Co-Designing Fashion with AI: A Small-Data Approach to Generative Garment Design
Paper type: Cultural Application

Ina Hrešć[1] **and Terence Broad**[2]
Creative Computing Institute
University of the Arts London
United Kingdom
[1]i.hresc0620231@arts.ac.uk, [2]t.broad@arts.ac.uk

Abstract

This paper details a case study of implementing a co-design framework for integrating artificial intelligence (AI) into fashion design through a collaboration between the researcher and a garment designer. Using the designer's fashion development as a practical example, this research investigates AI's potential to enhance creative workflows, improve efficiency, and facilitate marketing strategies that reduce the necessity for physical prototyping. The project implements a small-data approach, exclusively fine-tuning AI models on the designer's sketches and the researcher's photography to ensure highly personalised outputs. AI architectures, such as Stable Diffusion and ControlNet, are fine-tuned and used to generate sketches and photorealistic visualisations that work as an extension of the designer's artistic vision while promoting sustainable awareness and ethical data practices. Through continuous dialogue and a final interview, the designer's perspective was analysed, demonstrating AI's efficacy as a creativity support tool that streamlines design iteration by enabling rapid prototyping and ideation, while also addressing concerns about preserving artisanal expertise. This research contributes to the growing field of AI in the creative industries by demonstrating the value of co-design, small-data strategies, and ethical AI development to build tools that empower designers.

Introduction & Theoretical Context

Generative artificial intelligence (AI) has increasingly impacted the creative industries, not only as a tool for automation but as a means of supporting and enhancing human-led creative workflows. In fashion, AI's visual abilities present opportunities for human-machine collaboration, where generative systems can support designers in the creative process. This research explores the practical potential of diffusion-based models, specifically Stable Diffusion (Rombach et al., 2022) with ControlNet (Zhang, Rao, and Agrawala, 2023), in a close partnership with an independent fashion designer. Therefore, we position AI as a powerful assistive tool that facilitates iterative design exploration (Wu et al., 2021), thus enhancing the designer's workflow. The creative direction regarding the fashion designs was guided by the desire to avoid explicit gender associations, adopting a neutral perspective (Fan, 2023), to allow for the representation of fluid gender identities. This shaped the aesthetic goals and our approach to ethical data practices and inclusivity in design. By eliminating binary perceptions, we fostered an environment that allows the exploration of diverse identities, including those that are gender non-conforming and non-binary (Marquez-Gallardo and Rovira-Lorente, 2024). To achieve the desired gender neutral presentation, negative prompting, which functions by neutralising unwanted elements in the latent space of diffusion models (Ban et al., 2024), was used to mitigate clear gender identifiers. Furthermore, post-processing retained the mannequin aesthetic, removing potential biases related to facial features and hair, while also merging the designer's traditional sketching with AI's photorealistic garment renderings.

Generative AI tools, such as Generative Adversarial Networks (GANs) and diffusion models, have demonstrated potential in fashion design (Guo et al., 2023). Early work with GANs, as proposed by Goodfellow et al. (2014) and exemplified by FashionGAN, showcased AI's sketch-to-image generative abilities (Cui et al., 2018). However, GANs often suffer from limitations such as mode collapse (Thanh-Tung and Tran, 2020) and reliance on large datasets, constraining creative flexibility (Karras et al., 2018; Razavi, van den Oord, and Vinyals, 2019). By contrast, latent diffusion models (LDMs), like Stable Diffusion, have emerged as more flexible alternatives, operating in a compressed latent space and enabling diverse, high-quality image outputs that require less computational power (Rombach et al., 2022). ControlNet further enhances control by incorporating spatial conditioning, allowing for precise control of generated images (Zhang, Rao, and Agrawala, 2023). For instance, the Multimodal Garment Designer (MGD), proposed by Baldrati et al. (2023), enabled guided image generation through textual descriptions, body poses, and sketches. Zhang, Zhang, and Xie (2024) presented a two-stage approach using ControlNet for sketch-to-image and image-to-image refinement, showcasing precise control over garment structure and details.

