Demographic Influences on Ethical Perspectives in Human-AI Co-Creation: An Exploratory Study

Jeba Rezwana
Computer & Information Systems
Towson University
Towson, MD 21252 USA
jrezwana@towson.edu

Abstract

The rapid advancement and widespread adoption of generative AI in creative domains have boosted research in human-AI co-creation. The dynamics of human-AI co-creation are notably more complex than traditional human-computer interaction, as AI not only processes and interprets data but also creates new content and works alongside humans. Additionally, ethical concerns are more pressing when humans collaborate closely with AI as partners, underscoring the need for human-centered ethical AI research. This paper explores the relationship between user demographics and users' ethical expectations concerning human-AI co-creation. We conducted a study involving 115 participants to explore this relationship. Initial findings provide insights into design principles for developing human-centered ethical co-creative AI and personalized AI.

Introduction

Human-AI co-creativity, a subfield of computational creativity, involves humans and AI collaborating on a shared creative product (Davis 2013). Unlike general human-computer interaction (HCI), human-AI co-creation creates a more complex relationship between humans and AI (Rezwana and Maher 2023b) as AI 1) actively collaborates in the creative process, 2) assumes human-like roles of partner, evaluator, and generator (Negrete-Yankelevich and Zaragoza 2014), and 3) creates novel content which is blended with the user's contribution. The rapid advancement and widespread adoption of generative AI in creative fields have significantly accelerated research in human-AI co-creation. Yet, designing co-creative systems that effectively comprehend and align with human partners' values presents significant challenges (Kantosalo et al. 2014). These challenges are further compounded by the growing complex ethical concerns around human-AI co-creation (Muller and Liao 2017; Chopra and Singh 2018).

The effectiveness of co-creative AI depends on users, shaped by their social, cultural and demographic factors (Barile et al. 2021). These factors and demographics shape users' expectations of AI systems (Meurisch et al. 2020), ethical awareness and risk perceptions in human-AI co-creation (Yoon and Jun 2023). Rezwana et al. emphasize design principles for ethical co-creative AI from a human-centered perspective to improve user experience (Rezwana and Maher 2023b). To develop human-centered ethical co-creative AI, identifying how different user demographics influence users' expectations and ethical stances is crucial (Gero et al. 2020; Xu 2019) as people are more motivated to use AI when AI is congruent with their values, goals and needs (Brühlmann et al. 2018). Therefore, it is essential to understand how diverse user groups perceive and react to ethical dilemmas posed by co-creative AI.

This paper explores the association between basic demographic factors, such as age, gender, AI literacy etc., and users' ethical stances towards four ethical dilemmas in human-AI co-creation: *ownership* of the co-created product, AI *anthropomorphism*, *data collection* and *societal impact of co-creative AI*. We define the four ethical dilemmas as follows: *ownership* pertains to whether the human or the AI should own the co-created product. *Data Collection* by AI concerns the gathering of user data, such as visual or biometric information, during the co-creation process to enhance interaction. *Anthropomorphism* in co-creative AI involves attributing human-like characteristics to AI partners, including speech, personality, and embodiment. Lastly, *Societal Impact* addresses the profound social effects of AI on creative domains. Our research involves a study conducted with 115 participants who engaged with two existing AI systems, ChatGPT and Stable Diffusion, in creative contexts and completed survey questions. We present the initial findings from the study that reveal associations between user demographics and users' ethical stances, which can contribute toward design guidelines for human-centered ethical co-creative AI as well as personalized AI.

Related Works

In co-creative systems, humans and AI contribute as partners in the creative process (Davis 2013), distinguished from autonomous creative systems which generate creative products independently, and creativity support tools which support human creativity (Kantosalo et al. 2020). The role of co-creative AI changes from a lone decision-maker to a more complex one, depending on the collaboration. Designing co-creative systems has unique challenges due to the spontaneity of the interaction between the human and the AI (Kantosalo et al. 2014). Moreover, as AI takes on roles as social entities and interacts with us, questions of values and ethics become even more urgent and complex (Liao et

al. 2016). Rapid advancements in AI have heightened the challenge of ensuring that AI's behaviors align with human values (Russell et al. 2015). To maintain fairness across diverse user groups, ethically aligned design is essential for human-centered co-creative AI systems (Xu 2019).

