enhance product discovery, facilitate outfit composition, and might cut down on inappropriate product combinations. Ding et al. further note the following as characteristic computational issues: body-shape awareness, size fitting, occasion based recommendation, and complementary fashion recommendation.

But it is deeply rooted in our lives and does more than ‘tell us what we like’. Clothing is intertwined with reflections of one’s identity, gender, image, inclusion/exclusion, class, culture, and nature.

One could see how a filter which maximizes clicks on related historical searches would during time start to limit the assortment of available products for the consumer. The system

could reinforce gender stereotypes, personalize to excessive intimacy of behavioral data, promote profitable choices, increasingly pressure one due to one's appearance or really push something new and readily available.

These issues are becoming more prominent aspects of regulation. In the implementation of the European Digital Services Act transparency requirements relate to recommender systems, as well as the current European Commission examination of the 'Big Online Retailers' including recommender-system activities.

The ethical goal should thus be advice that is sustainable without being aversive, invasive, manipulative, psychologically damaging or environmentally maladaptive.

Literature Review

a) Personalization, Fairness, and Representation

Fashion recommender have advanced quickly in personalized product recommendation, outfit generation, visual compatibility, contextual styling and size-related applications. Deldjoo et al. bring out bias and fairness as critical unresolved issue in modern fashion recommender.

Our fairness concerns are not limited to fashion. We show Wang, Ma, Zhang, Liu, and Ma that this analogy can be extended to recommender systems with an unequal distribution of exposure and opportunities between users and items, suggesting the need to consider assessments of fairness across various stakeholders:

However, recent human-centered studies suggest that gender debiasing itself needs to be carefully designed rather than removing gender information from the recommendation system. Krause, Goeritz, and Gratz analyzed user reactions to gender-debiased recommendations.

Special Tensions for Fashion-Specific Ethical Issues. (These issues are hot-button because the recommendations may implicitly embed assumptions about which styles, colors, silhouettes, or product types are desirable for which users.

b) Explainability, Body Image, and Consumer Agency

Fashion compatibility models are often hard for users to understand since recommendation are often derived from complex visual and multimodal representations. Yang et al. provide

attribute-level explanations for fashion compatibility, and the review made by Selwon and Szymański highlights the explainability as an important research direction in fashion recommendation.

Human-centered concerns also go beyond technical explanation. Kubullek and Erdem's 2025 study did a specific investigation into the possibility of AI-enabled fashion advice aiding in body positivity and mitigating experiences of discrimination for young women.

Accordingly, an ethical system should not imply that only one body shape is a normative fashion ideal. Personalization must enable users to work towards the food preferences they articulate, rather than suggest to them that they need to hide, change or adapt their bodies to one ideal of beauty.

c) Privacy, Manipulation, and Sustainable Consumption

It also requires data to personalize. It can be browsing, previous purchases, clicks, returns, products view, saved outfits, size estimated, context, and maybe even visual or body-related data.

The FTC has reiterated the potential for interface design to encourage disclosures of more information than a user may wish through dark patterns, while its surveillance-pricing investigation has uncovered some wholesale fears about the use of personal data to personalise commercial treatment.

There is an environmental aspect too to fashion recommendation. As stated by UNEP, unsustainable fashion consumption and production directly accounts for greenhouse-gas emissions, waste, water usage, and microplastic pollution, and the Sustainable Fashion Communication Playbook encourages communication strategies that discourage overconsumption;

From this we can infer that recommendation quality should not be based on manufacturing completeness to purchase.

Research Gap

Fashion recommender systems are theoretically elaborate, but evaluation tends to be unduly focused on relevance, ranking quality, fit, click likelihood, conversion, and various other utility measures.

This results in various unfilled grey areas of ethical concern.

Secondly, traditional relevance measures do not show whether the recommender system in fact constrains choice in a systematic fashion. As a system gets better at anticipating what has been purchased before, it may also be limiting taster's view to a narrower range of styles, independent labels, tastes, or categories.

