Speculative Co-design With AI: An Artist-Friendly Prototype for Non-Human Avatar Creation

Mengzhi He*
Creative Computing Institute,
University of the Arts London
London, Uk
h.he0520241@arts.ac.uk

Stella Doukianou
Creative Computing Institute,
University of the Arts London
London, Uk
s.doukianou@arts.ac.uk

Nick Bryan-Kinns
Creative Computing Institute,
University of the Arts London
London, Uk
n.bryankinns@arts.ac.uk

Abstract

This technical demonstration responds to the growing importance of virtual identity. It introduces an artist-friendly, no-code prototype designed specifically to transparently integrate speculative AI processes into artistic co-design methods focused on virtual identity exploration. Artists and designers often struggle to effectively integrate AI into their creative practice because packaged commercial AI tools rarely accommodate their unique visual language and generally obscure creative control through black-box interfaces. These tools typically lack transparency for advanced users and are inconvenient when managing diverse creative inputs, especially in participatory co-design contexts involving multiple contributors. Addressing these challenges, this demonstration showcases a prototype developed using ComfyUI, a visual, node-based interface for building generative AI workflows without coding, combined with custom-trained LoRA models to ensure visual consistency and personal artistic expression in participant-driven avatar generation. Grounded in literature emphasising AI's potential to foster imaginative engagement through transparency and creative flexibility, this transparent AI tool supports co-design activities, encouraging a wider community of creative practitioners to confidently experiment with speculative AI-driven co-design.

CCS Concepts

• **Human-centered computing** → *Visualization systems and tools*; *Visualization systems and tools*; • **Social and professional topics** → Cultural characteristics; Cultural characteristics.

Keywords

Generative AI, Co-design, Virtual identity, No-code tools, Non-human avatars, Speculative design, Practice-based research

ACM Reference Format:
Mengzhi He, Stella Doukianou, and Nick Bryan-Kinns. 2025. Speculative Co-design With AI: An Artist-Friendly Prototype for Non-Human Avatar Creation. In *Creativity and Cognition (C&C '25), June 23–25, 2025, Virtual, United Kingdom.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3698061.3734414

*Corresponding author.

C&C '25, Virtual, United Kingdom

ACM ISBN 979-8-4007-1289-0/2025/06
https://doi.org/10.1145/3698061.3734414

1 Introduction & Aims

Co-design is widely employed in participatory design and research, integrating the perspectives and creativity of multiple stakeholders to develop solutions that are inclusive, practical, and sustainable [8, 14, 23]. In studies focusing on non-human avatars and virtual identities, enhancing participants' agency is especially critical. Participants in these contexts are not merely information providers but co-creators whose experiences and knowledge systematically inform the design outcomes, fostering greater identification with the final results [8, 13, 24, 25].

The identification process between users and non-human avatars often involves deep emotional attachment and projections of an idealised self, such as representing aspirational identities, hidden personalities, or cultural symbols [2, 10, 11, 18]. Traditional research methods like surveys or experiments often fall short in effectively capturing these subjective experiences, making co-design an optimal research approach.

However, the co-design methods commonly used by artists and designers, such as sketching, collage, or verbal discussions, often fail to clearly visualise diverse participant expectations in real-time. This limitation is particularly evident in collaborative settings, where synthesising multiple participants' unique interpretations into a unified visual language poses significant challenges. Additionally, the absence of immediate visual feedback during workshops further hinders effective communication and mutual understanding. Generative AI technologies offer effective solutions to these challenges.

Despite growing interest in integrating generative AI into artistic practices, substantial challenges remain [1, 9, 27]. Commercial generative AI tools (e.g. Midjourney [16] and DALL·E 2 [20]) typically use closed and opaque interfaces that obscure the generative process, limiting artists' creative autonomy [3]. Such 'black-box' systems restrict artists from fully expressing their distinctive visual language and complicate collaborative processes involving diverse creative inputs.