Current ethical guidelines for general AI technologies lack a focus on what they entail in the context of human-AI co-creation, which raises more complex ethical concerns (Llano and McCormack 2020), such as who owns the co-created product (Muller et al. 2022; Wang et al. 2021). Therefore, we should not assume that research on general AI ethics and HCI fully transfers to human-AI co-creation (Flathmann et al. 2021). Llano and McCormack suggest a common understanding of the ethical challenges in human-AI co-creative systems to devise ethical guidelines for co-creative AI (Llano and McCormack 2020). Recent research (Rezwana and Maher 2023b; Rezwana and Maher 2023a) demonstrates major ethical issues in human-AI co-creation, including ownership, AI anthropomorphism, data collection etc. Ethical principles based on empirical data are more likely to be translated into practice (Whittlestone et al. 2019). Rezwana and Maher highlight the importance of exploring ethical issues from the users' perspectives to develop human-centered ethical co-creative AI (Rezwana and Maher 2023b).

Meurisch et al. (Meurisch et al. 2020) found that users' expectations and perceptions of AI systems are influenced by human factors and demographics. The effectiveness of co-creative AI is influenced by the social and demographic factors that shape user experiences (Barile et al. 2021). Basic demographics of individuals, such as age, gender, and education, can be predictive of variability in their behavior (Wells 1975). Prior studies have found that demographic factors like age and gender significantly influence behaviors and interactions with AI technologies. (Bendell et al. 2021). For example, one study found a significant effect of age and gender on gamification in various human-computer interaction contexts (Tondello et al. 2019). Bendell et al. underscore the importance of identifying the influence of basic demographics on human-AI teams (Bendell et al. 2021).

Research on the impact of user demographic factors on users' ethical expectations of co-creation is scarce, creating a critical gap that must be addressed to develop human-centered ethical co-creative AI. Drawing on existing literature, this paper investigates the relationship between demographic variables and users' ethical perspectives through a study. The methodology of the study is outlined in the following section.

User Study

We conducted a study to explore the association between users' demographic factors and their ethical stances using Prolific, a platform for recruiting participants for online studies. Initially, participants engaged with two AI systems, ChatGPT and Stable Diffusion, performing a creative task with each of them to familiarize themselves with human-AI co-creation contexts. Participants completed a creative story-writing task with ChatGPT based on a given prompt. Subsequently, they used Stable Diffusion to create visual representations of their stories created with ChatGPT. Finally, participants responded to four questions about their ethical stances and provided their demographic data.

Participants

We recruited 155 participants, comprising 87 men, 64 women, and 4 non-binary individuals. The age range of participants spanned from 20 years old to 66 years old, resulting in an average age of 27.5 years (Med = 25, SD = 8.25). As our aim was to gather a diverse set of data across various demographics, we kept inclusion criteria to a minimum, requiring only that participants be at least 18 years old.

Data Collection

We used surveys [1] to collect data. Initially, we collected data on participants' familiarity with the AI systems we used by asking about their frequency of use (e.g., from never to every day) to explore whether familiarity influences any target variables. The following set of questions focused on gathering users' ethical stances within the broader context of human-AI co-creativity rather than being specific to the AI systems used. We asked 5-point Likert-scale questions (e.g., strongly agree to disagree) about four frequently discussed ethical challenges in the context of human-AI co-creativity: *ownership* of the co-created product, *data collection* by AI, *anthropomorphism* in co-creative AI and *societal impact* of AI, based on the literature (Smith 2022; Rezwana and Maher 2022).

In the end, we collected user demographics, including age, gender, ethnicity, highest level of education, first-generation college student status, disability status, annual income, AI literacy, field of work/study, profession, and political affiliation. Demographic data was collected at the end of the study to minimize the risk of stereotype threats influencing participants' responses to previous questions.

