Sketch-It: A Consent-First Creativity Support Tool for the Generation of Inspirational Sketches

Clàudia Pàmies
Linnaeus University, Sweden &
Universitat Autònoma de Barcelona, Spain
claudia.pamies@autonoma.cat

Garrit Schaap
Computer Science and Media Technology
Linnaeus University, Sweden
garrit.schaap@lnu.se

Maria M. Hedblom
Jönköping AI lab
Jönköping University, Sweden
maria.hedblom@ju.se

Abstract

Creativity support tools (CST) are designed to offer inspiration and help artists overcome creative block. For this, Generative AI tools offer quick and accessible methods to generate images, but are not adjusted to individual style, and they raise ethical concerns due to the copyright of the training data. Addressing these issues, we introduce "Sketch-It", a consent-first CST based on two generative adversarial networks (DCGANs) that is exclusively trained on the user's own artwork. To provide patterned structure, the first model is trained on greyscale sketches, while the second model is trained on photographs to acquire colour knowledge. The two models are bridged by a U-Net colouriser that applies the learned colour distributions to the structural forms. Intended for personal use, the system is packaged as a browser-based local application, requiring no coding knowledge for artists to train the system on their own artwork and use it in their creative process. We demonstrate that the system can generate abstract inspirational images by training the system on one artist's flower sketches and photographs. Our system offers three distinct contributions: technical – a dual-model form/colour separation architecture, ethical – a consent-first training paradigm on minimal datasets, and practical – an accessible tool designed for artists.

Introduction

Creative block is defined as the temporary inability to generate new ideas or initiate new work. It is a well-documented phenomenon among practising artists [3] that is independent of skill and defined as a momentary disconnection from one's own creative vocabulary and creative inspiration. In light of the growing involvement of technology in creative processes, creativity support tools (CSTs) aim to assist in rather than replace human creative capacity, by offering, for example, prompts, references or generative suggestions that reset the creative process back in motion [23].

Contemporary AI image generation systems (e.g., Stable Diffusion [15], Midjourney [12], DALL-E [14]) are superficially well placed to serve this need as they can rapidly generate images in a wide variety of styles based on textual descriptions. However, trained on other data, the generated images might substantially diverge from the artist's own style, leading to distractions in the creative process or a lack of genuine inspiration. Additionally, their adoption by the creative community has been met with significant resistance for principled reasons [10], since these systems are trained on vast datasets assembled via web scraping and, thus, contain the collected result of millions of copyrighted artworks that have been used without the knowledge or consent of their creators [21].

Our approach takes a different route by asking *"What if the AI system learned only from you?"* This addresses the style problem by ensuring the AI follows your uniquely expressed style while also removing any ethical issues of 'artistic theft.' However, it also introduces a new challenge: how to generate inspirational outputs of good enough quality with limited datasets.

To this end, we introduce "Sketch-It," a system that trains and combines the results of two small generative models, both of which the user has complete control over the training data. Thus, ensuring that exclusively their own work gets to influence the results. The system is presented as a browser-based interface, which requires no programming knowledge. The result is a creativity-support tool that is simultaneously ethical (fully consented data), computationally green [2] – due to the small datasets and local CPU training, and personally meaningful – the generated output is an extension of the artist's work. To demonstrate Sketch-It, we will present a prototype example trained on a single artist's sketches and photographs of flowers.

While creative block is the primary motivating scenario, CSTs more broadly support ideation, stylistic exploration and reflective practice [5], and Sketch-It's architecture is applicable across this wider range of creative support contexts.

Related Work and Positioning

Generative Adversarial Networks for Image Synthesis

Generative Adversarial Networks (GANs), introduced by Goodfellow [6] in 2014, have become one of the most influential methods for image generation. By framing generation as a game between a generator and a discriminator, GANs enable the synthesis of highly realistic images from random noise. Deep Convolutional GANs (DCGANs) further improved this paradigm by incorporating convolutional architectures, making it possible to generate photorealistic images with greater spatial coherence and training stability [16]. However, training GANs on small datasets remains challenging, as adversarial learning is often unstable in low-data regimes. Techniques such as spectral normalisation [13], instance noise [11], and label smoothing [19] have therefore been proposed to improve robustness and convergence when data is limited.