Technical advances in AI-assisted fashion design are well documented. However, the literature lacks a comprehensive examination of close collaborative relationships between developers and fashion designers. This research adopts a co-design framework, emphasising shared power, mutual learning, and open communication (McKercher, 2020). By utilising a small-data approach, fine-tuning AI models exclusively on personal data (designer's sketches, primary researcher's personal photography), this study promotes ethical data practices and sustainable AI development. We note that, in creative machine learning, goals often differ from those of traditional ML, as highlighted by Vigliensoni, Perry, and Fiebrink (2022). They argue that the assumption that "bigger data is better in generative AI systems" does not necessarily apply to more niche creative use contexts. DreamBooth, a fine-tuning technique, allows for learning of specific subjects and personalised style transfer with minimal data (Ruiz et al., 2022). Although LDMs are less computationally intensive than GANs, when trained on large datasets, they still require significant computational resources and increased model training times, leading to negative environmental impacts (Utz and DiPaola, 2023). By reducing model training times through small-scale fine-tuning, DreamBooth offers a potential reduction in these impacts. Further, the commercial use of AI models raises critical concerns for creatives, including potential job displacement and ethical issues such as copyright and privacy (Clarencia et al., 2024). The use of personal data and a co-design framework emphasises fairness and accountability in AI development, addressing crucial ethical considerations related to data governance and transparency (European Commission, 2019). Furthermore, the use of small personal datasets reduces the risk of bias, such as the inadvertent perpetuation of stereotypes, and increases the explainability of outputs.

This research aims to address the following questions:

- How can generative AI technologies, such as Stable Diffusion and ControlNet, enhance the creative processes of independent designers while preserving their distinct styles?
- To what extent can small-scale, personalised fine-tuning enable designer-led prototyping for sustainable workflows?
- Is the use of small, ethically sourced data a viable approach for ethically responsible AI-assisted workflows and for fostering creative collaboration between developers and designers in bespoke fashion projects?

This study contributes to the field of computational creativity by demonstrating the potential of using AI as a creativity support tool in fashion design, emphasising co-design, sustainability, and ethical data practices through practical experiments and collaborative work. It reflects an emerging form of human-AI co-creation, which positions AI as a complementary force that enhances and expands creative possibilities, as exemplified by AI-assisted poetry (Wu et al., 2021). This approach aligns with the foundational concepts of computational creativity, wherein AI can significantly aid the exploration and transformation of creative spaces (Boden, 2004). The following section details the implementation of and technical workflow used to explore these concepts.

Methodology

Figure 1 provides a high-level overview of the research workflow, which integrates co-design feedback loops. This section details our methodology from data acquisition and experimental design to post-processing and collaborative evaluation.

Data Acquisition & Preparation

All data was collected ethically and with informed consent, ensuring the designer was continuously updated on data usage and understood the aims of the project. The designer contributed 29 annotated sketches, each reflecting their unique style and specific project objectives. The researcher contributed 44 personal photographs, taken during various photoshoots, for backdrop generation, aiming for a consistent urban and rustic aesthetic.

Data preprocessing involved several transformations to prepare the images for model training. Designer annotations were removed, sketches were cleaned, and all images were resized and, where necessary, padded to 768x768 pixels. For data augmentation, all sketches were horizontally flipped. To serve as control inputs alongside prompts, edge maps were generated.

Experimental Design and Implementation

Stable Diffusion 2 was fine-tuned on the 58 sketches (29 original orientation, 29 flipped) using the DreamBooth technique (Ruiz et al., 2022). To guide the model in replicating specific stylistic details for each photorealistic visualisation, individual image prompts were crafted based on the designer's annotations of their corresponding sketch counterparts. A batch size of 4 was selected to maximise GPU memory usage while maintaining training efficiency. Gradient accumulation of 2 was used to simulate a larger batch size, improving training stability. A learning rate of 1e-6 was selected for training stability.

