"There is Beauty in the Small": A Study of Visual Artists' Impressions of Small Data and Model Crafting Approaches to Creative AI Work

Ahmed M. Abuzuraiq
School of Interactive Arts and Technology
Simon Fraser University
Surrey, British Columbia, Canada
aabuzura@sfu.ca

Arshia Sobhan
School of Interactive Arts and Technology
Simon Fraser University
Surrey, British Columbia, Canada
asobhans@sfu.ca

Alexandra Kitson
Department of Computer Science
University of Victoria
Victoria, British Columbia, Canada
akitson@uvic.ca

Philippe Pasquier
School of Interactive Arts and Technology
Simon Fraser University
Surrey, British Columbia, Canada
pasquier@sfu.ca

Abstract

Large-scale prompt-based AI generators have broadened access to media creation but often introduce friction in sustained creative practice. Visual artists must navigate limited personalization, narrow cultural coverage, homogenized aesthetics, and broader concerns around sustainability, bias, and data-scraping ethics. Small data and model crafting offer alternatives that center artists' agency over data and models. We examine this approach through Autolume, a no-code, locally run tool for crafting and navigating generative models in real time. Across five focus groups with 25 visual artists in three two-day workshops, Autolume was used as a research probe to explore small-data practices. Findings show that participants viewed small data as a personal and responsible alternative to large-scale AI, but identified model training as a technical barrier, underscoring the need for improved explainability and support. We highlight design opportunities for AI art tools, including low-level controls, mapping UIs to ML processes, ecosystem integration, and small data as a pedagogical tool.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**; • **Applied computing → Media arts**.

Keywords

Generative AI, AI Art, Visual Art, Small Data, Qualitative Research

ACM Reference Format:
Ahmed M. Abuzuraiq, Arshia Sobhan, Alexandra Kitson, and Philippe Pasquier. 2026. "There is Beauty in the Small": A Study of Visual Artists' Impressions of Small Data and Model Crafting Approaches to Creative AI Work. In *Designing Interactive Systems Conference (DIS '26), June 13–17, 2026, Singapore, Singapore.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3800645.3812958

ACM ISBN 979-8-4007-2563-0/26/06
https://doi.org/10.1145/3800645.3812958

1 Introduction

Artistic creation relies on expressive and controllable mediums—such as a painter who masters the brush and oils to realize an aesthetic vision. This tradition extends to artists who use algorithms and data as their medium, a practice variously known as algorithmic, generative, or AI art. Exploring the creative potential of generative AI models was pioneered by artists who combined artistic and technical expertise, adapting and training generative systems on personal or carefully curated datasets.

The proliferation of large-scale text-to-image generative models (LTGMs) such as Midjourney, Flux, or Nano Banana marks a significant shift in this lineage. These systems have broadened access to image generation through simple text prompts, but critical engagement by artists and researchers has surfaced substantial concerns for creative practice: privileging outputs over process [54], complicating or contesting artists' authorship [67], and biasing results toward dominant aesthetics [29]. In response, small-data and model-crafting approaches continue earlier traditions of AI art while offering alternatives that may mitigate some of the issues associated with LTGMs. A small-data approach focuses on training (or fine-tuning) models on personally curated or created small datasets to achieve specific intents as opposed to creating general, large-scale models [84], while model crafting refers to creating and manipulating generative models for creative intents [1].

However, there remains a gap in understanding how visual artists with varying technical expertise view alternatives to LTGMs, and how those views could inform the design of AI tools for art making. Prior studies of AI art practices either predate the rise of LTGMs [19, 65], focus on publicly available materials from prominent artists [71], or focus on AI-based music making [16]. Building on this work, we focus on visual artists and ground our inquiry in Autolume (Section 3.1) – a concrete no-code small-data and model-crafting tool. We use Autolume [58, 75] as a probe to elicit artists' impressions of small-data workflows and surface design opportunities through hands-on engagement. In particular, we ask:

RQ1 What are visual artists' impressions of the concepts of small data and model crafting?
RQ2 What considerations are relevant for artists when working with small-data AI tools for art creation (e.g., Autolume), particularly regarding aesthetics, authorship, creative control, and integration into their practice?
RQ3 What design opportunities can be derived from artists' views on small data and model crafting (RQ1) and their considerations when using small-data AI tools (RQ2)?

To address the above questions, we conducted 5 focus groups with 25 visual artists over three 2-day workshops. The workshops trained participants in using Autolume, introduced them to key technical concepts, such as AI training and latent spaces, to operate the tool, and invited them to train AI models on their own datasets. Following that, we used Autolume as a probe to solicit artists' impressions of the tool and the concepts underlying its design through a focus group discussion. Autolume is a no-code, open-source, and locally-run tool that enables artists to train and adapt their own generative models and interactively explore them in real time.

This paper's primary contribution is a set of qualitative insights derived from a thematic analysis of the focus group discussions. Artists' feedback reveals their impressions about Autolume and small-data approaches, and it surfaces potential benefits and challenges associated with small-data approaches to AI creative work. These insights also foreground design-relevant concerns around explainability, authorship, creative control, and aesthetics. We distill these findings into actionable design opportunities for AI art tools, contributing to a deeper understanding of AI practices in the arts and, more broadly, to HCI discussions on human-centred AI in the arts [15].

2 Background

This background provides an overview of the technical, empirical, and design concerns surrounding the interdisciplinary challenge of designing artist-centred interfaces for AI art. To this end, we review the history of the field, empirical studies of creative practices, and the current technological landscape. We also highlight artistic critiques of –and initiatives responding to– Large-scale Text-to-image Generative Models (LTGMs), while evaluating relevant work in tool design to situate our findings.

2.1 AI Art: Brief History and Approaches

Art made with Artificial Intelligence, though recently popularized, extends a longer history of generative, algorithmic and computer art, marked by technical and conceptual shifts that reshaped artists' relations with computation. Early examples of algorithmic arts include AARON, a rule-based system created in 1973 by Harold Cohen [20]. A major shift came with machine learning, particularly Generative Adversarial Networks (GANs) [33], moving the focus from coding rules to curating training data [20]. This "GAN era" fostered practices of dataset building and navigating latent spaces[1] to discover compelling outputs [19, 65, 71]. One example is Mario Klingemann's *Memories of Passersby I* (2018), an installation generating endless portraits using GANs [2]. More recently, large-scale text-to-image diffusion models (LTGMs) trained on massive internet datasets –such as Stable Diffusion, Midjourney, and DALL-E– have enabled the creation of increasingly high-fidelity visuals [26]. These technical advances have sparked a surge of interest in AI art[3], reflected in the rapidly growing online communities devoted to exploring and sharing generative imagery [4]. To better understand the current context of the study, we briefly describe next some of the critiques and responses to LTGMs.

2.2 Critiques from the Art World to LTGMs

The rise of LTGMs has prompted significant critique from artistic and research communities, sparking counter-initiatives aimed at fostering more personal, culturally representative, and creatively engaging AI-art practices. Reported concerns include a perceived lack of authenticity [57], limited creative satisfaction [67], a diminished sense of surprise [28], and reservations about authorship due to a perceived lack of effort or labor [67]. Additionally, practitioners have highlighted the friction of forced verbal articulation of tacit artistic intent through prompting [28, 47, 54]. This challenge is underscored by empirical findings demonstrating that artists must master a complex, style-specific and technical vocabulary to effectively align AI models with their creative intents [52, 61].

A central aesthetic critique is that LTGMs tend to produce generic outputs lacking a distinctive artistic voice [47] while reinforcing dominant visual styles [29]. LTGMs also frequently misrepresent or fail to generate culturally specific concepts—such as Arabic script [76] and other non-Western art forms [66]—which have prompted calls for community-based efforts to build more equitable, culturally specific datasets and models (e.g., [16]). Moreover, LTGMs' push for photorealism can stifle abstraction; they have been observed to exhibit an unintended bias toward figuration even when prompted with abstract inputs [88].

The accumulation of the above issues creates a sense of disillusionment among AI artists, who grapple with issues beyond their control; whether limited by the core design of large-scale models [9] or compelled to develop personal pipelines to circumvent them [9, 19]. Finally, the critiques listed above focus on artists' experience with creative work with AI or the frictions that arise from integrating AI into artists' creative processes. Those experiential and process-oriented critiques exist alongside broader concerns around sustainability, algorithmic bias, and data-scraping ethics [29]. These broader concerns continue to deter many artists from engaging with AI, as well as tainting how the work of AI artists is received and evaluated by others.