We presented the survey questions in each set in a randomized order to eliminate any potential response biases caused by the order of the questions. We included three attention-checking questions in the survey to check the diligence of the participants. Responses that failed these checks or did not complete the tasks with both AI as instructed were excluded.

Data Analysis

As most of the data we collected are categorical in nature, we used statistical analysis methods suitable for non-numerical data. We used Pearson's Chi-square test to identify associations between pairs of variables. Additionally, we conducted a cluster analysis to identify demographic groups and patterns of users' ethical stances, enabling more nuanced association analyses. We used K-modes clustering (Cao et al. 2012) due to its compatibility with categorical data. K-modes clustering is an extension of K-means, but instead of means, this algorithm uses modes to determine the cluster centroids.

[1] https://drive.google.com/file/d/1Ouq8chb0cvOXdHLzTR0HbILBSP1HMYVX/view?usp=sharing

Cluster number/Name	Who should Own	Attitude towards AI Anthropomorphism	Attitude towards Data Collection	Attitude towards Societal Impact of Co-creative-AI
1/Conservative-Positive Stance (114)	Human	Neutral	Somewhat good	Somewhat Positive
2/Liberal-Negative Stance (41)	Both	Somewhat Good	Neutral	Somewhat Negative

Table 1: Two major patterns of ethical stances represented by the cluster centroids.

	Identity Demographics						
Clusters	Age	Ethnicity	Formative community	Gender	Disability	Annual income	Political Affiliation
Cluster 1 (99)	21-24	White	A large city	Men	No	<$10,000	No political affiliation
Cluster 2 (39)	25-30	Hispanic/Latinx	A small city or town	Women	No	$10,000–$19,999	Liberal
Cluster 3 (17)	31-34	Black/African American	A large city	Women	No	$30,000–$39,999	Moderate

Table 2: Clusters of participants based on their identities (cluster centroids).

Results

In this section, we begin by detailing the clusters identified among users' ethical stances and demographic groups. Then, we present the associations between the clusters as well as the associations between individual pairs of variables.

Clusters of Users' Ethical Stances

We found two clusters (Table 1) of users' ethical stances towards four ethical dilemmas in the context of human-AI co-creation: *ownership* of co-created product, AI *anthropomorphism*, *data collection* and *societal impact of co-creative AI*. We applied the chi-square test to identify features that significantly influence cluster formation and found that all four ethical issues are significant factors (all p values < 0.05)

Participants in cluster 1, the larger cluster, predominantly hold the stance that humans should have sole ownership of the co-creative product. They also generally exhibit a positive outlook on the societal impact of co-creative AI. Given this combination of stances, we have labeled this cluster as ***Conservative-Positive*** ethical stance for referring to this cluster throughout the analysis. People in this cluster hold a neutral attitude toward AI anthropomorphism and a somewhat positive attitude toward data collection.

In contrast, cluster 2 consists of 41 participants who prefer the shared ownership of creative products between humans and AI while exhibiting a somewhat negative perspective on the societal impact of co-creative AI. Therefore, we label this cluster as the ***Liberal-Negative*** ethical stance. Participants in this cluster generally have a somewhat positive view of AI anthropomorphism and a neutral attitude toward data collection.

Clusters of User Demographics

We found three clusters (Table 2) of participants using demographics that are associated with their identities, such as their age, ethnicity, gender, annual income, disability, political affiliation, and childhood community as their formative background. To ensure these clusters accurately reflect our data, we conducted feature selection and applied the chi-square test to identify demographic variables that significantly contribute to forming more coherent identity clusters.

Identity cluster 1 is the largest cluster, consisting of 99 participants. This cluster predominantly consists of young white men aged 21-24 who grew up in large cities with no reported disabilities and generally having an annual income of less than 10,000 USD. Most participants in this group do not align themselves with any political affiliation.

Identity cluster 2 consists of 39 participants, mostly identifying themselves as Hispanic/Latinx women aged 25-30 with politically liberal views. Most participants in this cluster grew up in small cities or towns as their formative background and typically have an annual income ranging from $10,000 to $19,000.