For photorealistic rendering, *realistic-vision-v51* was selected due to its superior ability to produce high-quality photorealistic images (Saovana and Khosakitchalert, 2024). It was fine-tuned on 44 photographs from the backdrop dataset to achieve desired textures and lighting conditions in the image backdrops. ControlNet was used with edge maps from the sketches as control inputs (von Platen et al., 2023). We adjusted the *guidance_scale* and *controlnet_conditioning_scale* parameters to optimise results for each sketch. The *guidance_scale* balanced prompt adherence and image structure, with higher values emphasising prompt accuracy. The *controlnet_conditioning_scale* influenced ControlNet's strength, ensuring structural adherence to the edge maps without overpowering the backdrop style transfer.

Backdrop generation involved fine-tuning *realistic-vision-v51* on the 44 selected backdrop images. We explored two approaches:

- Fine-tuning with image captions (CSV descriptions).
- Fine-tuning without image captions.

Training parameters for backdrop generation were set to a learning rate of 5e-7, chosen to prevent catastrophic forgetting, while still allowing the model to learn the specific characteristics of the backdrop dataset. A train batch size of 3 and gradient accumulation steps of 2 were used to balance

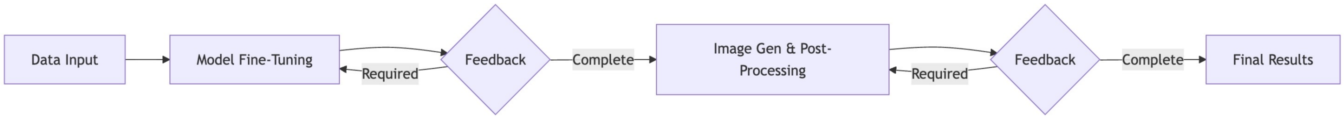

Figure 1: High-Level Workflow: Overview of the methodology, illustrating the integration of co-design feedback loops.

Experiment	Model Used	Dataset	Training Method	LR	Batch Size	Observations
SDv2 Sketch Generation	Stable Diffusion 2	58 sketches	DreamBooth	1e-6	4	Best: Checkpoints 600 & 800
RVv5.1 Backdrop Generation (Captioned)	Realistic-Vision-v5.1	44 backdrops	DreamBooth	5e-7	3	Suboptimal cohesion
RVv5.1 Backdrop Generation (Uncaptioned)	Realistic-Vision-v5.1	44 backdrops	DreamBooth	5e-7	3	Better cohesion, Checkpoint: 2400

Table 1: Summary of Training Experiments and Model Configurations

GPU memory usage and training stability, particularly important given the smaller dataset size.

Post-Processing and Collaborative Evaluation

Adobe Photoshop was used for post-processing, specifically to mask and blend AI-generated garments with the original mannequin sketches. This process was essential for achieving the desired neutral mannequin aesthetic and ensuring the final renderings retained the designer's distinctive drawing style in the mannequin body parts, resulting in outputs that merged AI-generated and traditional art.

Qualitative assessment played a key role in the project. Continuous dialogue with the designer throughout the implementation process allowed for iterative adjustments and refinements. A structured interview was conducted to gather detailed feedback on the model's performance and the collaborative process. Relevant quotes and remarks from the designer will be included in the results and discussion sections, providing insights into the computer-aided design process. Prior to the interview, informed consent was obtained from the designer, explicitly granting permission for the inclusion of interview data in published materials, alongside permission for the use of the designer's sketches for project purposes.

Results

Table 1 summarizes the configurations used for each training experiment. This section details the outcomes of our experiments, focusing on the outputs from sketch and backdrop generation, as well as the results of our post-processing procedures. Relevant qualitative assessments from the interview with the designer is presented here, with a more comprehensive analysis to follow in the discussion section. The aim of this structured presentation is to evaluate how the implemented AI techniques addressed our research questions and contributed to the project's creative objectives.

Sketch Generation

The fine-tuned Stable Diffusion 2 model demonstrated strong style adaptation, accurately replicating the designer's sketching style, as shown in Figure 2. It also exhibited contextual flexibility, rendering various materials and garment structures at different angles. For instance, it displayed pleats on clothing from both front and side angles, and paired with different garment types. Notably, the model successfully generated garment types not explicitly present in the training data, such as leather, showcasing its ability to adapt to diverse textures. The optimal model performance (as determined by the designer) was observed at checkpoints after 600 and 800 training iterations. Beyond these checkpoints, visual results showed a gradual loss of detail and deteriorated in clarity, indicating overfitting. The designer remarked that "the AI-generated sketches are impressive and closely mirror my style" and that they streamlined the design process. This, according to the designer, is particularly beneficial when working under tight deadlines, emphasising the practical value of this approach to assisting their workflow.