Sketch Generation and Colourisation

Automatic colourisation of line drawings has been explored through approaches such as conditional GANs, e.g. pix2pix [7], reference-based methods [24] and user-guided systems that allow partial human control over the output [25]. These methods typically require paired training data, such as a sketch alongside its corresponding coloured version, which is often expensive and labour-intensive to produce. This dependency limits their applicability when only small personal datasets are available. Our approach differs as it operates on unpaired data and uses a perceptual colour critic rather than pixel-level supervision. This makes the method more practical for artist-specific workflows, where the dataset may be small but stylistically consistent.

Protective Measurements Against Artistic Theft

The ethical debate around AI-generated art has been discussed extensively in recent years [4]. Key concerns include attribution, consent and economic impact. Proposals to address this issue include opt-in licensing systems and 'poisoning' tools such as Glaze [21], which adds imperceptible perturbations to artworks to degrade style extraction, and Nightshade [22], which poisons training data. This is effective as a defensive tool against unauthorised use of the images. However, such methods also block the artist from using their own images. By contrast, our work takes a different approach by protecting artists against unauthorised use of their art; it demonstrates that effective generative AI can be built on a fully consented, artist-controlled dataset and later used collaboratively with the artist to produce new art pieces that continue to contribute to the evolution of the artist's cultural background [20].

System Architecture

Sketch-It comprises two GAN models that are trained in parallel and a 'colouriser' that creatively merges the results from both models (see Figure 1. It is presented with a browser-based user interface that makes the full pipeline accessible to non-technical users.

Model 1: Personal Sketch Generator

Model 1 is a Deep Convolutional GAN (DCGAN) intended to be trained exclusively on the artist's own hand-drawn sketches. The generator maps a 128-dimensional vector to a 64×64 greyscale image through four upsampling transposed-convolution blocks. The discriminator uses spectral normalisation to improve training stability. To increase the effective size and diversity of the dataset, we

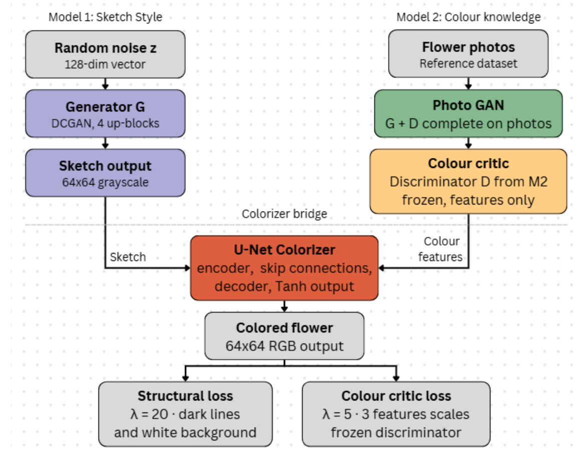

Figure 1: Representation of the system's architecture.

apply random flips, rotations up to $\pm 20°$, scale jitter and brightness variation.

For our flower example, the model was trained for 2000 epochs, see example output in Figure 2a. At convergence, the generator samples novel sketches that share stylistic characteristics with the training artist's work. In our example, this corresponds to petal curvature, line weight, compositional tendencies and the textural quality of pencil on paper. This means that the model is generative rather than retrieval-based; it produces configurations that are not present in the training set.

Model 2: Colour Knowledge from Nature

Model 2 is a separate DCGAN trained on photographs. Its purpose is not to influence the shape or to transform the output into something more photorealistic, but to encode colour knowledge in the discriminator's feature representations. After training, the generator is discarded, and the discriminator is frozen and repurposed as a colour critic. Trained on 2000 epochs in our example, its intermediate feature maps, extracted at multiple spatial scales, encode learned representations of what real flower colours look like at different levels of detail, see example outputs in Figure 2b.