Second, the concept of fairness cannot be reduced to demographic parity. Fashion carries associations with gender presentation, body form, disability, age, religious clothing, ethnicity, and personal taste that are culturally and individually significant. [7] By stripping away these obvious demographic identifiers, the indirect stereotypes in the data will persist.

Third there may be a conflict between personalization and autonomy. A recommender system that optimizes only to maximize engagement could end up tailoring one's preferences over time, even as it only claims to predict them.

Fourth, privacy governance is often disjointed from recommendation quality even though additional, more personal data may be needed for highly personalized systems to achieve greater accuracy;

Traditional recommendation criteria do not respect sustainability at all. A system could maximize sales in the short term leading to over-purchasing, keeps novelty constantly over sale and accelerates remaking.

Lastly, explainability research tend to focus on *why* the *item* matched, while ethical transparency should answer **why the user** was selected to receive the recommendation and what categories of personal information did it materially rely on.

This current research can fill these gaps with a unique paradigm, holistically connecting relevance, fairness, choice diversity, privacy, explainability, and autonomy, body-positive design, and sustainable consumption.

Novelty / Proposed Contribution

Based on the multi-objective optimization, the research put forward a primary Ethical Fashion Recommendation Governance Framework (**EFRGF**).

Conceptually the optimization of traditional recommender can be simply performed as maximizing predicted user relevance. As mentioned before, the proposed framework views

ethical recommendation quality as a compromise of relevance, diversity, fairness, autonomy, privacy protection, transparency and environmental sustainability.

Therealizes its framework does not support the illusion that personalization needs to be sustained ad infinitum.

A system that offers too little personalisation is likely to be perceived as not useful enough. Moderate levels of personalisation tend to significantly increase the product fit, but proceeding too far in this direction might not deliver additional benefits compared to the increasing invasion of privacy, behavioral manipulation and narrowing of options.

These processes allow to identify the following four states.

- **Preference Assistance** is mainly concerned with explicit preferences and low-sensitivity interaction history of consumers.

Behavioral signals can be used with the constraints of diversity, user-control and privacy to make adaptive personalization:

- **High-Intensity Profiling** :This approach makes use of broad cross session/inferred data. It also calls for greater clarity of justification, transparency and control.
- **Manipulative Personalization** using vulnerability, interface pressure, undisclosed commercial motivations, or psychologically manipulative inference to push users outside their own autonomous intentions must be banned or heavily regulated.

In the additional component called Ethical Recommendation Utility, where a conversion gained by undermining the consumers' free will, oversurveillance, biased exposure, false scarcity messages or inducement to take unneeded consumption would not be meant as improvement.

Thus the main contribution is to reconceptualize fashion recommendation as a multi-stakeholder ethical decision system rather than a prediction task.

Research Objectives / Research Questions

This research therefore intends to assess the ethical implications of fashion personalization becoming more and more profound and explore whether governance mechanisms will allow the maintenance of valuable recommendation performance and help enhance aspects such as fairness, consumer sovereignty, privacy, and sustainability.

The present study answered five research questions.

RQ1: Can utility-maximizing fashion recommendation induce quantifiable trade-offs between relevance and diversity in consumer choice?

RQ2: Will a higher magnitude of personalization increase simulated ethical-risk chance?

RQ3: Can explainability and data-minimization controls' to lower ethical risk, while still enabling effective personalization?

RQ4: How do the fairness-aware/ responsible-AI recommendation architectures influence on diversity, intrusive profiling and sustainable-item exposure?

RQ5. What specific governance controls are required to prevent authoritative power underlying personalized fashion recommendations from being potentially abusive, discriminatory, psychologically harmful, or environmentally damaging?