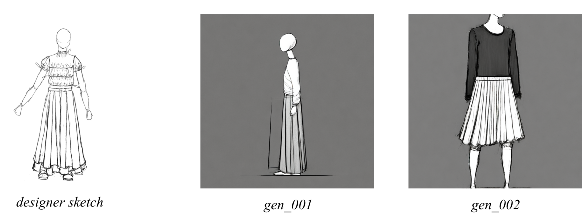

Figure 2: Comparison of the designer's original sketches and AI-generated sketches from the fine-tuned Stable Diffusion 2 model, demonstrating style replication and contextual flexibility, here regarding pleats.

Backdrop Generation

Prior to ControlNet's implementation, images were generated from captioned and uncaptioned training data. Using a fixed prompt, we compared the original *realistic-vision-v51* outputs against those generated from uncaptioned and captioned training. The fine-tuned outputs exhibited a significantly more worn and rustic appearance compared to the original model, aligning more closely with the training data. Notably, the uncaptioned dataset yielded cleaner,

designer sketch

fine-tuned outputs

realistic-vision-v51 outputs

Figure 3: Designer sketch next to photorealistic outputs generated by the original *realistic-vision-v51* model (upper row) versus the fine-tuned model (lower row), showcasing the impact of fine-tuning on backdrop structure, textures, and lighting.

designer sketch

realistic-vision-v51

captioned fine-tuning

uncaptioned fine-tuning

Figure 4: Designer sketch above comparison of original *realistic-vision-v51* (left), captioned training (middle), and uncaptioned training (right) outputs, illustrating the superior cohesion achieved with uncaptioned data.

more prompt-accurate backdrops, effectively capturing visual style without textual guidance. By contrast, captioned training posed challenges in consistently conveying complex visual and spatial information, highlighting the potential advantages of training solely on visual input.

Model Selection & ControlNet Implementation

To achieve our aesthetic goals of photorealism, we conducted a comparative analysis of multiple models before selecting *realistic-vision-v51* for fine-tuning. As shown in Table 1, *realistic-vision-v51* was fine-tuned using our backdrop dataset to achieve desired textures and lighting conditions in our results.

To ensure structural consistency between the designer's sketches and the photorealistic generative versions, we used ControlNet. Edge maps generated from the sketches served as the control input. ControlNet Scribble and Canny were used. Each showcased distinct strengths:

- *Scribble* excelled at handling sketches with less defined features.
- *Canny* was superior for structured edges and complex patterns.

While *Canny* successfully replicated the desired textures and lighting effects, structural differences in the backdrops were less pronounced, possibly due to stricter adherence to the detected edges. However, this subtlety aligned with the designer's preference for understated backdrops in fashion marketing.

Integrated Results: Fine-Tuning *realistic-vision-v51*

Here we provide a detailed account of the results obtained from the fine-tuned *realistic-vision-v51* model, focusing on the comparative analysis of fine-tuned versus standard *realistic-vision-v51* outputs and captioned versus uncaptioned dataset performance.

Fine-Tuned vs. Original *realistic-vision-v51* Outputs To determine the impact of fine-tuning on the model's ability to render desired visual styles within ControlNet, a comparative analysis was conducted between the generated images of the fine-tuned model and the original *realistic-vision-v51* model. The fine-tuning resulted in a noticeable impact on the visual characteristics of the outputs concerning general layout, textures, and lighting (Figure 3). Consistent with the training data, the textures exhibited a more urban and rustic aesthetic. The lighting appeared less diffused and more naturally integrated. Even in the absence of colour specifications in the prompt, the fine-tuned results displayed a greater variety of colours. While the original *realistic-vision-v51* renderings demonstrated slightly more coherent facial structures, this was deemed inconsequential, as the final images were intended to feature the original 2D mannequin aesthetic through post-processing. The designer's qualitative feedback indicated that the garments appeared more natural with the fine-tuned backdrops, and overall image aesthetics were perceived as "interesting and emotionally rich" compared to the original *realistic-vision-v51* results.

