Baby Lottery: An Interactive Demonstration of AI Bias in Generative Imagery

Cindy H Lin
Creative Computing Institute
London, United Kingdom
c.lin@arts.ac.uk

Rebecca Fiebrink
Creative Computing Institute
London, United Kingdom
r.fiebrink@arts.ac.uk

Abstract

The fertility industry has rapidly adopted AI as both a clinical and marketing tool. However, the biases embedded within these systems remain largely invisible to the people making intimate decisions based on the outputs of the AI models. Baby Lottery is an interactive demonstration that invites conference attendees to generate AI images of a hypothetical baby by combining their own photograph with a randomly selected AI-generated donor profile created from a dataset of over 1,000 sperm donor profiles collected from U.S. and European sperm banks. The results are printed on a thermal receipt printer, producing an immediate physical artifact. By making the generation process participatory and its outputs tangible, the demonstration invites critical examination of how AI systems encode bias in deeply personal contexts. This work extends the investigation of AI bias propagation in fertility marketing documented in our recent work, and engages attendees as active subjects in the ongoing inquiry.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Social and professional topics** → *User characteristics*; • **Computing methodologies** → Artificial intelligence.

Keywords

generative AI, AI bias, participatory art, fertility technology, critical design, reproductive technology

ACM Reference Format:
Cindy H Lin and Rebecca Fiebrink. 2026. Baby Lottery: An Interactive Demonstration of AI Bias in Generative Imagery. In *Creativity and Cognition (C&C '26), July 13–16, 2026, London, United Kingdom.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3803784.3816883

1 Introduction and Background

Field research conducted at fertility fairs and clinics in 2019 and 2024 [14] documents how the marketing of the fertility industry has shifted to a more polished presentation boasting the latest technology in AI. What were once thick, white physical binders of black-and-white donor profiles that sperm recipients had to sift through by hand are now digital platforms with instant filtering, AI-enhanced recommendation systems, and donor face-matching tools that let prospective parents upload photos of celebrities, friends, or strangers to find a genetic resemblance [10], with premium upsells to unlock features such as donors' handwritten letters, audio interviews, and adult photos. AI is now marketed as a technological achievement that can rapidly identify high-quality sperm [13] and optimize donor selection [7]. Meanwhile, consumer apps widely market the ability to visualize future children using AI [15].

What is less visible is what these systems actually do. Field observations suggest that much of the AI being promoted by fertility clinics amounts to AI-generated stock photography used in marketing materials [14]. Fairfax Cryobank, one of the largest sperm banks in the world [12], has trademarked a tool called FaceMatch, which claims to allow recipients to find sperm donors with similar facial structures by uploading a photo of any gender [10]. Fertility clinics leverage these visualizations to influence decisions about family formation, choices that are among the most intimate and consequential a person can make.

The commercialization of reproductive technology has created an industry in which genetic material is filtered, ranked, and marketed using the same logic as consumer products. The global assisted reproductive technologies (ART) market shows significant expansion, hovering around USD 26 billion in 2019 and was expected (in 2023) to reach USD 45 billion by 2025 [1]. Sperm bank databases now include detailed genetic profiles, personality assessments, childhood and adult photographs, voice samples, and demographic information, all presented through interfaces designed to facilitate consumer choice [11]. Access is tiered by cost: while recipients can view limited free data on donors, premium packages range from $250 to $499, bundling genetic data, personal media, and sperm vials as consumer products [9, 11].

AI systems trained on historically biased datasets do not simply reproduce existing data: they amplify and standardize outputs [5]. When applied to the generation of human facial imagery, these systems consistently erase racial and ethnic diversity in favor of dominant Western representational norms [3]. Recent audits of current-generation models have demonstrated that semantically neutral prompts consistently produce outputs skewed toward lighter skin tones and Western features [2]. Although the current pipeline relies on SDXL Turbo [17], a 2023 model, and required directive prompts to achieve coherent output quality, bias manifesting consistently across both semantically neutral prompts [2] and the directive prompts required by this ComfyUI pipeline suggests the bias operates independently of prompt specificity and persists across model generations.

This pattern echoes broader criticisms of computational sciences that perpetuate Euro- and Anglo-centric ideologies through data

C&C '26, London, United Kingdom

ACM ISBN 979-8-4007-2583-8/26/07
https://doi.org/10.1145/3803784.3816883

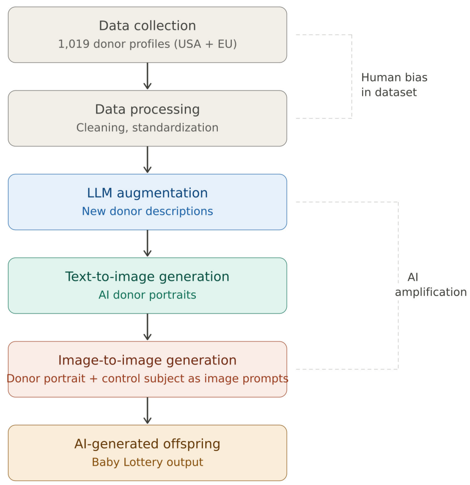

Figure 1: Multi-stage AI pipeline for Baby Lottery. Bias enters through human-created donor profiles and is amplified through successive AI generation stages. The findings were consistent: regardless of donor characteristics and ethnicities, AI-generated offspring defaulted to lighter skin tones, Western facial features, and commercial beauty conventions [14].

collection, classification, and categorization [4]. In the fertility industry specifically, the foundational dataset is already shaped by commercial logic: donor profiles are written as marketing copy, and some clinics have been documented encouraging donors to adjust self-reported information to increase demand [14]. The bias, in other words, begins before any AI model is trained.