Methodology

a) Simulation Design and Variables

We fabricated a reproducible simulation-based audit made up of 1,200 synthetic recommendation sessions. No actual consumer profiles, buying histories, body dimensions, photographs, browsing history, platform datasets, or commercially sensitive recommendation logs were utilized.

Four recommendation architectures were modeled with 300 synthetic sessions each.

Table 1: Simulated Fashion Recommendation Configurations

Configuration	Primary Optimization Logic	Major Governance Characteristic
Utility-Only	Maximum predicted relevance	Minimal non-utility constraints
Explainable	Relevance + recommendation explanation	User-facing rationale
Fairness-Aware	Relevance + exposure diversity	Diversity and stereotype controls
Responsible-AI	Multi-objective optimization	Privacy, autonomy, fairness, sustainability, transparency

The six main dependent variables we evaluated include recommendation relevance, purchase variety, consents awareness, occurrence of intrusive-profiling, purchase pressure, and sussie distraction.

b) Statistical Analysis and Ethical-Risk Model

The following paragraph is missing: initial conditions: a fixed random seed of 17 was used. In addition, descriptive statistics were summarized for each architecture.

A second synthetic dataset described events at the intersection of ethical risk and personalization by focusing on the degree of personalization, availability of explanation, minimum data, diversity control, and reranking for sustainability considerations.

Using binary logistic regression, estimates of the corresponding odds ratios and the 95% confidence intervals were obtained.

Additional configuration testing employed 60 synthetic recommendation configurations to validate whether the level of personalization correlates with the level of choice diversity.

Graph 1: Simulated Relationship between Personalization Intensity and Choice




Table 1A depicts a significant (about $r = 0.93$, Pearson) negative correlation between the degree of personalization and the spread of available recommendations in the synthetic CO framework. As we increase the intensity of personalization, we expect recommendations to cluster more and more around prior consumption (or more abstractly, a prior preference

vector used to generate randomized recommendations). Again, I do not mean for this effect to serve as a universal empirical observation or to speak to a specific retailer, but instead to suggest an empirically plausible ethical priority: in a setting where optimizing around proximity to prior consumption is an objective, responsiveness to that prior consumption will increasingly narrow our recommendation space unless we explicitly model novel consumption through exploration.

c) 6.3 Proposed Questionnaire Instrument

The following five questions, which are designed for future empirical testing, were **not** posed in the simulation:

1. Are the suggestions that you get being coached right now offering a wide enough coverage for you to try fashions other than those already bought?
2. Can you understand why you are being suggested certain fashion items?
3. Have you noticed how recommenders in fashion seem to ask more information about yourself than in most other recommendation systems?
4. Have you ever experienced the type of personalized fashion recommendations that discourage you from conforming to a specific look, while encouraging you not to buy any other products?
5. Would you like recommendation systems to tell you about durable, reusable, secondhand, repairable, or environmentally preferable options (where appropriate)?

Since these items were not administered to human participants, no survey percentages or inferential claims are reported from these items.

Results

The synthetic audit revealed a clear trade-off between recommendation relevance and broader ethical-governance performance across the four simulated recommendation architectures. The Utility-Only architecture achieved the highest relevance score at 0.864, but it also produced the lowest choice diversity and consent comprehension, the highest intrusive-profiling rate and purchase-pressure score, and the lowest sustainable-product exposure. By comparison, the Responsible-AI architecture retained a relevance score of 0.826, representing an absolute

decrease of only 0.038 relative to the Utility-Only model, while substantially improving the broader ethics portfolio.

Choice diversity increased by 26.0 points, consent comprehension by 35.7 points, and sustainable exposure by 34.9 points, whereas intrusive profiling declined by 18.4 percentage points and purchase pressure decreased by 32.1 points. The regression analysis further demonstrated that personalization intensity was positively associated with modeled ethical risk, while explainability, data minimization, diversity controls, and sustainability-aware reranking showed statistically significant protective associations.