Captioned vs. Uncaptioned Dataset Performance To evaluate the influence of captioned versus uncaptioned training data on model performance when used in conjunction with ControlNet, a series of photorealistic renderings were generated using fixed parameters. Although the model fine-tuned on the captioned dataset produced some coherent and high-quality images, it encountered inconsistencies with visual cohesion across multiple runs, even with simple text prompts. The model fine-tuned on uncaptioned images consistently delivered higher-quality and more cohesive renderings, as seen in Figure 4. The uncaptioned model retained structural similarities to the standard *realistic-vision-v51* output while effectively incorporating the learned textures and lighting from the training data. This resulted in a more cohesive aesthetic, whereas the captioned visuals appeared more incoherent, as the model attempted to reimagine visual elements based on textual descriptions during training. The fashion designer emphasised the importance of subtle backgrounds in fashion, stating that "typically, in fashion shoots and visualisations, it's important to keep the background quite subtle to emphasise the garments themselves". Consequently, the uncaptioned fine-tuning proved more aesthetically suitable for AI-driven fashion applications.

Implementation of Negative Prompting for Enhanced Control

Negative prompting, a technique involving the specification of undesirable elements in generative AI systems, was implemented in later stages of this research. This was instrumental in achieving finer control over several critical aspects of the generated images:

- Enhanced control over textures and colours, enabling more precise manipulation of visual attributes.
- Improved control over the backdrop and separation of specific garment sections, particularly in preventing the unwanted projection of patterns onto inappropriate areas when using ControlNet Scribble.
- The creation of a unisex or gender-neutral aesthetic for the generated figures, exemplified by prompting the model to remove gender-specific features.

Post-Processing & Final Image Refinement

In the concluding phase of this research, post-processing procedures were employed to refine the final results. Post-processing involved the superimposition of the original mannequin body parts onto the photorealistically rendered fashion designs. As confirmed by the designer, this achieved renderings in line with the creative vision for the project. We utilised Adobe Photoshop for this purpose. A selection of the final results is shown in Figure 5.

Discussion

This research demonstrates the significant potential of integrating AI as a co-creative mechanism within fashion design processes, particularly through collaborative efforts guided

Figure 5: Final post-processed images showcasing the integration of AI-generated garments with the original mannequin sketches, achieving the desired aesthetic.

by McKercher's (2020) co-design principles. This collaborative approach between fashion designer and AI technologist aligns closely with our research questions, addressing how AI can enhance and guide creative workflows, facilitate ethical data usage, and support sustainable practices. Using small data and reducing the need for physical prototypes, similar to Zhang, Zhang, and Xie (2024), this approach aims to minimise waste and promote environmentally conscious design practices.

AI as a Creativity Support Tool Central to this research was the investigation of AI's role in augmenting the design workflow by functioning as a creativity support tool, specifically through using fine-tuned Stable Diffusion models with ControlNet for personalised outputs. This process of fine-tuning and spatial conditioning can be understood as "adapting generative spaces" (Abuzuraiq and Pasquier, 2024), allowing greater creative control and highlighting the designer's agency, especially in contrast to using pretrained models. The project's success in producing photorealistic visualisations that closely aligned with the designer's intent illustrates the value of AI in supporting creative ideation and visual experimentation. The designer's feedback underscored this, noting that the AI enabled "quick iterations" and facilitated the exploration of diverse design concepts, streamlining the process of developing and presenting concepts to potential clients. The system's ability to replicate the designer's sketching style and generate high-quality visual renderings demonstrates its efficacy as a creativity support tool within a human artistic process. By alleviating the burden of extensive hand-drawing, the AI empowered the designer to focus on conceptual advancement, a critical factor of meeting demands of high-pressure design timelines. The visual culmination of this co-creative process (Figure 5) demonstrates the success of our co-design framework in producing a visually compelling outcome approved by the designer.