[1]Latent spaces are low-dimensional vector spaces a generator learns, where nearby points yield semantically similar images and certain directions align with interpretable attributes (e.g., pose, style, colour).
[2]More examples of GAN art can be found on https://aiartists.org/
[3]Examples of recent *exhibited* AI art can be found at https://thecvf-art.com/
[4]Community-shared prompts and models can be accessed at https://civitai.green/, https://tensor.art/, or https://promptbase.com/shop

2.3 Initiatives from the Art World to LTGMs

In response to these challenges, a number of artist- and community-led initiatives have emerged, focused on reclaiming creative agency and fostering a more critical and reflective engagement with AI. These initiatives include small-data approaches, the reinvigoration of AI craft, AI explainability through art practice [17], and research-creation as a means of exploring new or critical ways to engage with AI [69]. We focus below on the first two.

2.3.1 Small-Data Approaches. Studies with AI artists describe their practice as a *craft* over data and models [19], where generative systems are treated as artistic material [12]. A small-data mindset [16, 84] focuses on reusing, curating, or creating small datasets [65] to inject personal style, voice, or cultural elements into a generative model. In this context, the traditional ML pitfalls of overfitting and bias –typically avoided in models designed for generalization– are reframed as desirable traits. By leaning into these 'issues', artists can generate more recognizable outputs [84] and critically engage with how specific dataset features fundamentally shape large-scale AI systems [5, 8]. Finally, a small-data approach to creative AI work is not limited to training models from scratch, as it also encompasses personalizing large-scale models through fine-tuning techniques (e.g., Low-Rank Adaptation or LoRA [41]) and conditioning workflows such as image-based prompting (e.g., IP-Adapter [87]) or spatial controls (e.g., ControlNet [89]). However, personalizing large-scale models does not eliminate their underlying and emergent biases, nor reduce their energy consumption.

2.3.2 Model Crafting. Model crafting involves creating and manipulating generative models for creative intents [1]. This includes designing their architecture, training them, down to intervening in their normal functioning [12], and manipulating or mixing their parts to push them beyond the boundaries of their training [13]. Until recently, model crafting has been exclusively the purview of artists with coding expertise, but recent tools have made certain aspects of it accessible through low-code interfaces (cf. Section 2.5).

2.4 Studies on AI Art Processes

To situate our study, we turn now to studies that explored AI art practices before and after LTGMs. Early research into artistic engagement with machine learning (ML) focused on pioneering artists who treated AI as a complex, "stubborn medium" requiring craft and expertise [65]. These studies describe a practice centred on iterative exploration, where artists develop an intimate understanding of algorithms and datasets as materials for creation by tweaking them and embracing their unpredictable "semantic glitches" as sources of novelty [19]. AI can be understood as a responsive material in this process, where the artist and AI discover each other's affordances to achieve an "attentive unity" between human intent and machine capability [71]. The goal for these expert practitioners was not automation but augmentation, i.e., balancing control and surprise to push their work in new directions [65, 71]. This often involved the significant labour of curating or creating bespoke datasets, and modifying algorithms [19, 65], as well as the curation of outputs from the model's "latent space".

The studies mentioned above were predominantly conducted with AI artists before the rise of LTGMs [19, 65] or analyzed publicly available material of prominent artists [71]. Together, they point to a sense of *craft* that permeates how these artists view their creative practices with AI.

The rise of LTGMs shifted this landscape, introducing new challenges for a broader range of artists, which we covered earlier (Section 2.2). In response to large-scale AI, many artists are developing tactics of resistance by exploring "small-data" approaches (or "Little AI" [65]) to maintain ethical and creative control [83]. Working with small, personal datasets is proposed as a way to ensure artistic authenticity, reduce bias to dominant genres, retain attribution, and address privacy concerns [16, 32]. While small-data methods predate large-scale AI (e.g., [30]), their value has been renewed in the current context [65, 84].

Our study is situated in this recent turn toward small data. While prior work has examined the technical craft of pre-LTGM experts [19, 65, 71] and the tensions of creating with LTGMs [9, 47, 66, 67], we investigate how visual artists with varying AI expertise perceive small data and model crafting in a post-LTGM context. Building on Bryan-Kinns et al. [16], which found that small datasets and low-resource models can increase agency and transparency while reducing bias and energy use in AI-assisted music making, we extend this lens to visual artists and ground discussions in a specific small-data and model-crafting tool (Autolume). Prior studies have also explored small dataset fine-tuning for culturally specific needs [66] and how personalized LTGMs might fit into existing practices [67]; in contrast, we focus on training low-resource models from scratch and probe artists' impressions of a broad range of topics, including authorship, aesthetics, and explainability.

2.5 AI Art Ecosystem

Having outlined the social and conceptual landscape surrounding creative work with AI, we now turn to the technical landscape. Generative models, by definition, are models that learn the distribution of the training dataset in order to generate new, similar data. Two families of image generation models are most commonly used by AI artists; these include diffusion models [70] and Generative Adversarial Networks (GANs) [33]. In what follows, we describe some of the tools and interactions related to each.

2.5.1 (Text-to-Image) Diffusion Generation. Diffusion models generate images by learning to reverse the noise that is added to the training set. Latent diffusion models work similarly, but perform noising and denoising in a compressed latent representation rather than pixel space, which makes training and inference more efficient. Text-to-image systems combine diffusion with conditioning models, typically a text encoder (e.g., CLIP or T5) that maps text into embeddings that are used to guide generation. Diffusion models are currently the dominant approach for image generation. Some are proprietary and accessed through APIs or hosted products (Midjourney V6, Google Nano Banana), while others are publicly shared (Stable Diffusion [79] and Flux [10]).

Tooling for diffusion models generally falls into four groups: (1) local open-source UIs (Automatic1111 [7], ComfyUI [25]); (2) creative suites (Midjourney [59], Runway [72]); (3) multimodal chat (ChatGPT [60], Gemini [34]); and (4) integrated authoring tools (Adobe Firefly [4]). Across these environments, users typically specify prompts, adjust generation parameters, and iterate on outputs.

Interfaces may take the form of prompt-driven GUIs, node-based workflows, or canvas-based systems. Many tools also offer editing modes geared toward intentional authoring, including generation guided by style, structure, or composition references, inpainting and outpainting, object selection and replacement, other forms of localized editing, and adjustments to lighting and camera angle.

Because training diffusion models from scratch is resource-intensive, artists typically rely on pre-trained base models, community fine-tunes (hosted on HuggingFace [43] or CivitAI [24]), or personalization techniques like Low-rank Adaptation (LoRAs) [41]. While standard models rarely run in real-time, distilled versions (e.g., SDXL Turbo [80]) and specialized pipelines (e.g., StreamDiffusion [48]) can trade fidelity for speed.

2.5.2 GAN-based Generation. Generative Adversarial Networks (GANs) use a "generator" and a "discriminator" model that work against each other to produce images that closely resemble the training set. Compared to diffusion, GAN tooling offers fewer no/low-code paths. Options include: (1) artist-facing interfaces such as Playform [64], Autolume [58] and the (now-discontinued) GAN module on RunwayML [56]; (2) low-code Colab notebooks[5]; and (3) configurable research frameworks/SDKs [44, 62, 73].

While diffusion tools prioritize high-fidelity *intentional authoring*, GANs excel at interactive, real-time exploration of latent space. Consequently, GANs are well-suited for generating ambient content –textures and patterns that are *difficult to describe* concretely. In this study, we use GANs to examine small data and model crafting. GANs are faster and cheaper to train from scratch than diffusion models, making them a practical fit for small-data visual generation. Training from scratch also lets artists directly observe how a specific dataset shapes outputs, and allows us to study small-data practices in a context less entangled with broader concerns around large-scale AI systems (cf. Section 2.2). GANs have also been central to AI art practice and discourse (Section 2.1), making them especially relevant here. While GANs can be trained conditionally, including with text prompts [45], we focus on unconditional GANs because they foreground latent-space exploration and can invite surprise in both training and generation, a quality that may be de-emphasized in systems optimized for prompt adherence [28].

We chose Autolume [58] as our technology probe because it is the only tool that is free, open-source, feature-rich, runs locally, and trains GANs. Built on StyleGAN2 (a GAN variant), Autolume supports training on datasets as small as 100 images. It produces small models (often < 400MB) and generates images in milliseconds on consumer GPUs, making it suitable for interactive art purposes [78]. Autolume also enables real-time latent space navigation and model crafting features. See Section 3.1 for more details.

2.6 Design of AI Art Interfaces

The design of effective and enjoyable interfaces for AI art tools presents a unique set of challenges and opportunities for the HCI community. First, the expanding capabilities of generative AI have raised foundational questions about how artists relate to these systems, for instance, whether they should be framed as creativity-support tools [74] or as co-creative systems [68], and that framing may imply different design interventions.