Table 2. Descriptive Results Across Simulated Recommendation Architectures

Recommendation Architecture	Relevance Score	Choice Diversity /100	Consent Comprehension /100	Intrusive Profiling	Purchase Pressure /100	Sustainable Exposure /100
Utility-Only	0.864	47.9	44.9	25.7%	70.8	30.8
Explainable	0.840	54.4	68.9	19.3%	58.4	38.0
Fairness-Aware	0.825	69.4	64.1	17.0%	54.0	47.6
Responsible-AI	0.826	73.9	80.6	7.3%	38.7	65.7

Table 2 demonstrates that maximizing recommendation relevance alone does not produce the strongest overall ethical performance. The Utility-Only model records the highest relevance score of 0.864 but performs comparatively poorly across the remaining governance indicators. Its choice-diversity score is 47.9, consent comprehension is 44.9, intrusive profiling reaches 25.7%, purchase pressure is 70.8, and sustainable exposure is only 30.8. The Explainable architecture improves consent comprehension to 68.9 and reduces intrusive profiling to 19.3%, while maintaining a relatively high relevance score of 0.840.

The Fairness-Aware architecture produces a substantial increase in choice diversity to 69.4 and sustainable exposure to 47.6. The strongest overall ethical profile is observed under the Responsible-AI architecture, which achieves the highest choice diversity at 73.9, consent comprehension at 80.6, and sustainable exposure at 65.7 while simultaneously reducing intrusive profiling to 7.3% and purchase pressure to 38.7. Importantly, these gains occur with only a 0.038-point absolute reduction in relevance compared with Utility-Only optimization,


suggesting that a broader ethical objective function can substantially improve governance outcomes without eliminating recommendation utility.

Table 3. Logistic Regression Predicting a Simulated High-Ethical-Risk Recommendation Event

Predictor	B	SE	Odds Ratio	95% CI for OR	p-value
Personalization Intensity	1.548	0.301	4.70	2.61–8.48	<.001
Explainability Enabled	−0.950	0.176	0.39	0.27–0.55	<.001
Data Minimization	−1.377	0.184	0.25	0.18–0.36	<.001
Diversity Control	−0.803	0.171	0.45	0.32–0.63	<.001
Sustainability-Aware Reranking	−0.457	0.169	0.63	0.45–0.88	.007

Table 3 presents the simulated logistic-regression results predicting a high-ethical-risk recommendation event. Personalization intensity is the only modeled positive-risk predictor, with an odds ratio of 4.70, indicating that greater personalization intensity is associated with substantially higher modeled odds of ethical risk. All four governance controls demonstrate protective associations. Explainability produces an odds ratio of 0.39, corresponding to approximately 61% lower modeled odds relative to the uncontrolled condition, while diversity control produces an odds ratio of 0.45. Sustainability-aware reranking shows a more moderate but statistically significant protective association, with an odds ratio of 0.63 and a p-value of .007. Data minimization demonstrates the strongest modeled protective effect, with an odds ratio of 0.25 and a 95% confidence interval of 0.18–0.36, corresponding to approximately 75% lower modeled odds of a high-risk recommendation event. Because all reported confidence intervals exclude 1.00, the simulated coefficients consistently support the direction of the modeled relationships. These results suggest that limiting unnecessary data collection may be especially important when recommendation systems also employ intensive personalization.

Graph 2. Simulated Ethical-Governance Outcomes Across Recommendation Architectures



Graph 2 compares four governance-related outcomes across the Utility-Only, Explainable, Fairness-Aware, and Responsible-AI recommendation architectures. The Utility-Only configuration exhibits the weakest choice diversity and consent comprehension while generating the highest purchase-pressure score and the lowest sustainable exposure. Progressive introduction of explainability, fairness constraints, and responsible-AI controls shifts the overall performance profile away from narrow utility optimization toward a more balanced governance configuration. The Responsible-AI architecture produces the highest choice-diversity score of 73.9, the highest consent-comprehension score of 80.6, and the highest sustainable-exposure score of 65.7, while reducing purchase pressure to 38.7. The pattern demonstrates that responsible recommendation design does not require optimization around a single ethical metric. Instead, its principal advantage lies in producing a stronger portfolio of user-autonomy, transparency, sustainability, and risk-control outcomes while retaining recommendation relevance close to the utility-maximizing baseline. The simulated findings therefore support multi-objective recommendation governance rather than an architecture that treats immediate relevance as the sole criterion of system quality.