Identity cluster 3 consists of 17 participants, primarily comprising black/African American women aged 31-35 without any disabilities with a moderate political affiliation. Their typical annual income range from $30,000 to $39,000 and they have a background of growing up in large cities.

Association between Ethical Stances and User Demographics

Association between the Clusters Using Pearson's Chi-square test, we found a significant association ($p = 0.035$, Cramer's $V = 0.21$) between the clusters of *ethical stances* and participants based on their *identities*. The findings re-

veal that white young men aged 21-24 (*Identity cluster 1*) are more likely to hold the *Conservative-Positive* ethical stance. Among the participants, Hispanic/Latinx women aged 25-30 (*Identity cluster 2*) with a liberal political view also tend to have the *Conservative-Positive* ethical stance. On the other hand, Black/African American women aged 31-34 (*Identity cluster 3*) tend to have the *Liberal-Negative* ethical stance.

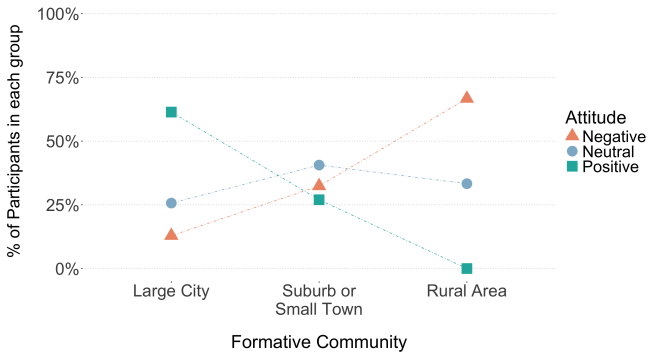

Figure 1: Distribution of users' ethical attitudes towards *data collection* by AI across formative communities.

Association between Pairs of Variables We examined the associations between all possible pairs of variables from ethical stances and participant demographics. No associations were found between users' familiarity with the systems used and any outcome variable. We report the significant associations from these analyses below.

We observed a significant association (Pearson's Chi-square p=0.028, Cramer's V = 0.22) between users' ethical stance towards ***data collection*** by AI and users' ***formative background*** (Figure 1). Participants who were raised in large cities generally exhibit a strong to moderate positive sentiment toward data collection by AI to improve the user experience during a co-creation. Conversely, those who were brought up in rural areas tend to have a negative stance towards data collection by co-creative AI. Individuals with a background in either small towns or suburban areas tend to adopt a more neutral attitude.

Additionally, our analysis revealed a significant association (Pearson's Chi-square p = 0.028, Cramer's V = 0.24) between ***gender*** and users' ethical stance towards ***ownership*** of the co-created product. 60.9% of men prefer exclusive human ownership of the co-created product, whereas 31% support shared ownership with AI. In contrast, 54.7% of women favor shared ownership, and 42.2% support exclusive human ownership. For non-binary participants, 50% prefer human ownership, with the remainder split between AI ownership and shared ownership.

We found a significant association (Pearson's Chi-square p = 0.022) between ***gender*** and users' ethical stance towards ***anthropomorphism*** in co-creative AI. Our data shows that, in general, both men and women predominantly feel either positive or neutral about AI anthropomorphism: 35.6% of men feel neutral, while 37.9% feel positive; 33.5% of women feel neutral, and 41.9% feel positive. This finding shows that men and women typically don't feel negative about giving AI human-like attributes in human-AI co-creation. In contrast, responses from non-binary participants were divided equally, with half feeling positive and the other half negative.