Secondly, machine learning is a notoriously difficult material to design around [27]. To address this, analytical frameworks and guidelines tailored specifically to generative systems [18, 35, 86] have become essential resources for tool designers. However, the design challenge is significantly amplified when the material – including both data and models– is intended to be crafted, shaped and manipulated by users, as seen in small-data and model-crafting artistic practices. This approach is frequently overlooked in HCI literature, which historically privileges user interaction with pre-trained (and increasingly large-scale) generative AI models over the act of crafting them [42]. Indeed, Xi et al. found that ~7% of the 189 surveyed systems supported such model crafting interactions.

Proper understanding of AI, tacitly and declaratively, is vital for effectively crafting it, which reinforces the importance of explainability for AI art practices [17], especially for a sustained artistic practice with AI [81]. In this context, analogies and metaphors play an important role in how users understand generative AI systems [37], and in framing artists' activities surrounding them (e.g., "tricking, taming, and breaking" [66]). Finally, when designing tools for artists, it is imperative to clearly define the qualities being prioritized —such as quality, efficiency, control, enjoyment, intuitiveness, and the potential for mastery— so that these tools meet artists within their natural workflows [53].

3 Methods

We conducted three workshops, each spanning 2 days, with a total of 25 participants joining 5 focus groups. Each workshop combined an educational and an optional research component. The educational part trained participants in using Autolume and introduced key technical terms (e.g., AI training, latent space). The research component consisted of a focus group discussion at the end of the workshop. The following sections describe the Autolume interface and educational component, which contextualize the research activities.

3.1 Autolume

Autolume is a software system, developed by the Metacreation Lab for Creative AI at Simon Fraser University, that enables artists to explore generative AI with small datasets [58, 75]. It combines offline processing modules and a real-time module to support end-to-end workflows, from dataset preparation and model training to live visual exploration and interactive performance. For more details, including a video of the system and sample artworks created with Autolume, we refer the reader to https://www.metacreation.net/autolume.

3.1.1 Offline Processing Modules. Autolume provides five offline processing modules for model creation, preparation, and post processing:

(1) Model Training. Users can train StyleGAN2 [46] models from scratch or fine-tune pre-trained ones using image datasets from a collection of stills or extracted video frames. The module includes pre-processing and small-data augmentation pipelines. With datasets as small as hundreds of images,

[5]A curated list of GAN-based Colab notebooks is available at https://pharmapsychotic.com/tools.html

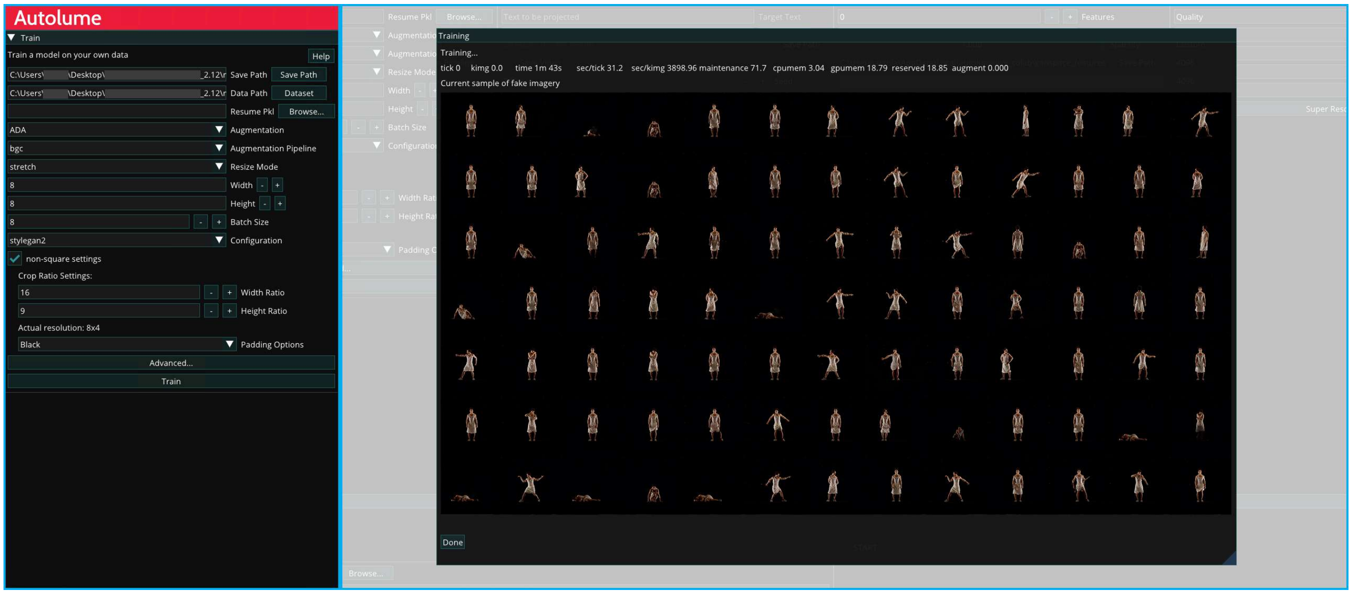

Figure 1: Left: Autolume's model training module. Right: Autolume's model training progress preview window. The dataset shown comes from *Longing + Forgetting*, an artwork by Philippe Pasquier, Thecla Schiphorst, and Matt Gingold. More at: https://www.metacreation.net/projects/longing-forgetting.

users can iterate toward desired results by adjusting hyperparameters (e.g., batch size, learning rate, gamma). Training runs on local GPUs, with meaningful outcomes in 12 hours on an RTX 3090, with progress monitored through real-time image grids[6].

(2) Latent Space Projection. Users can input an image, text prompt, or both to locate the nearest latent vector in a trained model. Output vectors can be saved for real-time use.

(3) Feature Extraction. Using GANSpace [38], interpretable semantic latent directions can be computed. These can be saved and explored in the real-time module.

(4) Super Resolution. Autolume uses Real-ESRGAN [85] to upscale outputs, with different modes for quality and speed.

3.1.2 Real-time Module (Renderer). The real-time module provides functionalities that allow artists to interact with trained models in real time, offering direct manipulation and fine-grained control of latent space, model bending and blending, and interactive features (Figure 2).

(1) Latent Space Navigation Features. Latent space can be explored through options such as fixed seed initialization, random walks with adjustable speed, and multiple interpolation techniques (Figure 2-a). Artists can control the visual features by adjusting parameters such as global noise and truncation (Figure 2-b). Users can mark points in the latent space and loop between them as keyframes, enabling curated explorations (Figure 2-c). Pre-calculated directions from the offline Feature Extraction can also be loaded and explored in real-time (Figure 2-d).

(2) Model crafting features [1]. Trained models can be manipulated by applying transformations to internal model layers (model bending [14], Figure 2-e) or mixing layers from multiple models (model blending or mixing [63], Figure 2-f), both of which help push outputs toward experimental or abstract aesthetics and beyond the original training distribution [13].

(3) Interactive Features. The real-time module supports Open Sound Control (OSC), a commonly used protocol that enables real-time data exchange and control. For example, this enables audio-reactive installations where interface parameters, like interpolation, are controlled through audio signals in real-time. The generated visuals can be sent through the Network Device Interface (NDI) protocol to other software for additional post-processing (Figure 2-g).

(4) History Keeping. Users can save presets that capture the full state of a session for reuse in future sessions (Figure 2-h).

3.1.3 User Workflow. A typical workflow in Autolume reflects this integration of offline and real-time processes. An artist may begin by using the Model Training module to train a model from scratch on a small dataset, or to fine-tune an existing model using the same interface. Once training is complete, the resulting model can be loaded into the Real-time Module (Renderer), where the artist explores and navigates its latent space through continuous interaction and direct manipulation. This phase also supports model bending/blending for additional visual exploration, and interactive features for live or installation contexts. In addition, tools such as Latent Space Projection, Feature Extraction, and Super Resolution can be used to extend this workflow for more advanced control, refinement, and preparation of outputs for final production.

[6]Training quality depends on the dataset's size and coherence.