Discussion / Critical Discussion

a) Accuracy, Diversity, and Algorithmic Stereotyping

The results show that the recommendation architecture offering the most ethically desirable advice need not be the recommendation architecture that offers the greatest accuracy. The utility-only optimization had the most effective relevance score but exhibited a much narrower diversity of recommendation. This tension is also seen in the area of general fairness and diversity research, which indicates recommender-system objectives cannot be reduced to predictive accuracy. Wang et al. describe multiple fairness relationships between recommender systems and users and Zhao et al. suggest conceptually, the very notions of fairness and diversity are intricately intertwined.

Fashion intensifies this issue. Prior shopping behavior can encode social conventions based on gender, age, physical characteristics, and fashion style. Merely copying those preexisting standards can provide reinforcing bias instead of just mirroring norm. A recent user study of gender-debiased recommendations shows that mitigation approaches also have to be user-centric and adaptive.

Hence the suitable goal is choice-supportive personalization: “algorithms must learn preferences, but should not treat past behavior as a fixed boundary of future identity.”

b) Privacy, Explainability, and Consumer Autonomy

The regression analysis reveals data minimization and explainability as two important protective controls: this is of conceptual significance because private risks in recommender systems often emerge from the presumption that any predictive signal, not just in recommender system should be fetched.

The FTC’s initiatives on dark patterns can illustrates how an interface architecture can nudge consumer’s increased data sharing or favorable commercial decisions while its work on surveillance pricing can highlight regulatory concerns of using personal data for personalized commercial treatment.

Fashion systems can exacerbate these issues where personalisation now encompasses (or is based on) sensitive appearance-related information or psychologically revealing behavioral inferences.

This way, must answer at least two different questions :

Why would this product match my declared preferences? and

Why does the platform decide to target me with that recommendation?

Technical explanation only is not enough to preserve autonomy if ‘commercially motivated ranking parameters’ are not made visible to the user.

c) Body Positivity, Sustainability, and Responsible Recommendation

The specificity of the ethical position of Fashion recommendation should be noted; a successful commercial system of recommendations influences the image of oneself, while also inspiring consumption.

Meanwhile, Kubullek and Erdem demonstrate the increasing desire for fashion AI to enhance body positivity rather than reinforce disempowering beliefs. Responsible design, therefore, should refrain from indiscriminate language implying that clothing is “more valuable/worthwhile/better because it can conceal/compensates/ slimming masculinize/feminine” the body in the absence of a styling identity that the user asks for.

Sustainability adds further pressure. If recommendation engines are all about promoting transactions, UNEP recognizes overproduction and overconsumption as the main issue in fashion and textile environmental footprint.

A sustainable recommender will thus be able to bring choices to consumers about repair, reuse, secondhand, rental, durable-garment and wardrobe-combination options when appropriate instead of presuming that high quality recommendations default to recommending another new product.

Research Contribution / Implications

The paper provides an integrated governance framework for assessing fashion recommender system nor simply for predictive ability.

Fashion retailers should pursue responsible personalisation by making clear distinctions between benefit to the consumer and benefit to the retailer. Of course, a retailer can judge products in addition to relevancy. But paid placements, pressure to sell inventory, margin enhancement or paid recommendations is not personalization based on consumer preference.