Sustainable Practices Through AI-Driven Design & Small Data This project underscores the value of using small data in AI-driven design. We have shown that AI models, fine-tuned on small datasets, can achieve aesthetic visions aligned with the designer's intent. Thus, we argue that small, personal data is not only a viable, but favourable, approach in design collaborations with AI, especially concerning bespoke creative projects. This finding supports the "Small-Data Mindset" proposed by Vigliensoni, Perry, and Fiebrink (2022), which emphasises that niche and personalised datasets often allow for greater creative control and efficiency in AI systems compared to large, generalised ones.

Furthermore, introducing prototyping in the form of photorealistic visualisations based on small data promotes sustainability in several ways:

- It allows for greater focus on developing ideas and visual representations with audiences. This decreases the need for physical prototyping and thereby reduces material waste.
- By leveraging small data and DreamBooth for fine-tuning instead of relying on big data to provide universal, yet

often less creative, solutions, less computational power is required. While more research on the connection between leveraging DreamBooth and sustainability is required, using small data theoretically leads to a reduction in the carbon footprint associated with AI development. This aligns with sustainable practices in the industry (Rombach et al., 2022; Utz and DiPaola, 2023).

These findings suggest that small-data approaches offer a promising pathway towards more sustainable and personalised AI-driven design practices in the fashion industry. However, it is important to acknowledge the complexities of comparing the carbon footprints of traditional physical prototyping and AI model training. Sustainability claims in our study should be viewed as emphasising the *potential* for reduced environmental impact, rather than a definitive assertion.

Ethical Data Usage, Collaborative Innovation, & Designer Empowerment This research highlights the significant value of using personal data within the creative process, particularly for bespoke projects. By concentrating on a compact dataset, derived solely from the researcher and the designer, whose professional work portfolio encompasses both fashion and costume design, the project demonstrates that extensive data repositories are not always necessary to achieve high-quality results. Beyond the advantages regarding sustainability mentioned above, this approach also addresses critical ethical considerations surrounding data obtainment and ownership. The findings advocate for a shift towards ethically sourced, small-scale data within the fashion industry. This positions our approach as a framework for future collaborations in creative sectors.

During the formal interview, the designer expressed concerns that advancements in AI could potentially diminish the need for human designers and their artisanal expertise across the design spectrum: "My primary concern is the potential for AI to replace artisanal jobs within the fashion industry". This is particularly relevant in a fashion market driven by large corporations, which often prioritise economic gains over the value of craftsmanship. Additionally, the designer noted that in the theatre industry, "the role of the costume designer is often overlooked", raising worries that "costume designers may be replaced by AI". This issue is particularly pressing for productions seeking to allocate higher budgets to other components, such as acting and stage design. The designer's concern is further supported by survey findings that highlight how budgetary constraints often disproportionately affect costume departments (Nelson, 2022).

To mitigate this risk, this case study serves as an illustrative example of generative AI not being used as a substitute for human creativity, but as an enabler of it, fostering a synergistic relationship between designer and technology. The designer's notion that "it's crucial to foster a relationship grounded in informed consent and collaboration, ensuring that designers are not exploited and that AI tools are utilised responsibly" suggests that original intent must remain central, ensuring that AI-generated outputs amplify rather than overshadow artistic visions.

Conversely, the designer highlighted how AI could be leveraged by independent fashion designers with limited production budgets. In this context, AI can save costs not only by reducing the need for physical prototypes but also by streamlining other design factors, such as pattern drafting. A good example of AI accentuating human creativity is the designer remarking how imperfections generated by AI can be turned into artistic elements like hand painting the garments, and as such reclaiming the AI-generated images as their own. In fact, they named this the biggest source of newly found inspiration through the project.

Throughout the project, continuous dialogue between the researcher and the designer fostered a reflective practice, ensuring that ethical considerations were woven into every phase of the creative process.

Inclusivity in Design While not the primary focus of the research, the project's exploration of genderless fashion opens up broader discussions about inclusivity in design. By eliminating traditional gender binaries from the design process, this research fosters an environment where diverse identities can be explored and expressed. The insights gained contribute to the ongoing discourse surrounding gender representation in fashion (Marquez-Gallardo and Rovira-Lorente, 2024), emphasising the need for more inclusive design practices that reflect a wider spectrum of identities. The designer noted being pleasantly surprised by the AI's representations of gender in the generated models wearing the garments, even prior to post-processing with the overlaid mannequin aesthetic. This observation reassured them that the initial designs successfully conveyed fluidity of gender.