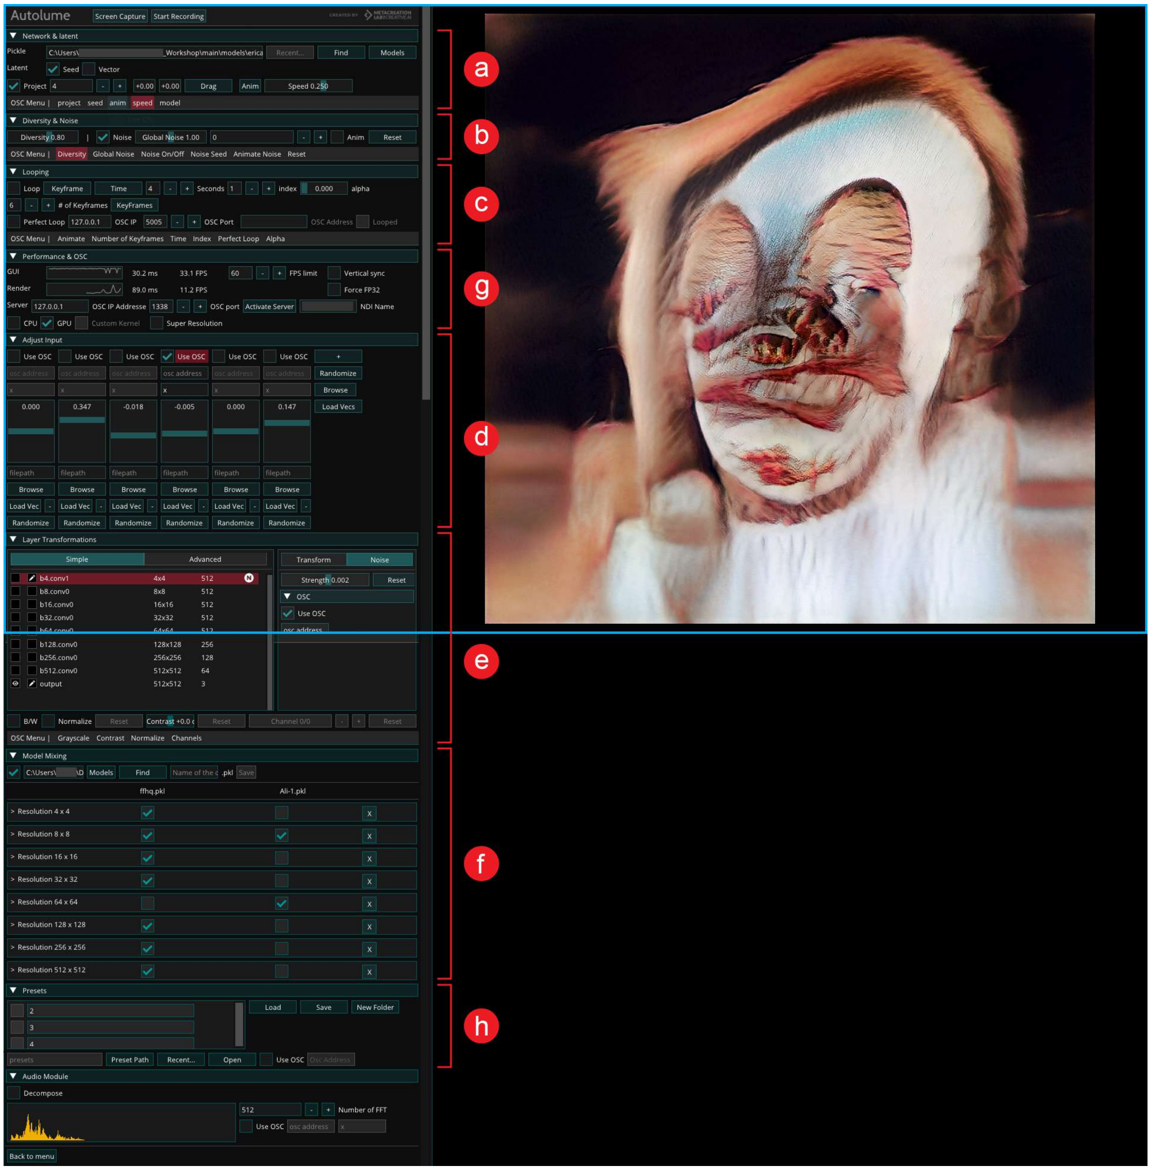

Figure 2: Extended view of Autolume's real-time module: (a) model loading and basic latent space navigation, (b) global noise parameters, (c) latent space interpolation, (d) Semantic feature directions, (e) per-layer model bending and manipulation, (f) layer-wise model mixing, (g) NDI or Network Device Interface protocol setup, and (h) preset saving. The image in the interface is generated from a model by Erica Lapadat-Janzen and Philippe Pasquier. More at: https://www.metacreation.net/projects/dreamscape

3.2 Workshop Design

Two in-person workshops in Toronto and Montreal (October 2024) and one online workshop (April 2025) were held, with 11, 14, and 42 attending, respectively. Each workshop spanned two consecutive days, with each day consisting of a 3-hour session.

3.2.1 Procedure. On the first day, facilitators (and participants during in-person workshops) briefly introduced themselves, followed by two introductory presentations (see 3.2.2 below). The session continued with demonstrations of Autolume's basic features, covering both offline and real-time modules, described in Section 3.1. The session concluded with a demonstration of the model training module, showing how to select datasets, adjust hyperparameters, and monitor training. Facilitators presented sample datasets, trained models, and training process snapshots illustrating successful training progression (e.g., similar to Figure 1). Participants then initiated their own training runs, using either personal datasets or provided samples. These sessions were set to run overnight to produce results.

On the second day, participants optionally shared training results, with both successes and failures discussed for collective learning. Facilitators then demonstrated advanced features, including feature extraction, latent space projection, and interactive control via OSC. The workshop concluded with an optional focus group discussion and an individual survey (see 3.2.3 below). Figure 3 shows pictures from the workshop.

Figure 3: Pictures from the workshops, showing participants discussing the results of training (left), and a presentation showing examples of GAN training (right).

3.2.2 Educational Component. The first presentation outlined recent research by the authors, highlighted issues with large-scale AI systems (some discussed in Section 2.2), and introduced small-data approaches as a potential alternative. Furthermore, the first presentation introduced Autolume and positioned it in relation to other AI systems and presented sample artworks created with it. The second presentation introduced Generative Adversarial Networks (GANs) for a non-technical audience, explaining core concepts such as unconditional generation, GAN architecture, and the training process. Special emphasis was placed on latent spaces, illustrated through artistic examples of latent space navigation to aid comprehension.

3.2.3 Research Component. At the end of the workshop, participants were invited to take part in an optional focus group and post-workshop survey. Those who chose to participate were asked to complete an online informed consent form on Simon Fraser University's Microsoft Forms. The workshops were offered as free educational sessions on Autolume, which is a free and open-source tool. Participation in the research components was optional and took place only after the educational portion was completed. All participants provided written consent prior to engaging in the study. An ethics approval was obtained for all workshops by the REB at Simon Fraser University, Canada.

3.3 Data Collection

Focus groups were employed to elicit impressions from a wide range of artists regarding the concepts underpinning Autolume: small data and model crafting. Rather than measuring direct effects, these discussions aimed to examine artists' considerations regarding small-data approaches to creative work —specifically authorship, creative control, and aesthetics— to inform design opportunities for future AI art tools.

To encourage engagement, groups were kept under eight participants. In-person workshops involved single-group discussions (7 in Toronto, 6 in Montreal), while the online workshop was split into three randomly assigned groups (6, 5, and 2 participants), each moderated by a researcher. A prepared question set (Appendix A.2) guided the semi-structured sessions, with follow-ups used when relevant to the research questions.

Five focus groups were recorded, online sessions via Zoom and in-person sessions with a Zoom H4n portable recorder, yielding 3.5 hours of audio. Recordings were transcribed using a locally run AI service, and the first two authors reviewed them for accuracy and familiarization.

To triangulate our focus group findings, we also administered an individual post-workshop survey. All respondents completed the survey after gaining hands-on experience with Autolume. The participant pool ($N = 25$) included focus group attendees (surveyed in conjunction with discussion sessions) as well as artists who attended only the workshop's educational component. The goal of the survey was not to conduct a summative evaluation, but to provide participants with a private space to share thoughts, thereby counterbalancing potential group dynamics or moderator/authority influence. Finally, the instrument included both qualitative and quantitative items (Appendix A.3) designed to mirror our research inquiries, covering subjective perceptions of usability, creative control, authorship, and aesthetic qualities.

3.4 Participants

We recruited participants for both in-person and online workshops through multiple channels. The Metacreation research lab's email newsletter (with 2950 subscribers) served as the primary announcement platform, while our partner institutions in OCAD University in Toronto, and NAD (École des arts numériques, de l'animation et du design) and SAT (Society for Arts and Technology) in Montreal also circulated the call through their networks. The workshop description emphasized accessibility for artists across all visual disciplines and technical backgrounds (Figure 4), from those exploring generative AI for the first time to those seeking to deepen their engagement with personalized and interactive generative AI workflows. We particularly encouraged artists interested in training AI models on their own datasets, asking them to prepare images or videos of their artworks or other ethically sourced content in advance.