In response, this technical demonstration introduces an artist-friendly, no-code prototype that transparently integrates traceable AI tools into co-design processes, supporting participatory, speculative, and identity-focused design research. The prototype is built upon ComfyUI [4] , a visual node-based interface clearly revealing and allowing direct manipulation of generative AI processes—and custom-trained LoRA [7] models, providing artists with precise control over outputs and ensuring visual consistency among participants. This method is particularly suited to real-time co-creation of non-human digital avatars, prominent within contemporary digital subcultures (e.g., furry fandom [22]), virtual social platforms (e.g.,

VRChat [26]), and virtual influencers on social media, all significantly influencing identity perceptions and self-expression among younger audiences.

This demonstration explicitly positions AI as a transparent, co-design partner, distinguishing itself from traditional commercial black-box AI tools. By enhancing transparency and artistic control, the tool empowers artists to confidently experiment, ensuring generated outcomes closely align with their unique aesthetic visions. Consequently, this approach significantly lowers barriers to creative engagement, fostering deeper and more meaningful exploration of AI-driven virtual identity creation and practice.

Developed as a prototype within practice-based research, this design method allows artists and designers to treat AI not as a black-box tool but as a transparent tool that supports co-design. In contrast to mainstream AI tools such as Midjourney [16] and DALL·E 2 [20], which prioritise ease of use by concealing complex generative parameters [3], this prototype offers clear and flexible access to the generative process without requiring technical coding skills.

This approach is grounded in literature indicating that generative AI tools significantly expand visual imagination, lower barriers to creative participation, and increase the confidence of non-experts in engaging with visually complex and abstract forms [1, 15, 21]. By integrating speculative AI processes at the early conceptual stages of research-through-design, this demonstration supports reflective engagement and encourages creative practitioners to explore new possibilities for collaborative and speculative identity inquiry.

2 Architecture & Implementation

2.1 Prototype Overview

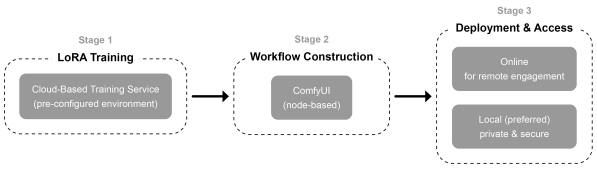

Figure 1: Overview of the prototype's architecture showing the three development stages: LoRA training, workflow construction in ComfyUI, and deployment either locally or online. Each stage supports artist-friendly interaction and control.

Figure 1 illustrates the three key stages of the proposed prototype. This prototype has been fully constructed and tested by the first author, demonstrating the feasibility and artistic flexibility of the proposed method, although it has not yet been evaluated through formal participant studies.

In the first stage, a no-code method for training custom LoRA [7] models are introduced. These allow artists to embed their own visual language into the generative AI process without any programming experience. The second stage involves building a tailored workflow in ComfyUI [4], which is a visual interface where model logic, textual prompts, and outputs are arranged using an intuitive node-based system. In the final stage, the prototype can be deployed locally on a personal device to support secure, private use in face-to-face settings, or accessed through an online platform when engaging with remote participants. Each of these stages is described in more detail in the following sections.

2.2 Stage 1: LoRA Training

In Stage 1, the system trains a lightweight LoRA model that reflects the artist's unique visual language. Artists can choose between a cloud-based, no-code training service (LiblibAI [12] is used for this demonstration) and a local training approach for sensitive content. Full details—including image curation, parameter settings, and model generation—are provided in the supplementary document.

2.3 Stage 2: Workflow Construction

Stage 2 involves constructing a visual, node-based workflow in ComfyUI that integrates Stable Diffusion, ControlNet, and the custom LoRA model. The workflow is saved as a JSON file that encapsulates model checkpoints, text prompts, and processing logic. The system offers a simplified interface for general users as well as adjustable parameters for advanced users. Detailed node configurations and system logic are available in the supplementary document.