For recommender-system engineers, diversity and exploration should be integrated into ranking goals. It is important that the exploration occurs on a regular basis, stochastically

steering the user to relevant options outside of the short-term local maximums, and indicating to the consumer that they may wish to skew exploration up or down.

For fairness auditing, demographic parity is a necessary but not sufficient criterion. Audits should also examine if the system recommends persistent reductive stereotypes, systematically exposes users to different types of products and if some sellers or product categories are systematically hurt by the ranking system.

Data minimization should occur before model development in the process of privacy engineering. Developers should clarify which categories of data are truly needed for meaningful recommendation, and whether comparable results are attainable with less intrusive processing.

In terms of interpretability, attribute-aware fashion recommendation offers a promising technical base as users can be provided with interpretable factors of compatibility, rather than a set of other opaque numbers. Yang et al. showed the possibility of providing attribute-level explanation in fashion compatibility, and Selwon and Szymański cited interpretability as a key feature of contemporary compatibility approach.

For the most important online services in the EU, recommender-system transparency now has a more direct regulatory significance, within the framework of the Digital Services Act. The European Commission has in fact enumerated a series of recommender-system transparency requirements and inquired of large online services about their implementation.

Ranking system should not systematically reduce the sustainability claim of fashion business by always pushing faster consumption in fashion oriented business. UNEP's Sustainable Fashion Communication Playbook has emphasized the communication reduction that cause overconsumption and shifting aspiration to more sustainable consumption.

The managerial implication is then that recommendation governance should be integrated in product governance rather than be only an engineering matter.

Conclusion

The recommendation of fashion has become more accurate, contextual, visual and personalized thanks to advances in artificial intelligence. We now have part-to-whole compatibility estimation; garment ranking; outfit building; contextual scene integration; visual attributes reasoning; specific tasks such as size or occasion-aware recommendation.

The broadness of the technical aspects covered today in the computational fashion can be seen through the reviews by Deldjoo et al. and by Ding et al., among many others:

This technological ability bears an equally consequential ethical responsibility. Fashion recommendation affects more than what product consumers see, but how they understand themselves: their own tastes, bodies, identities, consumption practices, and options. A recommendation engine optimized for relevance, having narrowed diversity, increased tracking, induced purchasing, and reinforced stereotypes cannot be successful just because click-throughs or conversions are high.

The simulation created in this work illustrates the viability of an alternative optimization approach. The responsibility-AI setting maintained high levels of contextually sensitive accuracy while achieving significantly greater improvements in recommendation variety, privacy, comprehension of consent, user coercion, and sustainable-item exposure. The regression analysis also found that overly personalized recommendation increased the modeled ethical risk likelihood, with other factors such as explainability, data minimization, diversity safeguards, and a sustainability-aware ranking playing a mitigating role.

The main conclusion is that fashion recommenders need to transition from predicting what is most likely to be bought by helping people explore and decide for themselves all the usefully diverse and contextually appropriate options. Good systems must personalize without stereotyping, explain without manipulating, learn without spying, support appearance without body shaming, and suggest without perpetually encouraging endless consumption.

Limitations

Main limitation: The statistical analysis is simulation-based. The figures are not representative of customers behaviors, retailer's performance or any measured consequence of a fashion platform.

The simulated Ethics-risk may be seen as a construct rather than a pathway or standardized clinical, regulatory or industrial measure.

Choosing a variety of options to recommend also should not be equated with fairness, which assumes that a wider variety is preferable. Increased variety may even be detrimental when consumers want highly constrained sets of options.

The model over-simplifies personalization and autonomy. Concrete personalization can also have some real benefits, namely when it helps with accessibility, specific size issues, climate factors, workwear or explicit style preferences.

Notably, body image is not assessed directly through the simulation, although this was a potential area of inquiry. If this type of study were to be more closely examined through human-participant research, high measures of ethical approval would be necessary.