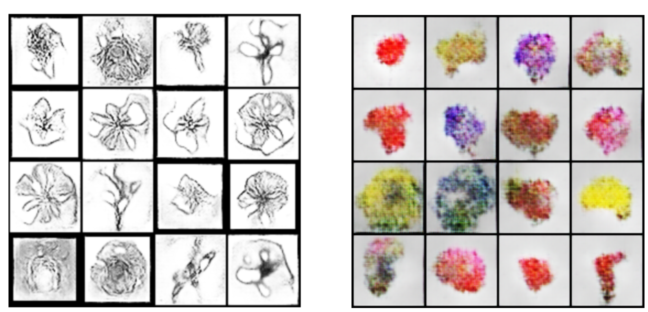

(a) Output from Model 1 after 2000 epochs

(b) Output from Model 2 after 2000 epochs

Figure 2: Example of model outputs.

The Colouriser: Bridging Form and Colour

The colouriser is a U-Net [18] that maps a greyscale sketch $[1 \times 64 \times 64]$ to an RGB coloured image $[3 \times 64 \times 64]$. Training is governed by three losses. The line computes a soft mask over pixels darker than a given threshold and enforces that the greyscale projection of the coloured output matches the original sketch in those regions, ensuring that dark lines remain visible. The background loss computes a soft mask over bright pixels and constrains the output to remain light where the sketch is white. The colour critic loss computes L_1 distance between the colouriser's output features and real photograph features extracted by Model 2's frozen discriminator. Mid-grey petal interior pixels are left unconstrained structurally, allowing the colour critic to fill them freely with naturalistic flower colours. This three-zone approach (lines, interior, background) is the key design decision that enables the system to colour within the sketch rather than replacing it with a photograph.

User Interface

The three models are wrapped in a browser-based local web application built with a Flask back-end (Python) and a single-page front-end developed with HTML and JavaScript[1]. The interface is designed for artists with no prior programming knowledge, presenting the full pipeline as a four-step workflow:

Step 1: First, a legal agreement is presented on screen, in which the user confirms that all uploaded artwork is their own original work and that they hold the intellectual property rights to it. This is a deliberate design choice to ensure data consent and user accountability.

Step 2: Second, the user is presented with a drag-and-drop zone to upload sketched reference images and colour mapping photographs. This is presented together with live file counts and minimum thresholds to prevent underpowered training runs.

Step 3: Third, a "Train Your Model" button is used to trigger all three training stages sequentially shown with a unified progress bar.

Step 4: After training is complete, the system offers the user the ability to use the system to generate new inspirational images.

The application runs entirely locally on the artist's own computer, meaning that no data is shared with a third party. This ensures that the artist retains full ownership and control of their training data, their trained models and all the generated outputs.

Prototype Experiment and Results

The prototype was trained on 104 hand-drawn pencil flower sketches and 34 cropped flower photographs, all produced by a single artist.

Training was performed on a common consumer CPU (13th Gen Intel(R) Core(TM) i5-1335U [1]). Model 1 required approximately 3.5 hours of training (2000 epochs), Model 2 approximately 3-4 hours (2000 epochs), and the colouriser took approximately 1 hour (500 epochs).

[1]Link: `https://github.com/clauudus/SketchIt`

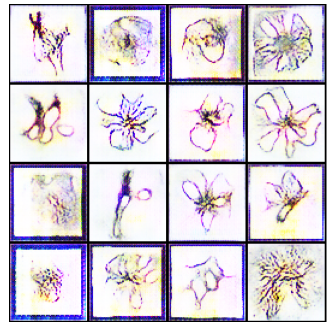

Figure 3: Example of coloured drawing outputs

The generated sketches at epoch 2000 show recognisable flower structures such as daisy-like forms, rose-like clusters or tulip silhouettes, with visible stylistic consistency across seeds, see examples in Figure 2a. Similarly, Model 2 provides flower-like colour patterns from the reference photographs, see Figure 2b. The characteristic line weight and petal curvature of the training artist are preserved across a diverse range of generated forms. The colouriser successfully preserves sketch line work while applying flower-appropriate colour palettes, see Figure 3.