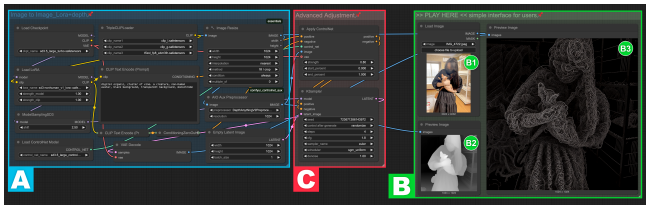

Figure 2: Visual layout of the proposed ComfyUI workflow showing its modular structure. Section A (blue) contains the core AI logic, Section B (green) provides a simplified upload and preview panel for general users, and Section C (red) offers adjustable parameters for intermediate users seeking more control.

As illustrated in Figure 2, the ComfyUI workflow is organised into three clearly defined sections:

Group A: Core Architecture and Model Logic (Blue Section). This section contains foundational elements of the workflow including base model, LoRA style injection, activator words, ControlNet setup and additional utility nodes. They are structured to remain hidden from general users but modifiable by experienced artists. These elements are foundational to maintaining visual consistency and generating repeatable outputs aligned with specific aesthetic goals.

Group B: Simplified User Interface (Green Section). Designed for users with minimal technical knowledge, this group includes intuitive controls such as upload buttons and output previews. It enables novice participants to engage with the system without being overwhelmed by complex parameters. This design decision supports participatory workshops where accessibility and immediacy are crucial. B1 is uploading panel; B2 is depth map preview; B3 is the output preview.

Group C: Adjustable Parameters for Intermediate Users (Red Section). This section provides deeper artistic control through adjustable values such as combinations of Classifier-Free Guidance [17] (CFG) scale and sampling steps. It supports users with moderate experience in visual design or generative workflows, enabling them to experiment creatively while maintaining system structure. This tiered approach allows facilitators to tailor interactions to participants' skill levels, encouraging engagement across experience levels.

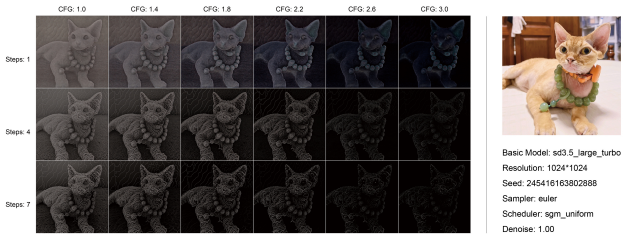

Figure 3: Comparison of generated avatar results using different combinations of Classifier-Free Guidance (CFG) scale and sampling steps. The same input image (bottom right) was used across all generations. The visual differences demonstrate how higher CFG and step values impact stylisation, detail clarity, and contrast.

As shown in Figure 3, During testing, the first author identified that setting 4 sampling steps with a Classifier-Free Guidance (CFG) scale of approximately 1.8 provided an optimal balance between stylistic consistency, detail clarity, and colour contrast.

For a comprehensive explanation of the technical requirements and the detailed construction of the workflow, please refer to the supplementary document submitted alongside this demo.

2.4 Performance & Flexibility

The proposed workflow was tested on two different local devices to assess its responsiveness and suitability for co-design workshops where real-time generation is essential.

- Desktop PC (GeForce RTX 3090, 32GB RAM, Windows 10 X64) generated a 1024×1024 image in approximately 36 seconds.
- MacBook Pro (M4 MAX, 128GB RAM, macOS 15.4) completed the same task in 45 seconds.

While both setups are capable of running the workflow efficiently, the NVIDIA RTX GPU environment provides faster inference, which is expected since the ecosystem around Stable Diffusion and ComfyUI is largely optimised for CUDA [19]. This highlights that while this approach is accessible across platforms, participants using NVIDIA GPUs may experience smoother performance, especially in real-time or iterative workshop settings.