Nuances of Human-AI Creative Processes The designer noted that while many of the generated designs resonated with their creative vision and were a source of idea generation, some "overshadowed" the original intent (these were not included in any of the post-processed and finished visualisations), occasionally leading to a disconnect between designer and AI. Although this can be beneficial for developing new ideas, it necessitates critical evaluation, as the research underscores the importance of retaining artistic integrity.

Limitations, Future Research & Industry Implications This study is limited in scope due to its single-participant design and qualitative methodology, which, while offering rich insights into one co-design collaboration, does not allow for broad generalisability. The subjective evaluation based on one designer's feedback reflects individual aesthetic preferences and creative priorities, which may not translate directly to other practitioners or design contexts.

Future research should therefore explore the scalability and adaptability of the co-design framework across a wider range of participants, including designers of varying expertise, backgrounds, and stylistic orientations. Studies could test the method within distinct fashion sectors, such as haute couture, streetwear, costume design, or mass-market production, to assess how effectively generative AI supports diverse creative processes.

Additionally, optimising ControlNet's fine-grained control over colours, materials, and details through alternative modules or integrated AI techniques, such as semantic segmentation, would be beneficial. Evaluating how well these improvements preserve the designer's intent across styles would be a valuable contribution.

To further validate the sustainability potential of small data and AI-driven prototyping, future works should conduct comprehensive life cycle assessments, comparing it directly with traditional methods. Further studies should also examine the feasibility of broader industry adoption, including on-demand production systems. Logistical, economic, and environmental impacts must be thoroughly examined.

Finally, further research should address job displacement concerns by exploring AI integration into educational programmes and fostering synergistic collaborations between technologists and designers, ensuring a future where AI and human creativity coexist.

Ethical Considerations

Ethical considerations are central to this research, particularly in relation to data usage, representation, environmental impact, and the role of AI in creative collaboration.

A primary ethical concern addressed in this project is data provenance and consent. By sourcing data exclusively from the designer and researcher, we ensured informed consent and mitigated bias, emphasising ethical data provenance. We critically evaluated AI-generated representations to align with inclusivity and respect for diverse identities, ensuring AI amplified, rather than misrepresented, the designer's vision.

Advocating for small, ethically sourced datasets, we acknowledge the carbon footprint of AI models and encourage future life cycle assessments to compare AI-driven versus traditional prototyping.

We emphasise AI as a creativity support tool and enabler of human creativity, with the co-design framework ensuring the designer's intent remains paramount. Continuous dialogue throughout the project embedded ethical considerations into every phase, contributing to the discourse on AI's impact on the fashion industry.

Concluding Thoughts

This research demonstrates the transformative potential of human-machine collaboration in fashion design, achieved through a co-design approach. By utilising small, ethically sourced data, we produced high-quality outputs, while also promoting sustainable awareness and personalised projects. The study advocates for ethically responsible data utilisation and highlights AI's role in inclusive design by helping dismantle restrictive binaries and enabling broader forms of diverse representation.

AI served as a catalyst for idea generation, enabling rapid prototyping and streamlining the design process, all while reinforcing the indispensable role of human creativity. Continuous dialogue ensured that ethical considerations were embedded in every phase of the AI-aided creative process, safeguarding artistic intent. To advance this work, future research should incorporate diverse designers and styles to further explore the nuances of human-AI collaboration across different design contexts.

Ultimately, this research supports the development of innovative, inclusive, and sustainable practices in fashion design by offering a roadmap for collaborations that harness AI's generative capabilities to support, rather than replace, human creativity.

Acknowledgments

We would like to thank the collaborating fashion designer for their creative input and feedback throughout the project. This work has been supported by the Creative Computing Institute, University of the Arts, London. Gratitude is also extended to the reviewers for their constructive feedback.