The full generation pipeline (sketch generation followed by colourisation) runs in under one second on CPU, enabling real-time interactive use through the web interface.

Conceptual Evaluation Study

To assess the reception of Sketch-It, we conducted a short anonymous survey with 35 participants. The study assessed the perceived usefulness and its ethical framing from the intended target audience rather than measuring creative outcomes or task performance. Details of the survey and the results can be found in the Appendix.

The results show that the majority of respondents (68.6%) experience creative block *sometimes or more* (mean 2.94/5), supporting the usefulness of CSTs such as Sketch-It. Simultaneously, attitudes toward existing AI tools were predominantly negative or cautious; 60.0% held a *negative or concerned view* with explicit mentions of how AI image generation had already negatively affected their practice or livelihood. Here, creator consent when training AI was rated highly important, receiving a mean of 4.11/5, with 74.3% scoring 4 or 5. For tools, like Sketch-It, which train exclusively on the user's own artwork, 62.9% of respondents reported a positive perspective, see Figure 11 (in the Appendix). The likelihood that the participants would try the tool had a mean of 3.03/5, with 71.4% scoring 3 or above. The most frequently recurring concern was over-reliance of the tools. Respondents worried that using an AI tool during a creative block might inhibit the development of internal creative resources. A second recurring concern was residual security: even with local storage, one respondent noted their images could still be compromised. Both concerns directly inform the future work directions described below.

Discussion

The Ethical and Environmental Contribution

One of the most important contributions of this work is the demonstration that consent-first generative AI can be practically realised. The system produces useful creative outputs (ie., novel flower sketches in a personal style, plausibly coloured) from a dataset of around 100 images created by and belonging to a single individual.

As a consequence of the small dataset, our approach also represents a "green, cost-efficient generative AI system. Large foundation models for image generation require training runs of thousands of GPU-hours on datasets of billions of images [17]. Our solution's total training energy is several orders of magnitude smaller, achievable on a consumer laptop overnight. For a field increasingly scrutinised for its environmental footprint, small-dataset, consent-first systems represent a meaningful alternative paradigm.

Existing responses to AI ethics problem in art take a predominantly defensive stance, tools such as Glaze [21] and Nightshade [22] protect artists by making their work resistant to style extraction, while opt-in licensing frameworks seek to regulate access[9]. Sketch-It takes a complementary but distinct affirmative approach; rather than protecting against misuse of existing systems, it demonstrates that genAI systems can be built from the ground up on fully consented data.

Creativity Support, Not Creative Replacement

Designed as a creativity support tool, our system is not a replacement for artistic production. In comparison to more commercial AI image generators, the generated output does not represent complete or polished images. Instead, the generated sketches are intended as creative starting points that are to be interpreted as inspirational visual prompts that reflect the artist's own vocabulary in new configurations. The artist does not press a button and receive finished artwork; rather, they receive a suggestion from which they can diverge, react or build. This positions the system within what Kantosalo [8] calls co-creative AI: a collaborative partner that augments rather than supplants human creative agency.

The cultural implications are significant. Trained on the artist's own work, such tools do not homogenise culture towards the statistical average of a billion training images; instead, it amplifies the individual artist's particular perspective. In this sense, our solution is not only a tool for overcoming creative block, but also an argument about what AI-assisted culture can look like when the human at the centre of the system retains authorship of its foundations.

Generalisability: Beyond Flowers, Beyond Sketches

Flowers were chosen as the demonstration domain because they offer sufficient visual variety to test the system while remaining tractable for small-dataset training. The architecture is designed to be domain-agnostic – an architect could train Model 1 on their sketched floor plans and Model 2 on photographs of built spaces, or a character designer could train on their character sheets with colour reference from comics or frames from animation. However, these remain theoretical extensions. A full multi-domain evaluation is identified as an important direction for future work.