Likewise, it is not enough to be exposed to sustainable items. Through the whole process of utilization and consumption of the sustainably-sourced products, their contributions to the overall emissions and impacts are affected by durability, supply-chain processes, material content, production source, labor policies, repairability, transportation, rates of return, and tons of total consumption.

The paper moreover does not assess any of the commercial recommendation websites nor make any assertion about how fair individual fashion stores are.

Lastly, there are regulatory duties may be imposed under different jurisdiction. Although the European recommender-system transparency requirement is covered by Digital Services Act, it only takes effect according to the nature of platform and legislation situation.

Future Work

Further using ethically approved human-user studies and properly governed fashion-recommendation datasets, the above framework should also be experimentally validated with the aid of “real” human users in upcoming investigations.

Longitudinal assessment should be given precedence. Short-term justification measures alone cannot identify whether low-level individualized exposure over several months reduces style variety, decreases purchase rate, supports stereotypes, or shifts consumers’ notion of the appropriate. How would you rewrite this in a more human-sounding academic English?

In future research, the design of body-positive recommendations should be examined more explicitly. Kubullek and Erdem’s findings offer a valuable starting point for building a case as to how AI fashion consultation can foster inclusive and non-discriminative experiences.

User-adjustable personalization is a further possibility. Consumers may have the opportunity to adjust exploration, fine-tune their weights on sustainability, decide on commercial

personalization, set data retention period or constrain the extent of influence of past preferences.

Research is ongoing in this area, as well. Attribute explanation may help users understand why two garments work well together, but future systems should also be able to differentiate relevance to the user due to preferences from relevance due to paid placement, targeted inventory, quantity pressures, common trends, and other ranking metrics from the business perspective.

Sustainable-aware recommendation has to be assured with relevant lifecycle evidence, not ungrounded labels. Recommendation systems could embed longevity, maintainability, rental and resale possibilities, usage of current wardrobe, lower impact alternatives if credible data exists.

Researchers need to also examine provider side fairness. Independent designers, small businesses, culturally specific apparel manufactures, and less popular products may not get the attention they need if ranking algorithms overly emphasize historical engagement and popularity.

Finally, the challenging task of defining what really responsible recommendation is will require interdisciplinarity among recommender-systems engineering, consumer psychology, fashion research, sustainability science, privacy law, human-computer interaction, and ethics.

Declarations

- **Conflict of Interest:** The author has no conflicts of interests. Please note that the author identity, designation, affiliation, e-mail id and ORCID shown above are dummy sample metadata and should be verified before submission.
- **Funding:** None was received from outside sources for this simulation research.
- **Ethics Statement:** No human participants, customer profiles, photographs, biometric measures, browsing histories, purchase histories or PII were used in this research. No institutional human-subjects ethics approval was required because this synthetic analysis did not use human-subjects data.
- **Data Availability:** The statistical results presented above were obtained using made-up data for demonstration purposes only. The made-up data is not representative of

any specific retailer, recommendation website, consumer group or commercial AI system.

- **Proposed Questionnaire Items:** The five questionnaire items listed in the Methodology section are suggested items for future research and were not included in the study. There are no questionnaire results reported herein.
- **Sustainability Statement:** In this paper we treat sustainability-aware recommending as one part of the responsible fashion AI. We do not assume that recommend products labeled sustainable will always impose less overall environmental impact.
- **Author Contribution:** The concept, design, the simulation methods, statistical analysis, ethical issues and paper writing were implemented for this simulation study.
- **AI Transparency Statement:** AI-supported applications could recommend research writing/analysis processes but humans remain accountable to ensure the accuracy of methods, verification of facts, appropriate citations, sound conclusion drawing, and publication ethic compliance.