Interested artists completed a registration form about their backgrounds and prior experience with generative AI (Appendix A.1). In-person workshops were capped at 11 participants in Toronto and 15 in Montreal, due to the limited availability of GPU-equipped computers at partner institutions. We prioritized artists with relevant backgrounds and who brought their own datasets. The online workshop had no participant limit, as attendees used their personal computers.

We received 42 sign-ups for the in-person workshops and 76 for the online session. Ultimately, **25** attended in person and **42** attended online. As focus groups were optional and held on the second day, only a subset participated: **6** in Toronto, **5** in Montreal, and **13** online. Participants' artistic practices spanned photography, printmaking, painting, sculpture, dance and animation to work in immersive experiences (AR/VR), audio-visual performances (including VJing), interactive installations, game development, storytelling and creative coding. Most had used at least one LTGM such as Midjourney or DALL-E (Figure 4). While most identified primarily as artists, about half also identified as educators or researchers. In terms of generative AI experience, the distribution formed a bell curve skewed toward higher experience and use: most participants knew the basics and used genAI occasionally, about a third were experienced or frequent users, and a smaller group had little or no technical experience.

3.5 Data Analysis

We employed a reflexive thematic analysis (RTA) approach, as outlined by Braun and Clarke [11], to systematically identify, analyze, and report patterns of shared meaning within our dataset. We chose this method over a codebook approach due to the complex and often conflicted relation of generative AI to artists and their practices, requiring an approach grounded in reflexivity and humility. The interpretive flexibility of RTA allows us to capture the diverse and sometimes contradictory opinions voiced by our participants. We used NVivo 14 for data management and initial coding and Miro for collaborative theme development, as detailed below.

(1) Data Familiarization and Coding: Analysis began with data familiarization. The first two researchers, responsible for data analysis, independently listened to all five workshop recordings to capture nuances and verify transcript accuracy. Following transcription, the data were imported into NVivo 14 for coding. To foster reflexivity and develop a rich, multi-perspective interpretation, we undertook an iterative and collaborative coding process. Both researchers independently open-coded all workshop transcripts, meeting after each to compare our codes, discussing similarities and differences in our interpretations. These discussions helped us refine our understanding of the data and challenge our individual assumptions. This process resulted in over 450 codes when the codes of both researchers were aggregated. This approach enabled us to use our differing interpretations as an analytical tool, revealing nuanced patterns from the granular data (i.e., in Miro) that might have been lost in an earlier reconciliation (i.e., in NVivo 14). During our discussions, we paid close attention to where our interpretations converged and diverged. When a text segment's ambiguity prevented reliable interpretation, we excluded it to preserve the robustness and clarity of our findings.
(2) Collaborative Theme Development: Once the coding was complete, we exported data to Miro for theme development. We eliminated redundant codes and organized the rest into high-level categories: aesthetics, authorship, effort/control, interactivity, model/interface understanding, complexity, small data, training, crafting, hardware, comparisons, plans, bugs, suggestions, liked, and unclear. We then clustered codes by shared meaning to visualize connections and formulate themes. Miro and NVivo 14 facilitated quick reference to the underlying text, ensuring the analysis remained grounded in the data.
(3) Analytical Approach and Refinement: Our analysis combined inductive (data-driven) and deductive (theory-driven) elements: themes primarily emerged from the data, but our research questions guided initial high-level categorization. We analyzed discussions primarily at a semantic level, as the focus group data offered limited context on individual participants. When latent interpretation was required, it was informed by broader debates and developments in generative AI. Theme development was recursive, with continual checks between codes, themes, and quotes to ensure a cohesive, accurate account. The quotations reproduced here were lightly edited for clarity.

3.5.1 Researchers Positionality. The first two authors conducted the analysis in this paper under the guidance of a senior supervisor and an expert in qualitative research.

- **First Author**: I approach this research as a PhD student with a background in computer science and specialized training in Human-Computer Interaction and AI. My prior research involves building tools for various creative practitioners, including game designers, architects, and visual artists. While I am not formally trained as an artist, my personal engagement with creative coding and generative art projects situates me as an "insider" to the practices and discourses of creative AI. This positioning provides me with a shared vocabulary and a deep familiarity with the tools and values of these communities, which inevitably influenced how I formulated research questions and interpreted participant responses.
- **Second Author**: I am a PhD researcher with a background in physics, which allowed me to learn machine learning through technical literature. Over time, I adopted an interdisciplinary approach bridging arts and technology, using AI both as a research subject and as a medium for artistic practice. This dual orientation gives me technical and experiential perspectives as an active AI artist. I have worked extensively with Autolume and am an experienced user. In facilitating workshops and co-analyzing transcripts with the first author, my expertise inevitably shaped how I framed content, guided participants, and interpreted data. While participants' words formed the basis of the analysis, themes were co-constructed through the interplay of their accounts and my interpretive lens.

As members of the Autolume team, the first and second authors possess an intimate technical familiarity with the platform. While this dual role enabled them to provide expert guidance and present realistic, practice-based case studies, their proximity to the tool may have influenced participant engagement and data interpretation. To maintain qualitative rigour and mitigate potential bias, the research process incorporated independent supervisory oversight, collaborative coding, and multiple rounds of discussions. Additionally, an independent follow-up survey offered a space for candid responses away from the influence of the researchers or other participants.

4 Findings

The analysis produced three major themes around impressions of small data, creative practice considerations, and the learning process. Overall impressions of Autolume were very positive (even thankful, some called it *"historic"*), with participants appreciating the easy installation, quick visual feedback and the ability to get personal results quickly. Many highlighted model training as an

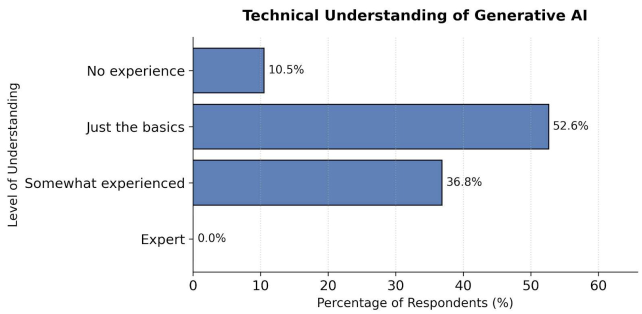

(a) Distribution of Technical Understanding

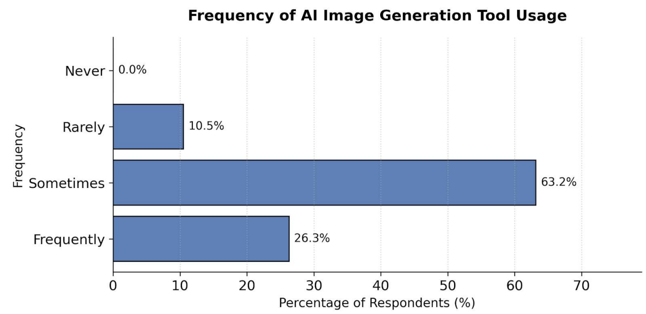

(b) Frequency of AI Tool Usage

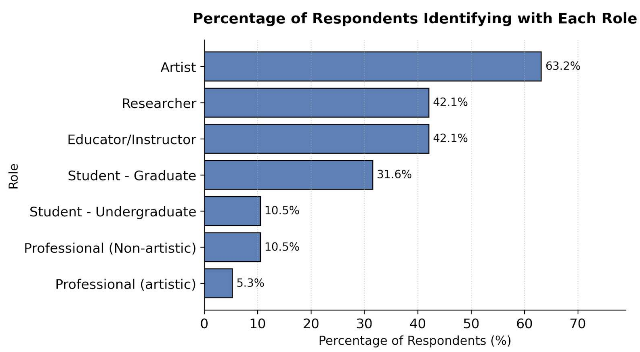

(c) Distribution of Respondents’ Role(s)

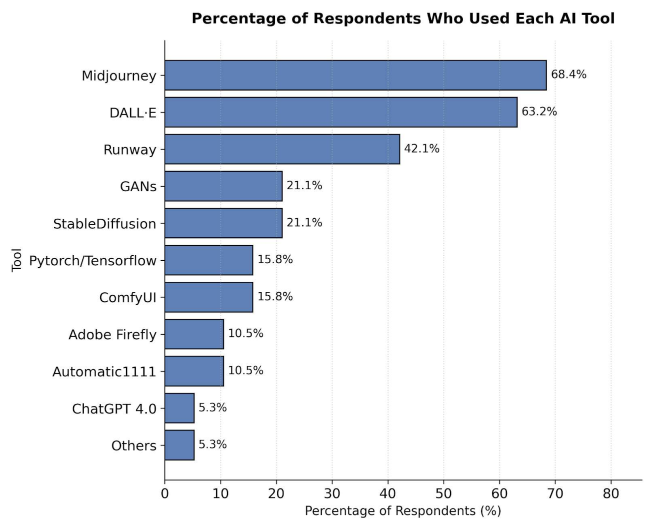

(d) Most Popular AI Tools

Figure 4: The backgrounds of the artists who participated in the focus groups. Covers their current role, and their generative AI expertise, tools of choice, and frequency of use.

advantage over other tools, noting the freedom that training provides, and a recognition that it covers most bases needed for artistic practice (pre-processing, post-processing, and interactivity).