Our experiments suggest a minimum of 50 images for the drawings and 20 photographs for the colourisation. This threshold depends largely on the artist's needs and the level of abstraction they expect; the smaller the dataset, the more abstract the generated output is likely to be.

Limitations

The current prototype was trained on a single artist's flower sketches and photographs. Claims about generalizability to other artistic domains, styles, or levels of stylistic consistency remain theoretical and untested. Additionally, the evaluation measured attitudes and perceived reception rather than creative effectiveness, meaning that it cannot be concluded that Sketch-It demonstrably reduces creative block or improves creative output. These limitations should be addressed in future work on multi-artist, multi-domain experiments and task-based evaluations.

Conclusion

We have presented a creativity support system that generates novel artworks in an artist's personal style using only their own consented drawings. By separating the learning of artistic form (a personal sketch GAN) from the learning of colour (a nature-photograph GAN), and connecting them through a learned colouriser guided by a three-zone structural loss, we demonstrate that ethically grounded generative AI can produce meaningful creative output from minimal data. A preliminary survey of 35 participants confirms a strong interest in consent-first approaches: 74.3% rated dataset consent as highly important, and 62.9% responded positively to Sketch-It's framing relative to existing tools.

We hope that this solution contributes to a shift in how the field thinks about AI-assisted creativity, away from systems that extract value from artists without consent, and towards systems that amplify individual creative voices while keeping the human firmly in the centre. The best creative AI, we argue, is not one that replaces artistic thinking; it is one that says: *"Here is something new, in your own voice. What do you want to do with it?"*

Future Work

The legal consent model, while carefully designed, is self-reported. Future work could explore integration with digital provenance tools or blockchain-based artwork certification. The interface could also be extended with real-time training visualisation, showing the model improving epoch by epoch. Regarding security beyond consent, while the local architecture ensures training data never leaves the artist's machine under normal use, future versions could integrate adversarial perturbation techniques [21] [22] applied automatically before local storage, offering a dual guarantee; the data is not shared by design, and even if compromised, would be resistant to exploitation.

Authors contribution

As lead author, Pàmies designed the system and performed the survey. Schaap and Hedblom acted as supervisors who aided in project guidance and co-authorship.

References

[1] Intel® Core™ i5-1335U Processor (12M Cache, up to 4.60 GHz) - Product Specifications — Intel — intel.com. [Accessed 19-04-2026].

[2] Ardalan Arabzadeh, Tobias Vente, and Joeran Beel. Green recommender systems: Optimizing dataset size for energy-efficient algorithm performance. October 2024.

[3] Mihaly Csikszentmihalyi. *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins, New York, 1996.

[4] Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. CAN: Creative adversarial networks, generating "art" by learning about styles and deviating from style norms. June 2017.

[5] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. Mapping the landscape of creativity support tools in HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, May 2019. ACM.

[6] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. 2014.

[7] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. November 2016.

[8] Anna Kantosalo and Tapio Takala. Five c's for human–computer co-creativity: An update on classical creativity perspectives. In Amílcar Cardoso, Penousal Machado, Tony Veale, and {João Miguel} Cunha, editors, *Proceedings of the Eleventh International Conference on Computational Creativity*, pages 17–24, Portugal, September 2020. Association for Computational Creativity. International Conference on Computational Creativity, ICCC ; Conference date: 07-09-2020 Through 11-09-2020.

[9] Yasuto Komada. New law responding to imitations of designs that intersect real and virtual space: Amendment of the unfair competition prevention act in japan. *GRUR International*, 74(2):153–159, January 2025.

[10] Lin Kyi, Amruta Mahuli, M Six Silberman, Reuben Binns, Jun Zhao, and Asia J Biega. Governance of generative AI in creative work: Consent, credit, compensation, and beyond. January 2025.