Prior creative work has examined how computational systems interpret and generate human genetic information. Heather Dewey-Hagborg's Stranger Visions (2013) reconstructed facial composites from DNA found in public spaces, interrogating both the claims and the limits of biological prediction [8]. Anna Ridler's Mosaic Virus (2018) foregrounded the hidden labor of AI training data, making visible the human effort that machine learning typically conceals [16].

Lin's recent thesis work [14] situates itself within this lineage, using the fertility industry as a lens through which to make AI bias legible and experiential. That investigation created a dataset of 1,019 sperm donor profiles from U.S. and European sperm banks, then fed the dataset through multiple stages of AI processing: from commercial large language models generating new donor descriptions, to text-to-image models generating donor portraits, to image-to-image models generating hypothetical offspring with a single control subject's photograph as input.

2 System Overview and Attendee Experience

The demonstration uses ComfyUI, a node-based interface for Stable Diffusion, running locally on a laptop. The image generation pipeline combines three models: SDXL Turbo for single-step image generation [17], ControlNet MediaPipeFace for facial feature control [6], and IP-Adapter for image prompt blending without modifying the base model. We utilized IP-Adapter [18] to perform image-conditioned generation. Unlike standard image-blending techniques, IP-Adapter allows the participant's facial features to act as a visual prompt, guiding the diffusion process of the donor portrait without the need for computationally expensive fine-tuning or DreamBooth training. The current workflow is being refined to ensure image inputs are weighted appropriately against text prompts in the generation pipeline. No bias mitigation techniques were applied to the pipeline. This is an intentional design choice: the demonstration is not an attempt to correct for bias, but to make its unmitigated propagation visible and experiential.

This demonstration invites attendees to step into that process themselves one at a time. The attendee experience proceeds as follows. After reviewing the participant consent form, participant sits down at the demo station and takes a photograph using the laptop webcam. The system randomly selects a donor portrait from a pre-existing pool of AI-generated donors created during Lin's recent work [14]. The participant's photo and the selected donor portrait are fed into the ComfyUI workflow, which produces a generated offspring image. The entire generation process takes approximately 30 to 45 seconds. Once complete, the output is printed on a thermal receipt printer, producing a small physical printout the participant keeps. The receipt format mirrors the transactional logic of the fertility industry and produces a physical artifact that can be compiled as physical documentation of their participation.

The physical setup consists of a table with a laptop, a webcam, and a thermal receipt printer. A poster documenting the pairings between a standardized control subject and various donor profiles will be displayed alongside the demo station, presenting results in a grid format similar to a catalog. This provides immediate visual context and enabling comparison for participants.

In the recent work by Lin [14], a single control subject's photograph served as the sole input across all donor pairings to establish a baseline. For this demonstration, conference attendees replace the baseline subject as the input. Each attendee's photo and a randomly assigned donor portrait are weighted equally as image prompts, standardizing the output for comparison across a diverse participant pool. This shift from a single-subject baseline to multiple participants allows patterns of bias to be observed across a wider range of skin tones, gender, facial structures, and ethnic backgrounds.

2.1 Participant Consent

Before participating, attendees will be briefed on what the demonstration involves and asked to give written consent. Participants may opt out at any point. No photographs of participants will be stored on the laptop after the generation process is complete. The generated offspring images will only be retained if participants give explicit consent. Participants will also be asked whether they consent to their generated printout being included in a compiled zine documenting the demonstration results. Art supplies will be provided for those who want to contribute their reflections into the compiled zine. Those who decline will still receive their printout to keep.

Figure 2: Representative sample of synthetic offspring generated in Lin's recent work [14]. Despite diverse donor inputs, the outputs consistently exhibit a standardized commercial aesthetic, characterized by homogenized facial features and studio-style lighting. All images shown are AI-generated and do not depict real infants.

Prior to generation, participants are informed that outputs will likely reflect the model's dominant aesthetic conventions regardless of their own features or the donor's characteristics. This briefing frames the result as evidence of systemic bias rather than a meaningful prediction of genetic outcome. Following the generation, participants are invited to reflect on their experience through prompt cards, which ask how the output compares to their expectations and what they notice about the result. This debrief structure is designed to convert individual outputs into collective evidence, reinforcing the demonstration's function as participatory research rather than personal prediction.