4.1 THEME 1: Small Data and Autolume’s Promise

Many participants viewed "small data" as a strength, linking it to creative nuance and the development of a personal artistic voice. Multiple participants highlighted the ethical benefits of local processing, while several specifically valued the safety of offline systems compared to cloud-based alternatives. Regarding authorship, some felt increased ownership through model training, whereas others valued the partnership felt in text prompting. While one participant questioned the authenticity of AI compared to traditional media, others saw the tool as a way to revisit existing work. Finally, many participants expressed frustration with text systems and preferred direct visual control, though some acknowledged the specific creative control text inputs offer.

4.1.1 "There is beauty in the small". Many participants expressed a generally positive sentiment toward the concept of small data, viewing training with small datasets as a strength and a rich area for exploration, often connecting it to the ability to achieve greater creative nuance. Others associated it with increased freedom, attributing this to the ability to train their own models and explore endless permutations. One participant stated *“there is beauty in the small”*, in reference to advantages they see in small-data approaches in contrast to the prevalent large-scale AI systems.

Participants also linked the small-data approach to the development of a more personal artistic voice, especially in contrast to text-to-image systems, and saw it as a means to produce more specific and nuanced results. In this way, they described model training through Autolume as a tool that could support and amplify artists’ unique voices. One participant remarked: *“what’s nice is* *that it feels like you can get really hyper specific, especially when it comes to thinking about a place or person. It kind of allows you, in some ways, to inject a little bit more of yourself into the model"*. However, regarding the term *small* itself, a few participants expressed uncertainty about the scale of small data *"is it 50 samples to 500? or 50 samples to 50,000? For me, if it's under a billion data points, it's small"*.

4.1.2 Ethical Resonances of "Small". Many participants highlighted the positive ethical dimensions of the small-data approach. Some described it as an *"anti-capitalist"* alternative and a way to address the energy consumption concerns associated with large-scale models. One participant remarked: *"the beauty of the small is the amount of data processing and energy being used by the system, if it's local rather than cloud rendering"*.

The local nature of Autolume was also viewed positively. Several participants expressed concerns about cloud-based AI systems, citing perceived privacy and safety issues related to uploading their data. One participant mentioned: *"It definitely didn't feel safe with uploading pictures online"*. Some noted that they would only upload data already available online, considering local AI to be a safer and more trustworthy option. One participant noted that for political work or high-tension issues, keeping data offline is a matter of *"security, even survival"*.

4.1.3 My Data, My Model! Some participants described an increased sense of authorship, which they associated with being able to train their own models and having greater control over the data and model parameters. One participant mentioned: *"using your own corpus gives one a sense of authorship"*. Another participant said: *"I do [feel more ownership], because I'm controlling what goes into it and how it thinks"*. For some, this sense of authorship was greater than through prompting, while for others it was still not as strong as experienced with coding, as one participant noted:

> *"...the prompt is not my creation. But I don't have this feeling in your product because I have the capability to train the machine by myself."*

However, it was noted that prompting can offer more control over outputs and foster a stronger sense of partnership in the creative process, whereas Autolume was often seen as providing machine learning as a tool rather than as a competitor to artists.

> *"It's an interesting balance between prompt pasting, where I feel like I have more control over the output being created. I feel a little bit more involved in the process of what comes out. But with small data, it feels like what I'm putting into it is really my own. So I guess it's on which level of interaction does ownership come in?"*

4.1.4 A Qualified Authorship. A sense of authorship was also associated with participants' understanding of model training and with having more control over inputs. As one participant put it: *"I think understanding how to ... get that sense of authorship, I think it's understanding how to train better first"*. At the same time, at least one participant framed authorship as *"distributed between the model and the outputs"*. Finally, one participant suggested that Autolume should take a clear stance on the legalities of authorship to help users "align with the ethic" of the tool, noting that a clearer understanding of intellectual property and the legalities surrounding authorship could influence their perception of it.

4.1.5 The Old, the New, and the In-between. One participant reflected on the authenticity of creating with AI, comparing it to more traditional media. Although the group expressed general excitement about AI, this participant questioned whether it can support a genuinely authentic form of expression. They also described a perceived gap between AI (and other computer-mediated) art and traditional art, and tended to associate greater authenticity with the latter.

> *"From an artistic point of view, I'm still struggling with whether it feels genuinely authentic to me deep down, as perhaps traditional painting."*

While some questioned the authenticity of the experience of creating with AI in contrast to traditional media, others saw in model training an opportunity to revisit an existing body of work or a previous artistic practice, or expressed a desire to revisit a traditional medium, knowing that they could feed what they created to Autolume.

> *"It has definitely made me want to sketch more on paper with watercolour. I really miss working like that and generating new things with Autolume, and exploring new things that I can create with that."*

These participants saw Autolume as offering new potential for derivative work based on their existing artworks (e.g. recorded videos or paintings) through latent space exploration or the discovery of connections between data points through navigating the "in-between" spaces in the latent space.

> *"If we can use our own data, we can explore kinds of connections between that. Maybe reveals some of the gaps within that as well, or that we can explore it sort of critically."*

4.1.6 Direct Manipulation is Worth a Thousand Words. Frustration with prompt-based systems was a shared sentiment among many participants. They described challenges such as lengthy prompting, the difficulty of precisely describing what they meant, and misalignment between inputs and outputs. Many participants expressed positive feedback on Autolume's direct manipulation controls and indicated a preference for the interface's interactive and visual control over prompting. There was a recurring sentiment that refinement through prompting reached a point of diminishing returns when additional prompting no longer improved results.

> *"With the large language models, I feel like there's a ceiling where, after a certain point, no matter how much you rearrange your text prompt, it's not going to get any better or closer to what you want. And I didn't feel that here."*

However, some participants mentioned the level of control possible over the outputs in prompt-based systems (*"I want to bring more details, I can add keywords"*) and described them as offering a stronger sense of co-creation.

4.2 THEME 2: Creative Practice and Artistic Potential

Many participants adopted the term "crafting" to describe their practice with AI and showed excitement for it, while emphasizing the importance of community resources. In terms of aesthetics, several described GAN visuals as familiar "morphing" that Autolume makes accessible. Others felt unable to form a conclusive opinion due to limited experience or a desire for longer engagement, while some argued that using the tool is itself a distinct aesthetic choice. Participants often appreciated the visual cohesiveness of the results, though one noted a trade-off in reduced diversity compared to large-scale models. Several participants highlighted the advantage of real-time interaction, particularly in combination with small data, while many saw potential in the OSC feature for interactive installations.

4.2.1 Model Crafting as Evolutionary, Creative, and Communal Practice. Participants found "crafting" a rich analogy for imagining possibilities with Autolume, and multiple participants adopted the term in their responses, often extending it to encompass both model crafting and dataset crafting. One participant described model crafting as an evolutionary process, where models can be mixed, mutated, and branched. Model mixing, in particular, was described as *"smart"* and enabling *"going beyond imagination"*. Mixing and bending models were highlighted by some as especially promising for gaining greater creative control. Participants also emphasized the importance of community and shared resources, such as models, datasets, and metadata. They drew parallels to existing LTGM repositories where artists share images along with the prompts, models, and parameters used to generate them. Similarly, participants expressed interest in accessing others' training hyperparameters, seeing it as useful in successfully training their own models. Some viewed access to trained models as a key step in lowering barriers to entry and making the remaining features on the interface more accessible. Beyond technical exchange, participants highlighted the value of community for learning, noting that seeing others' work helped them in thinking about their own process. Finally, some envisioned *co-operative* model training as a way to overcome hardware limitations and reduce reliance on *"corporate structures"*.

4.2.2 GAN Aesthetics: Familiar, Yet Unsettled, and a Choice. Participants described Autolume's aesthetics as resembling *"morphing"*, *"interpolation"* or *"fractals"* –a familiar, predictable look, but one that Autolume made accessible and controllable: *"The artist has more control over lots and lots of parameters, which they wouldn't have had before"*. At the same time, several participants noted that they could not yet form a conclusive opinion on Autolume's aesthetic, fearing that such judgment might *"be born of inexperience"*. We acknowledge these views may have been shaped by the workshop materials, which largely featured morphing visuals and smooth interpolations.