[11] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? January 2018.

[12] Midjourney, Inc. Midjourney, 2024. Version 6.0.

[13] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. February 2018.

[14] OpenAI. DALL-E 3, 2023. Image generated with prompt: "prompt".

[15] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. July 2023.

[16] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. November 2015.

[17] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. December 2021.

[18] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. May 2015.

[19] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. June 2016.

[20] Kelsey Schoenberg. How Generative AI Endangers Cultural Narratives - issues.org. https://issues.org/generative-ai-cultural-narratives-rettberg/. [Accessed 19-04-2026].

[21] Shawn Shan, Jenna Cryan, Emily Wenger, Haitao Zheng, Rana Hanocka, and Ben Y. Zhao. Glaze: Protecting artists from style mimicry by Text-to-Image models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 2187–2204, Anaheim, CA, August 2023. USENIX Association.

[22] Shawn Shan, Wenxin Ding, Josephine Passananti, Stanley Wu, Haitao Zheng, and Ben Y Zhao. Nightshade: Prompt-specific poisoning attacks on text-to-image generative models. October 2023.

[23] Ben Shneiderman. Creativity support tools: accelerating discovery and innovation. *Commun. ACM*, 50(12):20–32, December 2007.

[24] Lvmin Zhang, Yi Ji, and Xin Lin. Style transfer for anime sketches with enhanced residual u-net and auxiliary classifier GAN. June 2017.

[25] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. May 2017.

Appendix

Additional Data from the Study

Experimental Setup and Participants

The study was conducted as an anonymous online questionnaire. Participants were invited through convenience sampling through social network distribution.

The majority of respondents identified as hobbyist creators (60.0%), with smaller numbers describing themselves as non-practitioners interested in art (20.0%), semi-professional artists with occasional commissions (17.1%), students in artistic training (5.7%), and two professional artists whose art constitutes their primary income.

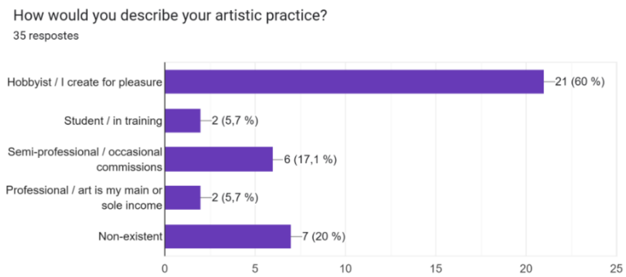

Figure 4: How did participants identify in an artistic spectrum

Respondents worked mostly with pencil and ink drawing, digital illustration, painting, and music. Two respondents identified dance and writing as their primary practice, reflecting a broader creative community beyond visual art.

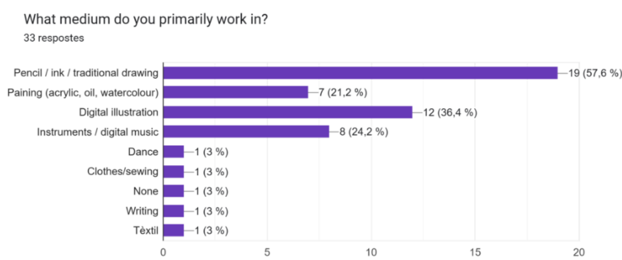

Figure 5: What kind of materials do artists tend to use?

Questionnaire Questions

The study asked 10 questions related to the use of creativity support tools like Sketch-It.

Likelihood of Creative Block and Solutions The participants rated the frequency of a creative block on a five-point scale from 1 (never) to 5 (very often). The mean response was 2.94, and 24 of 35 respondents (68.6%) scored 3 or higher, confirming that creative block is at least a sometimes occurrence for most of the sample. Only three respondents (8.6%) reported never experiencing it.