3 Discussion

The interactive format enables participants to experience bias firsthand. Rather than presenting findings as static documentation, the demonstration makes the propagation of bias from a human-generated dataset to AI-generated outputs something participants can see happening to their own image. By having the poster available for immediate comparison, participants can observe how the same visual tropes repeat regardless of the ethnicity and gender of the user input.

Lin [14] documented a consistent pattern: AI-generated offspring defaulted to lighter skin tones, blue eyes, and Western features regardless of the diversity of donor inputs. Attendees at the demo will generate new data points in real time across a diverse participant pool, extending the single-subject baseline of Lin [14] to a wider range of skin tones, facial structures, and ethnic backgrounds. Because the conference will take place approximately 18 months after Lin's original work [14], the demonstration also functions as an informal comparative analysis of whether bias patterns persist in current model versions.

Through active participation, conference attendees become collaborators in an ongoing investigation rather than passive viewers of a fixed result. The demonstration also surfaces questions about consent and the ethics of using one's own image as input to a generative system. By placing attendees in the position of the research subject rather than the observer, the work reframes the typically abstract discourse around AI bias as a personal encounter. Future iterations will develop the project into a larger interactive installation examining how technology and legislative frameworks together constrain reproductive autonomy, with outputs varying by the participant's selected country of origin to reflect how reproductive rights differ across legal contexts.

4 Conclusion

Baby Lottery demonstrates how creative practice can function as a mode of visual research, making systemic bias experiential rather than merely legible. By positioning conference attendees as active participants rather than observers, the demonstration generates comparative data across a diverse participant pool in real time while opening critical questions about consent, representation, and the role of AI in mediating intimate decisions. As AI systems become further embedded in reproductive contexts, work that makes their assumptions visible becomes increasingly necessary.

References

[1] Joao Felipe Aderaldo, Beatriz Helena Dantas Rodrigues De Albuquerque, Maria Teresa Fialho Câmara De Oliveira, Mariana De Medeiros Garcia Torres, and Daniel Charles Ferreira Lanza. 2023. Main topics in assisted reproductive market: A scoping review. *PLOS ONE* 18, 8 (2023), e0284099. doi:10.1371/journal.pone.0284099

[2] Roberto Balestri. 2026. Neutral prompts, non-neutral people: Quantifying gender and skin-tone bias in Gemini Flash 2.5 image and GPT image 1.5. arXiv preprint arXiv:2602.12133. https://arxiv.org/abs/2602.12133

[3] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. ACM, New York, NY, USA, 1493–1504.

[4] Abeba Birhane and Olivia Guest. 2020. Towards decolonising computational sciences. arXiv preprint arXiv:2009.14258. https://arxiv.org/abs/2009.14258

[5] Kate Crawford and Trevor Paglen. 2021. Excavating AI: the politics of images in machine learning training sets. *AI & Society* 36, 4 (2021), 1105–1116.

[6] CrucibleAI. 2023. ControlNetMediaPipeFace. Hugging Face model repository. https://huggingface.co/CrucibleAI/ControlNetMediaPipeFace

[7] Cyprus IVF Centre. 2025. AI assisted sperm selection. Web page. https://cyprusivfcentre.com/ai-assisted-sperm-selection/

[8] Heather Dewey-Hagborg. 2013. Stranger Visions. Project website. https://deweyhagborg.com/projects/stranger-visions

[9] Danny Elliot. 2024. You wouldn't believe how difficult it is to buy sperm. *The Guardian* (2024). https://www.theguardian.com/lifeandstyle/article/2024/may/28/buying-sperm-fertility-clinic

[10] Fairfax Cryobank. 2025. Fairfax FaceMatch. Web page. https://fairfaxcryobank.com/fairfax-facematch

[11] Fairfax Cryobank. 2026. Donor Unlimited Access Packages. Web page. https://fairfaxcryobank.com/donor-unlimited-access-packages

[12] Fairfax Cryobank. 2026. Fairfax Cryobank: Find your ideal sperm donor. Web page. https://fairfaxcryobank.com/

[13] Jacqueline Howard. 2025. A couple tried for 18 years to get pregnant. AI made it happen. *CNN* (2025). https://edition.cnn.com/2025/07/03/health/ai-male-infertility-sperm-wellness

[14] Cindy Lin. 2024. *Baby Lottery: Examining Bias Propagation in AI Generative Models*. MSc Thesis. University of the Arts London, Creative Computing Institute.

[15] Overchat AI. 2025. What is the best baby generator app in 2025/2026? Web article. https://overchat.ai/ai-hub/best-baby-generator-apps

[16] Anna Ridler. 2018. Mosaic Virus. Art project. https://annaridler.com/works/mosaic-2018

[17] Stability AI. 2025. SDXL-Turbo. Hugging Face model repository. https://huggingface.co/stabilityai/sdxl-turbo

[18] Hu Ye, Jun Zhang, Sibo Liu, Xiao Han, and Wei Yang. 2023. IP-Adapter: Text Compatible Image Prompt Adapter for Text-to-Image Diffusion Models. arXiv preprint arXiv:2308.06721. https://arxiv.org/abs/2308.06721