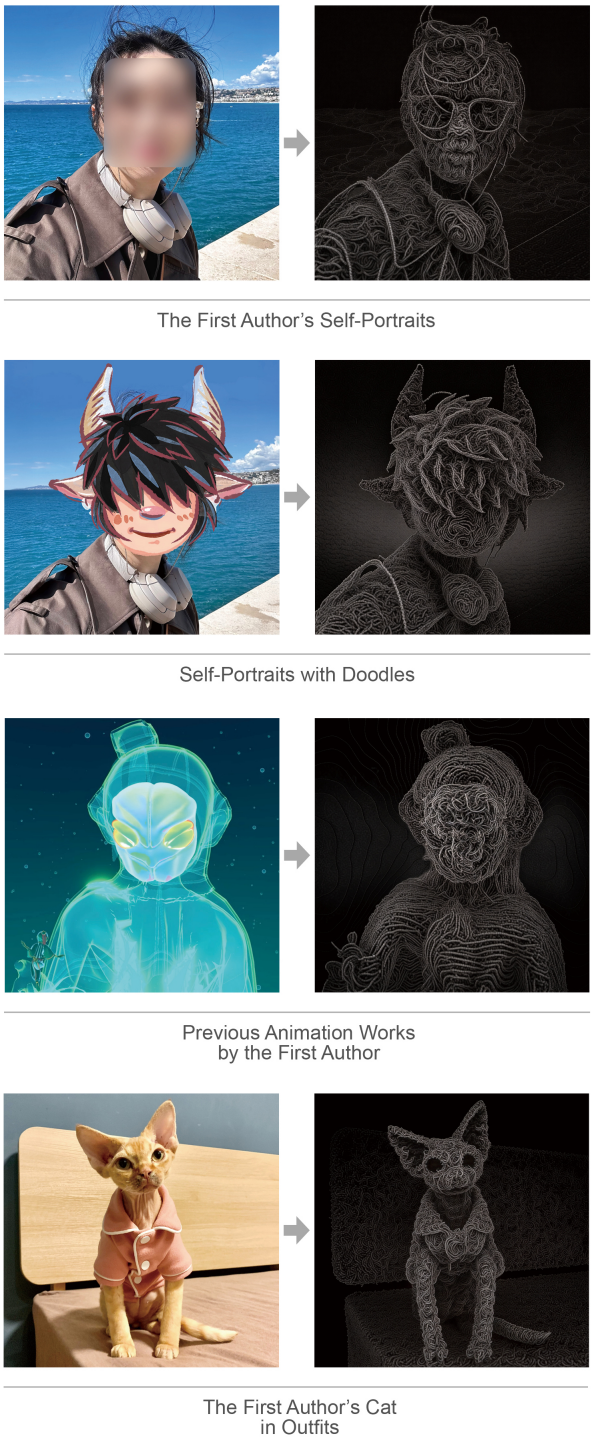

Figure 4: Examples of input image types and their generated outputs using the proposed workflow. All images were processed using the same LoRA model, demonstrating its robustness across a wide range of styles.

As shown in Figure 4, the workflow was tested using a wide range of input types, including realistic self-portraits, hand-drawn doodled additions, digital artworks, and non-human subjects. In every case, the outputs maintained a consistent and visually coherent aesthetic, demonstrating the LoRA model's robustness across diverse inputs. All generated images can be interpreted as non-human avatars. This is achieved through abstract transformation, where key biological features are replaced or reinterpreted using textured, organic visual elements. These forms encourage participants to imagine new bodies and identities beyond the human norm.

These tests demonstrate that the workflow supports diverse artistic media, from digital images and drawings to animations and documentation of physical art, enabling artists to seamlessly integrate AI into their practice. Its broad format support and no-code interface empower creators to experiment freely while maintaining stylistic cohesion.

In future use, the prototype can support personal and speculative identity exploration [21]. Participants might begin by uploading an unedited self-portrait, or creatively modifying their own image—for example, painting in animal ears, longer hair, or imagined features. These modified portraits are then transformed through the workflow into non-human avatars, offering an opportunity to explore alternative expressions of self and reimagine the boundaries of the body.