The most common coping strategies were looking at others' work for inspiration and searching for reference images online. Only one respondent reported using an AI image tool

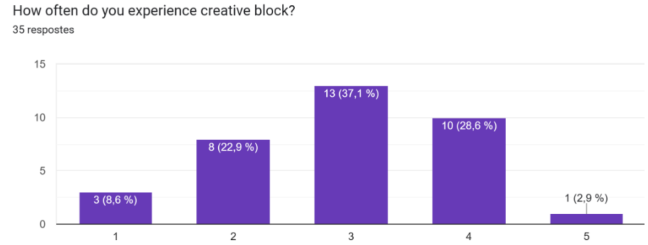

Figure 6: Participants and creative block

as a strategy during creative block, suggesting that the space for ethical AI creativity support remains largely unoccupied.

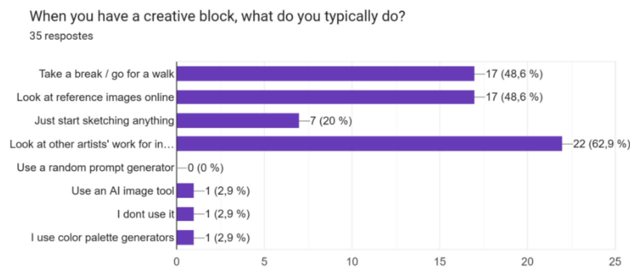

Figure 7: How artists normally overcome an Art Block

Attitudes Toward AI Image Generation Attitudes towards existing AI tools were predominantly negative or cautious: 37.1% described themselves as opposed to AI-generated imagery on principle, 22.9% expressed concern about exploitation of artists' work without consent, 34.3% were neutral, and only 5.7% responded positively. In total, 60.0% had a negative or concerned view.

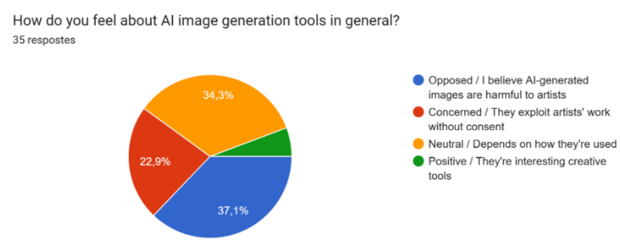

Figure 8: How do Participants feel about AI?

Eleven respondents (31.4%) reported that AI generation had already negatively affected their creative practice, compared to two participants who reported a positive effect.

The Role of Consent When asked to rate how important creator consent is for the use in the training dataset, on a five-point scale, the mean response was 4.11 out of 5. Twenty-one respondents (60.0%) assigned a maximum score of 5, and 26 of 35 (74.3%) rated the importance of consent at 4 or greater.

When told that Sketch-It trains exclusively on the user's own artwork and asked whether this changes their feelings

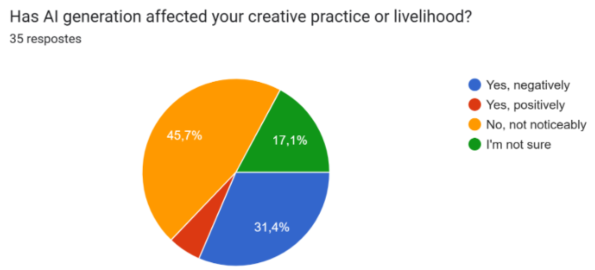

Figure 9: Perceived effect of AI Generation tools on the participants' lives

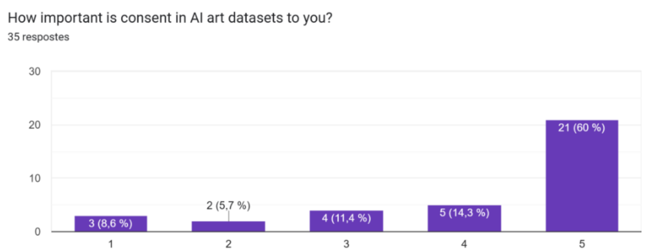

Figure 10: Graph indicating the concern towards consent in AI training datasets

compared to other tools, 62.9% responded positively, as seen in Figure 11 (in the Appendix). 14.3% stated that consent makes a significant difference, and 48.6% said that it makes some difference, although they retain reservations. Nine respondents (25.7%) said that the consent-first architecture does not change their concerns, and four (11.4%) reported having no prior concerns.