Some participants associated Autolume with a distinct GAN aesthetic, to the point that, according to one participant, working with Autolume constitutes a creative aesthetic choice in itself.

> *"It feels like this is almost just a different medium ... I make the decision to use this, it's saying a very specific thing visually"*

4.2.3 The Trade-off Between Cohesiveness and Diversity. Several participants appreciated that the results appeared visually cohesive and consistent with their training datasets, as one participant stated: *"It's so interesting that Autolume creates the same works with the consistent colour palette and emerge them. And somehow it saves the textures for you. And it's so interesting to see them consistently emerge in one another"*, while another appreciated how they could control the diversity of the results by varying the training set:

> *"Then if we wanted to add more dynamics to it, we bring in other images with different compositions. So I think the idea that we can use our own data is really nice."*

One participant noted a trade-off: while small data allows for achieving specific outputs, the diversity is not "comparable to the kinds of diversity of outputs I can get from ChatGPT."

Participants also expressed a general concern that AI tools like Autolume, when constrained for ease of use, risk producing homogeneous results. One participant noted that a tool's signature style, while *"special"*, *"could lead to similarity in the long run"*. Another articulated a trade-off they perceived: design constraints intended to make a tool *"easier"* might ultimately risk funnelling many artists toward similar-looking work.

4.2.4 The Best of Both Worlds: Small-Data and Direct Manipulation. Several participants appreciated the real-time interaction in Autolume and the many ways it allowed them to explore a trained model. They described the ability to navigate the latent space interactively as a key advantage over other approaches, such as code-based GAN training/inference or the now-discontinued GAN module in Runway.

> *"Yeah, I think the most important part of this program is the connectivity with interaction. Because in that case [Runway], you don't have this possibility and it's very important for like visualization."*

This direct manipulation was particularly valued by some in contrast to prompting systems. One participant stated: *"I really liked the affordances of the control over the dataset and being able to quickly get feedback with the way you've parameterized everything"*. Another participant mentioned: *"I mean, you can just change anything with a click of a button or dragging some items that seem to be working at the end"*. Finally, participants felt the benefits of real-time interaction were especially clear in combination with the small-data approach.

> *"If you're trying to understand what's going on with your model, you can build your small data model or your sort of boutique model, and then be able to see those changes right away."*

4.2.5 Autolume in Action: Inspiring Performances and Installations. The OSC feature sparked many ideas among participants, who saw it as a way to create interactive installations, enable embodied latent space navigation (e.g., through dance), or position Autolume as an intermediary interface that communicates with other software. Many of the ideas participants mentioned for future use focused on Autolume's real-time and interactive capabilities, including concepts for live visualizations and its integration into installations, performances, and multimodal setups where different inputs could be connected through OSC to control the generated outputs.

4.3 THEME 3: The Learning and Training Process and Its Challenges

Several participants found the underlying technical terminology difficult to grasp, though many credited the workshop materials with clarifying core concepts. Some felt the workshop lowered entry barriers, while others emphasized the value of peer learning. Participants with teaching roles suggested the tool could support pedagogy in their own classrooms, particularly because the interface components map to the ML process. To improve the training experience, participants proposed features such as tooltips and glossaries, along with clearer system feedback. Many also noted that training takes significant time, which could present a barrier to adoption; at the same time, some felt this time for experimentation was necessary to learn how to align the ML process with their creative goals. Finally, some drew on analogies to make sense of the interface and the ML process behind it.

4.3.1 Building Confidence through the Learning-Doing Cycle. Several participants mentioned that the underlying ML concepts and technical terminology were difficult to grasp. One participant noted: *"I don't understand any of the terminology. It is all gobbledygook in a way"*. Another participant mentioned: *"I haven't gone through trial and error to find out the difficulties, not just of Autolume, but actually the AI concepts behind it, which are sometimes quite sophisticated"*. It was also mentioned: *"I don't agree that the UI is that much easier. I mean, you have to be actually aware of the whole process behind the machine learning"*.

Many participants said workshop materials (Section 3.2.2) clarified core ML concepts (GANs, latent space). Some have used Autolume before, and the workshop improved their use of the tool. They emphasized that understanding these concepts was key to navigating the interface and credited the workshop material for contributing to that understanding.

> *"I've been learning more intellectually what machine learning is and all these terms, GANs, VAEs, whatever. I've heard of them, but I don't know what it means. But then with this workshop, I can kind of piece it together."*

Some participants even noted that the workshop helped lower the barrier to getting started with Autolume. As one participant stated: *"It just seemed incomprehensible without having had this training"*. The workshop gave these participants the confidence to experiment with the system. Similarly, others noted that seeing the work and perspectives of fellow participants enhanced their own learning and influenced their creative process. One participant stated: *"it was really important to see how other people ... were thinking, to change how I was thinking about my own process with the software"*.

4.3.2 Autolume as a Pedagogical Tool. Some participants said Autolume clarified ML concepts, as its interface embodied the algorithm: *"I find it fascinating how you embody this whole machine learning algorithm that is behind this"*. Participants with teaching roles suggested Autolume could serve as a pedagogical tool for teaching latent space or integrating into courses they teach, noting that it bridges theoretical and practical ML knowledge, with one planning to use it for "collaborative art projects" due to the quick turnaround time.

4.3.3 ML as Material: Hard to Train, Easy to Play with. The dependency on understanding machine learning concepts was especially evident when it came to model training. Many participants found the training process (including dealing with training data) challenging overall, describing it as confusing, mysterious, or generally less clear than working with models in the real-time module. As an example, one participant said: *"... for me, it's the training model part, which was like, oh, what's going on?"*. Some also noted that other processes and functionalities became clearer after going through model training once.

Participants suggested that additional training examples and discussions of failure cases would help them better grasp the training process. To improve explainability, they proposed including default or recommended hyperparameters, tooltips, links to resources, and diagrams or glossaries for machine learning concepts like training and latent space. They also recommended novice-expert modes to support varying levels of expertise.

To support model training, one participant envisioned a script that could scan datasets and suggest optimal subsets for training, while another expressed a need for system feedback to indicate whether the training process was progressing effectively.

Finally, multiple participants described the real-time module as easier to grasp than training, as it enabled experimentation without deep technical knowledge. Training, by contrast, required additional conceptual understanding.

4.3.4 It Takes Time, and I Need Time! Multiple participants noted that both training and learning to control it required significant time. Some were surprised by its length, and others saw it as a barrier to adopting Autolume: *"... if it takes two weeks of running my computer, you know, every night before I see results, then I don't think it's like a jump on the bandwagon kind of moment"*. They also raised concerns about limited GPU access for some artists as an equity issue.

Several participants mentioned they needed more time with the interface before judging its aesthetics or creative control. One participant stated: *"... there's so many more opportunities there that I can't judge it simply off of my inexperience, but there are many more aesthetic possibilities than I think"*. Some participants found it difficult to steer training toward desired outcomes and felt they needed more experimentation and experience with training and data curation to better align the process with their creative goals and explore the system's affordances. This reflected an overall learning curve: *"it takes a little bit to understand what you can do with it"*.

4.3.5 Analogies for the Unknown. Some participants utilized analogies to understand the interface, likening the sliders to DJing or audio layering. Overall, participants found analogies useful for understanding both the interface and the underlying concepts, as one mentioned: *"Anytime there's a practical, physical example that could be an analogy that would really help people on board"*.

The analogy of crafting was seen as a rich framework for imagining possibilities with Autolume. Many participants adopted the term *"craft"* in their responses, using it to refer to both data and model crafting. Model crafting was also described by one participant as an evolutionary process, where models are mixed, mutated, and branched. The discussion around small data and crafting led some participants to speculate about the future of *"artisanal datasets"* or *"craft AI"*.

> *"... this is the flavour of craft AI from 2025. This is what ... artisanal dataset would look like that you wouldn't have a company generating."*

4.4 Survey Results

Post-workshop survey results largely corroborated the qualitative analysis from the focus groups.

4.4.1 Qualitative Questions. The survey results generally align with the qualitative themes, highlighting both the perceived benefits of small-data practices and the practical frictions of working with them. Consistent with the findings regarding increased sense of authorship in 4.1.3, many respondents valued the *"creative ownership"* and *"precision"* associated with training on personal datasets, and several contrasted this with the more opaque nature of mainstream text-to-image tools, which some explicitly described as a *"different medium"*. The survey also mirrors the pedagogical dualities in Theme 3: while participants praised the system for *"lowering the entry barrier"* and demystifying ML concepts (highlighted in 4.3.1 and 4.3.2), they also reported a *"learning curve"* around dataset selection and difficulty anticipating or diagnosing training failures, reinforcing the idea of model training as a major challenge in 4.3.3. Finally, consistent with the need to spend more time with Autolume in 4.3.4, respondents frequently withheld judgment on aesthetics due to a need for extended engagement.