2.5 Stage 3: Deployment & Access

Once constructed, the workflow is deployed for demonstration in two modes: a local deployment that provides secure, private use, and an online Dockerized version hosted on Hugging Face Spaces for remote exploration. The online version [6] allows users to inspect the ComfyUI interface and download the JSON workflow file [5]. Comprehensive deployment instructions and dependency details are included in the supplementary document.

3 Discussion & Conclusion

This technical demonstration presents a transparent, no-code design prototype that addresses pivotal challenges in creative AI applications, specifically designed for artists and designers without technical expertise. Within co-design design contexts, these considerations become particularly salient when dealing with abstract visual styles and speculative identity representation.

For demonstration purposes, an on-line Docker version [6] has been deployed via Hugging Face Spaces. While platform GPU limitations preclude real-time image generation, participants can fully inspect the ComfyUI interface - examining node architectures, input-output configurations, and employed models/modules. To enable hands-on experimentation, we provide downloadable JSON workflow files for local GPU-enabled ComfyUI implementation of non-human avatar generation.

Unlike opaque commercial platforms, our prototype facilitates experimentation aligned with individual aesthetic objectives. By framing AI as a tool that supports co-design, it fosters not only technical engagement but also identity-driven speculative exploration. As a design prototype, this approach offers adaptable scaffolding rather than definitive solutions - creating conceptual space to investigate how AI can support preliminary design inquiry through controlled experimentation and accessible interaction, including speculative narratives, non-human embodiment and digital self-expression.

While currently focused on non-human avatars, the framework remains extensible to broader applications such as self-perception and identity construction studies. Future research could examine how transparent generative processes influence creative decision-making and emotional identification, potentially employing theoretical lenses like the Proteus Effect [18]. Ultimately, this work contributes a practical, accessible and reflective instrument to the field of AI-augmented co-design.

4 Responsible Practice Statement

All visual content (e.g., photos, artwork) is original and owned by the first author, with no external data used. This demonstration explicitly showcases a technical prototype focused on transparency and creative control within AI-assisted artistic prototype, rather than participant-generated outcomes. Any future co-design studies or broader participatory applications will undergo rigorous ethical review and institutional approval prior to deployment.

References

[1] Oliver Hauser Anil Rajnikant Doshi. 2023. Generative Artificial Intelligence Enhances Creativity. *SSRN Electronic Journal* (2023). doi:10.2139/ssrn.4535536

[2] Laura Aymerich-Franch. 2012. Can we identify with a block? Identification with non-anthropomorphic avatars in virtual reality games. In *Proceedings of the international society for presence research annual conference*, Vol. 6.

[3] Gregory Thomas Croisdale, John Joon Young Chung, Emily Huang, Gage Birchmeier, Xu Wang, and Anhong Guo. 2023. DeckFlow: A Card Game Interface for Exploring Generative Model Flows. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 89, 3 pages. doi:10.1145/3586182.3615821

[4] ComfyUI Developers. 2024. *Getting Started - Introduction.* Retrieved April 4, 2025 from https://docs.comfy.org/get_started/introduction

[5] Mengzhi He. 2025. *non_human_abstract.json Workflow File.* Retrieved April 4, 2025 from https://huggingface.co/spaces/haha9771/nonHuman_Demo/blob/main/non_human_abstract.json

[6] Mengzhi He. 2025. *nonHuman_Demo: Speculative Co-Design Prototype on Hugging Face Spaces.* Retrieved April 4, 2025 from https://huggingface.co/spaces/haha9771/nonHuman_Demo

[7] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. *arXiv preprint arXiv:2106.09685* (2021). https://arxiv.org/abs/2106.09685

[8] Toni Robertson Jesper Simonsen. 2012. *Routledge International Handbook of Participatory Design.* Routledge, Routledge.