Perspective on fully consented dataset To determine if Sketch-It could be a well received tool among artists, we presented an explanation on how it would work to then know if the perspective on generative AI would change. 25.7% stated that they still had the same concerns, while 62.9% had a positive shift on the matter.

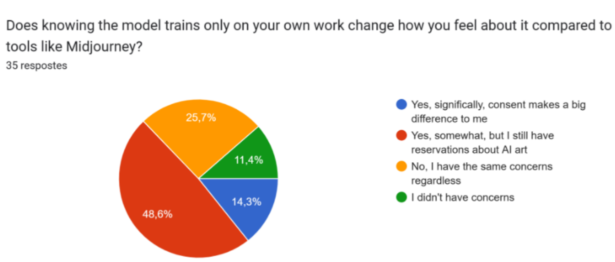

Figure 11: Results of perspective on user-art only trained AI systems.

Receptiveness to Sketch-It Participants rated their likelihood of trying Sketch-It on a five-point scale (mean 3.03). Eleven respondents (31.4%) gave a score of 4 or 5, while then (28.6%) gave 1 or 2, indicating reluctance. With 71.4% scoring 3 or above, the majority of the sample is at minimum open to trying the tool; a positive baseline given that 60.0% hold negative or concerned views on AI image generation in general.

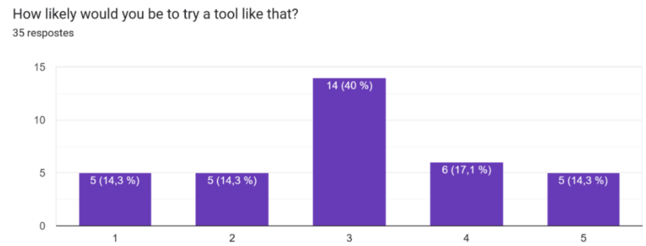

Figure 12: Participants trying Sketch-It

The two most frequently selected appealing attributes were the tool's small environmental footprint (18 respondents, 51.4%) and its consent-first training using only the artist's own drawings (17, 48.6%). Local data storage was selected by 14 respondents (40.0%), personal style generation and creative block assistance by eight each (22.9%). The prominence of ethical and environmental attributes over functional ones suggests that for this audience, the conditions of trust matter as much as utility.

Concerns and Open Responses The most frequently recurring concern was over-reliance. Several respondents worried that using an AI tool to escape creative block might inhibit the development of internal creative resources. One participant wrote that a great way to overcome art block is by purely drawing, while another participant noted that "it can be a thin line between using it as a tool or becoming dependent on it". A second recurring concern was residual data security. One respondent noted that their images, stored on their own computer, "could still be compromised". This directly motivated the adversarial perturbation direction described in Future Works. Several open responses pointed toward productive extensions, where two of the participants explained that having such a tool to help artists overcome art block in a very interactive sense would motivate them to try it out.

Summary

The study confirms three of Sketch-It's core: creative block is real and common (68.6% experience it sometimes or more), consent is a central concern in attitudes towards AI art tools (mean 4.11/5, 74.3% rating it 4 or above), and the consent-first approach meaningfully improves reception among artists otherwise sceptical of AI (62.9% positive

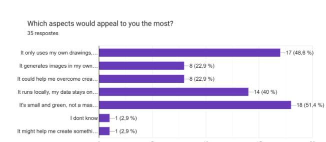

Figure 13: Picked attributes among the participants

shift). Primary concerns, including over-reliance, residual data security, and questions of artistic authenticity, are genuine and inform the future work directions described in the main paper. We note that the sample is small (n=35) and that a longitudinal study with a larger, more diverse participant group would be needed to draw stronger conclusions.