References

Abuzuraiq, A. M., and Pasquier, P. 2024. Seizing the means of production: Exploring the landscape of crafting, adapting and navigating generative ai models in the visual arts. *ArXiv* abs/2404.17688.

Baldrati, A.; Morelli, D.; Cartella, G.; Cornia, M.; Bertini, M.; and Cucchiara, R. 2023. Multimodal garment designer: Human-centric latent diffusion models for fashion image editing. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* 23336–23345.

Ban, Y.; Wang, R.; Zhou, T.; Cheng, M.; Gong, B.; and Hsieh, C.-J. 2024. Understanding the impact of negative prompts: When and how do they take effect? In *European Conference on Computer Vision*, 190–206. Springer.

Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge.

Clarencia, E.; Tiranda, T. G.; Achmad, S.; and Sutoyo, R. 2024. The impact of artificial intelligence in the creative industries: Design and editing. In *2024 International Seminar on Application for Technology of Information and Communication (iSemantic)*, 440–444. IEEE.

Cui, Y. R.; Liu, Q.; Gao, C. Y.; and Su, Z. 2018. Fashiongan: Display your fashion design using conditional generative adversarial nets. *Computer Graphics Forum* 37.

European Commission. 2019. Ethics guidelines for trustworthy ai, shaping europe's digital future. Accessed: 30 September 2024.

Fan, S. 2023. A kernel inquiry into genderless fashion: The social concepts it conveys. *Communications in Humanities Research* 15:67–72.

Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial networks.

Guo, Z.; Zhu, Z.; Li, Y.; Cao, S.; Chen, H.; and Wang, G. 2023. Ai assisted fashion design: A review. *IEEE Access* PP:1–1.

Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive growing of gans for improved quality, stability, and variation. In *Proceedings of International Conference on Learning Representations (ICLR) 2018*. International Conference on Learning Representations, ICLR ; Conference date: 30-04-2018 Through 03-05-2018.

Marquez-Gallardo, S., and Rovira-Lorente, A. 2024. In genderqueer closets: Challenging gender binarism through embodied narratives of affect and style. *Fashion Theory* 1–24.

McKercher, K. 2020. *Beyond sticky notes: Co-design for real: Mindsets, methods and movements*. Sydney, NSW, Australia: Beyond sticky notes.

Nelson, A. G. 2022. *Creating safe and inclusive costuming spaces*. American University.

Razavi, A.; van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with vq-vae-2. In *Neural Information Processing Systems*.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models.

Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2022. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 22500–22510.

Saovana, N., and Khosakitchalert, C. 2024. Assessing the viability of generative ai-created construction scaffolding for deep learning-based image segmentation. In *2024 1st International Conference on Robotics, Engineering, Science, and Technology (RESTCON)*, 38–43.

Thanh-Tung, H., and Tran, T. 2020. Catastrophic forgetting and mode collapse in gans. In *2020 international joint conference on neural networks (ijcnn)*, 1–10. IEEE.

Utz, V., and DiPaola, S. 2023. Climate implications of diffusion-based generative visual ai systems and their mass adoption. In *International Conference on Innovative Computing and Cloud Computing*.

Vigliensoni, G.; Perry, P.; and Fiebrink, R. 2022. A small-data mindset for generative ai creative work. In *Generative AI and HCI - CHI 2022 Workshop*.

von Platen, P.; Patil, S.; Lozhkov, A.; Cuenca, P.; Lambert, N.; Rasul, K.; Davaadorj, M.; Nair, D.; Paul, S.; Liu, S.; Berman, W.; Xu, Y.; and Wolf, T. 2023. Diffusers: State-of-the-art diffusion models.

Wu, Z.; Ji, D.; Yu, K.; Zeng, X.; Wu, D.; and Shidujaman, M. 2021. Ai creativity and the human-ai co-creation model. In Kurosu, M., ed., *Human-Computer Interaction. Theory, Methods and Tools. HCII 2021. Lecture Notes in Computer Science*, volume 12762 of *Lecture Notes in Computer Science*. Springer, Cham. 171–190.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* 3813–3824.

Zhang, Y.; Zhang, T.; and Xie, H. 2024. Texcontrol: Sketch-based two-stage fashion image generation using diffusion model.