Discussion, Limitation and Future Work

The findings from the study revealed significant associations between users' ethical stances toward ethical dilemmas in human-AI co-creation and demographic factors, including user identity, gender and formative background. Association with users' ethical stance towards *data collection* by AI and their formative background prompts further investigation into the factors influencing regional user preferences, aiming to develop ethical co-creative AI that is considerate of a diverse user base. The association between identity clusters, ethical stances, and factors such as gender with views on *ownership* and AI *anthropomorphism* highlight the need for personalized AI in human-AI co-creation contexts. These results highlight the need for ethical design principles and policies that support the development of human-centered co-creative AI that are inclusive and respectful of all users, regardless of their identities and demographic backgrounds. Shi et al., in their recent literature review on co-creative AI, advocate for exploring the ethical implications of AI to cater to diverse user groups (Shi et al. 2023). These findings might provide initial insights for developing human-centered ethical co-creative AI and can be transferred into AI personalization in the context of human-AI co-creation. Comprehensive ethical principles based on empirical data are more likely to be translated into practice (Whittlestone et al. 2019). Our findings underscore the significance of cultural sensitivity and inclusivity in co-creative AI by considering multiple viewpoints and avoiding presumptions regarding ethical expectations, thus addressing ethical challenges in a human-centered manner. Future research is needed to establish human-centered ethical guidelines that account for the nuanced differences influenced by users' backgrounds, enhancing user experience.

Limitations The sample used in our study may not fully represent the broader population, which could limit the generalizability of our findings. Further research is necessary to examine how demographic factors influence users' ethical stances and expectations more comprehensively. Nonetheless, we believe that the findings from our study serve as valuable preliminary insights that can guide future research directions in relevant domains, such as value-sensitive AI and human-centered AI. Additionally, while our preliminary findings highlight associations between variables, they do not explore the details and complexities of these relationships in depth. To develop comprehensive and nuanced design principles for human-centered ethical AI and AI personalization, a more rigorous and in-depth analysis is required to fully understand the complexities of these associations.

Conclusion

This paper reports on an exploratory study involving 115 participants that examines the associations between user demographic factors and their ethical viewpoints towards ethical dilemmas in human-AI co-creation. The preliminary findings from the study are presented in the paper, which provide insights into designing human-centered ethical co-creative AI and personalization in human-AI co-creation.

References

[Barile et al. 2021] Barile, S.; Bassano, C.; Piciocchi, P.; Saviano, M.; and Spohrer, J. C. 2021. Empowering value co-creation in the digital age. *Journal of Business & Industrial Marketing* (ahead-of-print).

[Bendell et al. 2021] Bendell, R.; Williams, J.; Fiore, S. M.; and Jentsch, F. 2021. Supporting social interactions in human-ai teams: Profiling human teammates from sparse data. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 65, 665–669. SAGE Publications Sage CA: Los Angeles, CA.

[Brühlmann et al. 2018] Brühlmann, F.; Vollenwyder, B.; Opwis, K.; and Mekler, E. D. 2018. Measuring the "why" of interaction: Development and validation of the user motivation inventory (umi). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13.

[Cao et al. 2012] Cao, F.; Liang, J.; Li, D.; Bai, L.; and Dang, C. 2012. A dissimilarity measure for the k-modes clustering algorithm. *Knowledge-Based Systems* 26:120–127.

[Chopra and Singh 2018] Chopra, A. K., and Singh, M. P. 2018. Sociotechnical systems and ethics in the large. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 48–53.

[Davis 2013] Davis, N. M. 2013. Human-computer co-creativity: Blending human and computational creativity. In *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[Flathmann et al. 2021] Flathmann, C.; Schelble, B. G.; Zhang, R.; and McNeese, N. J. 2021. Modeling and guiding the creation of ethical human-ai teams. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 469–479.

[Gero et al. 2020] Gero, K. I.; Ashktorab, Z.; Dugan, C.; Pan, Q.; Johnson, J.; Geyer, W.; Ruiz, M.; Miller, S.; Millen, D. R.; Campbell, M.; et al. 2020. Mental models of ai agents in a cooperative game setting. In *Proceedings of the 2020 chi conference on human factors in computing systems*, 1–12.

[Kantosalo et al. 2014] Kantosalo, A.; Toivanen, J. M.; Xiao, P.; and Toivonen, H. 2014. From isolation to involvement: Adapting machine creativity software to support human-computer co-creation. In *ICCC*, 1–7.

[Kantosalo et al. 2020] Kantosalo, A.; Ravikumar, P. T.; Grace, K.; and Takala, T. 2020. Modalities, styles and strategies: An interaction framework for human-computer co-creativity. In *ICCC*, 57–64.