[9] Christian Remy Michael Mose Biskjaer Peter Dalsgaard Jonas Frich, Lindsay MacDonald Vermeulen. 2019. Mapping the Landscape of Creativity Support Tools in HCI. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (May 2019), 1–18. doi:10.1145/3290605.3300619

[10] Dominic Kao. 2019. The effects of anthropomorphic avatars vs. non-anthropomorphic avatars in a jumping game. *Proceedings of the 14th International Conference on the Foundations of Digital Games* (Aug 2019), 1–5. doi:10.1145/3337722.3341829

[11] Connor Leshner. 2017. Theories and implications of the Proteus Effect on cosplay. *The Phoenix Papers* 3, 1 (2017), 46–55.

[12] LiblibAI. 2025. *Training Walkthrough Documentation (Feishu Internal).* Retrieved April 4, 2025 from https://zdorongbe5.feishu.cn/docx/QiHzd8wW7o5eecxQ0AOc55uLn1e

[13] Ann Light and Yoko Akama. 2012. The human touch: participatory practice and the role of facilitation in designing with communities. In *Proceedings of the 12th Participatory Design Conference: Research Papers - Volume 1* (Roskilde, Denmark) *(PDC '12)*. Association for Computing Machinery, New York, NY, USA, 61–70. doi:10.1145/2347635.2347645

[14] Ezio Manzini. 2015. *Design, When Everybody Designs.* MIT Press, MIT Press.

[15] June Kyong Trondsen Mathias Østvold Ek, Magnus Paulsen. 2024. SPECULATIVE DESIGN THROUGH THE LENS OF AI. *Proceedings of the International Conference on Engineering and Product Design Education, EPDE 2024* (2024), 627–632. doi:10.35199/epde.2024.106

[16] Midjourney. 2022. *Midjourney Home.* Retrieved April 4, 2025 from https://www.midjourney.com/home

[17] Alex Nichol and Prafulla Dhariwal. 2021. Improved Denoising Diffusion Probabilistic Models. *arXiv preprint arXiv:2102.09672* (2021).

[18] Nicolas Ducheneaut Nick Yee, Jeremy N. Bailenson. 2009. The Proteus Effect. *Communication Research* 36, 2 (Apr 2009), 285–312. doi:10.1177/0093650208330254

[19] NVIDIA. 2024. *CUDA Toolkit Documentation.* Retrieved April 4, 2025 from https://docs.nvidia.com/cuda/index.html

[20] OpenAI. 2022. *DALL·E 2 Overview.* Retrieved April 4, 2025 from https://openai.com/index/dall-e-2/

[21] Daniel Pillis, Pat Pataranutaporn, Pattie Maes, and Misha Sra. 2024. AI Comes Out of the Closet: Using AI-Generated Virtual Characters to Help Individuals Practice LGBTQIA+ Advocacy. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) *(IUI '24)*. Association for Computing Machinery, New York, NY, USA, 686–698. doi:10.1145/3640543.3645213

[22] Courtney N Plante, Stephen Reysen, Camielle Adams, Sharon E Roberts, and Kathleen C Gerbasi. 2023. *Furscience: A Decade of Psychological Research on the Furry Fandom.* Stephen Reysen.

[23] Elizabeth B.-N. Sanders and Pieter Jan Stappers and. 2008. Co-creation and the new landscapes of design. *CoDesign* 4, 1 (2008), 5–18. doi:10.1080/15710880701875068 arXiv:https://doi.org/10.1080/15710880701875068

[24] Ina Wagner Tone Bratteteig. 2016. Unpacking the Notion of Participation in Participatory Design. *Computer Supported Cooperative Work (CSCW)* 25, 6 (Dec 2016), 425–475. doi:10.1007/s10606-016-9259-4

[25] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 429–438. doi:10.1145/2470654.2470716

[26] VRChat. 2024. *Community Guidelines.* Retrieved April 4, 2025 from https://hello.vrchat.com/community-guidelines

[27] Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G. Michael Youngblood. 2018. Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. In *2018 IEEE Conference on Computational Intelligence and Games (CIG)*. 1–8. doi:10.1109/CIG.2018.8490433