References

1. Y. Deldjoo, F. Nazary, A. Ramisa, J. McAuley, G. Pellegrini, A. Bellogín, et T. Di Noia, A review of modern fashion recommender systems, *ACM Computing Surveys* 56 (4) (2023) 1–34..
2. Y. Ding, Z. Lai, P. Y. Mok and T.-S. Chua, “Computational Technologies for Fashion Recommendation: A Survey”. *ACM Computing Surveys*. 2023. 56 (5). <https://doi.org/10.1145/3627100>
3. W.-H. Cheng et al., “Fashion meets computer vision: A survey,” *ACM Computing Surveys*, vol. 54, no. 4, 2021. <https://doi.org/10.1145/3447239>
4. K. Selwon, J. Szymański, A. G. Khan et al., ‘A survey of approaches to compatible fashion modeling that include transparent concepts’, *ACM Computing Surveys*, vol. 56, no. 11, pp. 1–29, 2024. <https://doi.org/10.1145/3664614>
5. X. Yang, X. Song, F. Feng, H. Wen, L.-Y. Duan, and L. Nie, “Attribute-wise explainable fashion compatibility modeling,” *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 17, no. 1, Art. 36, 2021. <https://doi.org/10.1145/3425636>

6. T. Ye, L. Hu, Q. Zhang, Z. Y. Lai, U. Naseem, and D. D. Liu, 'Show me the best outfit for a certain scene: A scene-aware fashion recommender system,' Proc. The Web Conference 2023, pp. 1172–1180, 2023. <https://doi.org/10.1145/3543507.3583435>
7. Y. Wang, W. Ma, M. Zhang, Y. Liu, and S. Ma, "A survey on the fairness of recommender systems," in ACM Transactions on Information Systems, 2023, vol. 41, pp. 1–43. <https://doi.org/10.1145/3547333>
8. Y. Zhao, Y. Wang, Y. Liu, X. Cheng, C. C. Aggarwal and T. Derr. Fairness and diversity in recommender systems: A survey. ACM Transactions on Intelligent Systems and Technology, 2024. <https://doi.org/10.1145/3664928>
9. T. Krause, L. Göritz and R. Gratz, "The effect of gender debiased recommendations A user study on gender-specific preferences," Proceedings of the CHI Conference on Human Factors in Computing Systems 2025. <https://doi.org/10.1145/3706598.3713155>
10. A.-K. Kubullek and D. Erdem, "AI vs. human fashion consultation for young women: A quantitative survey on expectations, formats and the ability to deliver positive body images," in Proceedings of OZCHI 2025 (2025), p. 159. <https://doi.org/10.1145/3764687.3769949>
11. European Commission, The impact of the Digital Services Act on digital platforms , 2026, available at: <https://digital-strategy.ec.europa.eu/en/policies/dsa-impact-platforms>
12. European Commission How the Digital Services Act brings better transparency online 26. <https://digital-strategy.ec.europa.eu/en/policies/dsa-brings-transparency>
13. European Commission: "Commission requests information from Shein on illegal products and its recommender system," 2025. <https://digital-strategy.ec.europa.eu/en/news/commission-requests-information-shein-illegal-products-and-its-recommender-system>
14. Federal have to Commission, Bringing dark patterns to light, Washington, DC, United States, 2022, <https://www.ftc.gov/reports/bringing-dark-patterns-light>
15. Federal Trade Commission, "FTC surveillance pricing study indicates wide range of personal data used to set individualized consumer prices," 2025. <https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-surveillance-pricing-study-indicates-wide-range-personal-data-used-set-individualized-consumer>

16. United Nations Environment Programme, The Sustainable Fashion Communication Playbook, 2023. <https://www.unep.org/resources/publication/sustainable-fashion-communication-playbook>
17. United Nations Environment Programme – “Sustainable fashion to take centre stage on Zero Waste Day”, 2025. <https://www.unep.org/technical-highlight/sustainable-fashion-take-centre-stage-zero-waste-day>
18. United Nations Environment Programme, “Sustainable and circular textiles” 2026. <https://www.unep.org/topics/chemicals-and-pollution-action/circularity-sectors/sustainable-and-circular-textiles>