[Liao et al. 2016] Liao, Q. V.; Davis, M.; Geyer, W.; Muller, M.; and Shami, N. S. 2016. What can you do? studying social-agent orientation and agent proactive interactions with an agent for employees. In *Proceedings of the 2016 acm conference on designing interactive systems*, 264–275.

[Llano and McCormack 2020] Llano, M. T., and McCormack, J. 2020. Existential risks of co-creative systems. In *Workshop on the Future of Co-creative Systems 2020*. Association for Computational Creativity (ACC).

[Meurisch et al. 2020] Meurisch, C.; Mihale-Wilson, C. A.; Hawlitschek, A.; Giger, F.; Müller, F.; Hinz, O.; and Mühlhäuser, M. 2020. Exploring user expectations of proactive ai systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4(4):1–22.

[Muller and Liao 2017] Muller, M., and Liao, Q. V. 2017. Exploring ai ethics and values through participatory design fictions. *Human Computer Interaction Consortium*.

[Muller et al. 2022] Muller, M.; Ross, S.; Houde, S.; Agarwal, M.; Martinez, F.; Richards, J.; Talamadupula, K.; Weisz, J. D.; Human-Centered, A.; Suneja, S.; et al. 2022. Drinking chai with your (ai) programming partner: A design fiction about generative ai for software engineering.

[Negrete-Yankelevich and Zaragoza 2014] Negrete-Yankelevich, S., and Zaragoza, N. M. 2014. The apprentice framework: planning and assessing creativity. In *ICCC*, 280–283.

[Rezwana and Maher 2022] Rezwana, J., and Maher, M. L. 2022. Identifying ethical issues in ai partners in human-ai co-creation. *arXiv preprint arXiv:2204.07644*.

[Rezwana and Maher 2023a] Rezwana, J., and Maher, M. L. 2023a. User perspectives of the ethical dilemmas of ownership, accountability, leadership in human-ai co-creation.

[Rezwana and Maher 2023b] Rezwana, J., and Maher, M. L. 2023b. User perspectives on ethical challenges in human-ai co-creativity: A design fiction study. In *Proceedings of the 15th Conference on Creativity and Cognition*, 62–74.

[Russell et al. 2015] Russell, S.; Hauert, S.; Altman, R.; and Veloso, M. 2015. Ethics of artificial intelligence. *Nature* 521(7553):415–416.

[Shi et al. 2023] Shi, Y.; Gao, T.; Jiao, X.; and Cao, N. 2023. Understanding design collaboration between designers and artificial intelligence: A systematic literature review. *Proceedings of the ACM on Human-Computer Interaction* 7(CSCW2):1–35.

[Smith 2022] Smith, J. 2022. Human-ai partnerships in generative music. In *International Conference on New Interfaces for Musical Expression*. PubPub.

[Tondello et al. 2019] Tondello, G. F.; Mora, A.; Marczewski, A.; and Nacke, L. E. 2019. Empirical validation of the gamification user types hexad scale in english and spanish. *International Journal of Human-Computer Studies* 127:95–111.

[Wang et al. 2021] Wang, D.; Maes, P.; Ren, X.; Shneiderman, B.; Shi, Y.; and Wang, Q. 2021. Designing ai to work with or for people? In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–5.

[Wells 1975] Wells, W. D. 1975. Psychographics: A critical review. *Journal of marketing research* 12(2):196–213.

[Whittlestone et al. 2019] Whittlestone, J.; Nyrup, R.; Alexandrova, A.; and Cave, S. 2019. The role and limits of principles in ai ethics: towards a focus on tensions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 195–200.

[Xu 2019] Xu, W. 2019. Toward human-centered ai: a perspective from human-computer interaction. *interactions* 26(4):42–46.

[Yoon and Jun 2023] Yoon, H., and Jun, S. 2023. Ethical awareness of uxers in the loop: Ethical issues in the uxer-ai collaboration process from a ux perspective. In *Proceedings of the 25th International Conference on Mobile Human-Computer Interaction*, 1–6.