4.4.2 Quantitative Questions. Participants reported a strongly positive experience with Autolume. Similar to the overall positive sentiment toward the small-data concept reported in 4.1.1, 56% of participants reported Very Satisfied and 44% Satisfied with Autolume, with no dissatisfaction reported. Aligned with the training challenges highlighted in 4.3.3, the biggest friction point was Training Models (45.8%). This was followed by Model Mixing (25%) and the OSC/Audio Module (25%) as other challenging features. This also aligns with more mixed usability ratings: only 40% found the interface Intuitive/Very Intuitive, while most rated it Moderate (44%) and 16% Complex. Even so, ease-of-use skewed neutral-to-positive (50% Neutral, 42% Easy/Very Easy, 8% Difficult), and participants reported strong creative outcomes, with 72% agreeing they had creative control and 80% feeling authorship over outputs, which is in agreement with the increased sense of authorship reported in 4.1.3. Perceived fit and value were also high: about two-thirds said the tool aligned with their artistic practice (33% Neutral, none negative), most felt outputs were at least Somewhat Aligned with their aesthetic vision (64% somewhat, 20% well), and 83%+ agreed it was useful, culminating in a strong adoption signal, with 92% saying they would integrate Autolume into their practice, reinforcing the inspirational role of Autolume mentioned in 4.2.5.

Overall, the post-workshop survey results are largely consistent with the findings from the focus group discussions, particularly around authorship, creative control, and the challenges of model training, with no major discrepancies observed. While the survey reinforces key themes, it does not introduce significant new insights. The primary additional contribution lies in more explicit feedback on interface usability, highlighting areas for potential refinement in ease of use and intuitiveness.

5 Discussion

In this section, we interpret our themes in light of prior work as they relate to our research questions: Theme 1 "Small Data and Autolume's Promise" mapping to **RQ1**; Theme 2 "Creative Practice and Artistic Potential" mapping to **RQ2**; and Theme 3 "The Learning and Training Process and its Challenges" mapping to **RQ3**.

5.1 RQ1: Impressions of Small Data and Model Crafting

There is beauty in the small, as some participants stated, but the beauty may be hidden by the demanding and technical nature of model training.

5.1.1 Personal, Unique. The concept of small data resonated strongly with artists in our study as an appealing alternative to large-scale models. They valued training on their own datasets for emphasizing authorship and enhancing creative control, noting that models crafted from small data reflected their unique aesthetic and context. This aligns with prior arguments that giving creators ownership of training data and models can strengthen their sense of agency in art-making [2, 16].

5.1.2 Locality, Privacy, and Ethics. Ethical concerns strongly shaped artists' views of small-data AI. While small data does not inherently require local training or inference, Autolume's local and small aspects often intersect in participants' reflections. Training on personal computers with personal datasets was seen as a way to avoid exploitative data practices and power asymmetries of large-scale AI. Some participants referenced negative experiences with cloud platforms, such as Runway discontinuing GAN training and deleting users' data, which led to the loss of generated work and original datasets (For more, see [56]). However, participants highlighted accessibility challenges in local, small-data workflows. The need for high-performance GPUs and long training times discouraged spontaneous experimentation, raising equity concerns by potentially excluding the same marginalized groups this approach could empower. Although the GAN models in Autolume can run on lower-end hardware relative to the hardware requirements of diffusion models, slower training may shift the experience from fluid exploration to a constrained process. Ultimately, tool designers should address hardware and knowledge barriers to ensure small-data AI remains an inclusive alternative.

5.1.3 Challenges to Small Data. The small-data approach presents notable challenges in model training. Despite Autolume's no-code interface, participants found the process conceptually and technically demanding due to a lack of understanding of core ML concepts (e.g. how GANs work and how to effectively train them). A simple user interface that abstracts complexities away may not lessen this knowledge barrier for non-technical artists. These findings align with prior research identifying ML as a difficult "design material" [27] and warning of co-creative "pitfalls" without system transparency [18]. Furthermore, training generative models even on small datasets remains a time-intensive and computationally demanding task, requiring access to high-end GPU hardware for the necessary iterative refinement. The above-mentioned technical difficulty and time/compute demands are heightened by Autolume's use of GANs, which are notoriously unstable to train and prone to mode collapse [82], which potentially requires more iterations to obtain desired results. However, while Autolume currently uses GANs, small-data approaches are not tied to them, and future generative models that are easier or faster to train, while offering compelling aesthetics, may become a new *de facto* standard.

5.1.4 What is Small? Some participants questioned what counts as small, noting that *"for me, if it's under a billion data points, it's small"*. From a computational perspective, this assessment may indeed depend on factors such as the number of trainable model parameters or the amount of available compute. However, small data relates less to dataset size and more to the personal or artistic purpose behind collecting and using a dataset. This is evidenced by the workshop study by Bryan-Kinns et al. [16] with musicians, which suggests that small datasets are composed of self-consistent samples, are not defined by genre, and are not compiled for generalization purposes. Similarly, an art project focused on generating Arabic calligraphy [77] would be considered small data even if the dataset contains over 100,000 items, because it is designed to foreground underrepresented art forms within the context of large-scale AI.

5.1.5 Crafting Communities of Practice. Artists in our study were interested in engaging with a community around crafting. For example, they were excited by the idea of sharing models and datasets to collaboratively expand their creative toolsets. Multiple participants suggested features such as a community model repository, reflecting a desire for *cooperative* training, support, and learning from others' experiments, particularly around how to train models successfully. This emphasis on sharing echoes the role of community in DIY and maker cultures [49], and prompt-based art practices [22, 52] (e.g., success recipes like prompts and prompt templates). At the same time, ML datasets can be sensitive because they can be learned from and incorporated into models in ways that are difficult to undo. It is therefore understandable that participants both emphasized the safety and ethical benefits of local, small-scale processing (Section 4.1.2) and expressed a strong desire for community and sharing (Section 5.1.5). In Bryan-Kinns et al. [16]'s study, participants wanted dataset sharing primarily for learning and exploration rather than for model training. Prior work examining online digital art distribution similarly highlights the role of communities and stewards in encouraging reuse that returns value to artists [6].

5.1.6 No-Code Crafting. Participants also showed interest in advanced crafting techniques in Autolume, such as model mixing and model bending (Section 4.2.1), which can be performed without additional model training, through a visual interface with real-time feedback. Participants saw them as promising avenues for artistic exploration, but had limited time during the workshop to experiment with them in depth. Future studies could examine no-code, low-level manipulation of model weights as additional tools for creative exploration. This includes mixing [63], bending [3, 14], grafting [21] models or Hidden Layer Interaction [36]. These approaches support tinkering with AI models as material and can enable creative expression with little or no model training.

5.2 RQ2: Artist Considerations for Small-Data AI Tools

We presented Autolume as a small-data tool, then anchored our discussions around key concepts (Appendix A.2) that our past engagements with artists and the literature revealed as pertinent to creative practices with AI. The discussions show a diversity of opinions around authorship and creative control, strong stances around ethics, a preference for local and low-resource AI tools, a desire for more explainability, and an attunement to the distinctive aesthetics of AI tools.

5.2.1 Control and Surprise. In our study, many artists leaned toward maintaining control, valuing personal input, yet they also acknowledged what might be *"lost"* in the small-data approach: the expansive freedom that large-scale prompt-based systems afford. A few even described prompting large-scale models as a kind of partnership yielding unexpected ideas, versus the more deliberate, author-driven interaction with Autolume. In this context, under-specifying prompts can introduce surprise, with the system generating outputs beyond the artist's initial intent. Prior work has examined *prompt journeys*, where prompting is approached as an iterative craft-like process of writing, evaluating, and refining prompts [52]. Drawing a parallel, both small data and model crafting, as well as prompt crafting, highlight how artists negotiate between exerting control and leaving room for serendipity. This highlights a key design tension for creative AI tools: how to support artists in asserting control and personal style, while still enabling moments of surprise that can fuel creativity. This echoes prior accounts of striving for a balance between control and surprise in co-creation [65].

5.2.2 Authorship. Participants' views of authorship were closely tied to control. Many felt that training on their own images strengthened their claim over the outputs, with some stating that they felt more ownership of Autolume's results than of images produced by prompting large-scale AI models. Notably, a few contrasted Autolume with coding or traditional art. While they felt that small-data and model crafting improved their sense of authorship relative to generic AI outputs, it still did not always match the complete creative ownership they experience when coding algorithms or working with traditional media such as painting. These perspectives suggest that authorship in AI art might be varied, emerging at different *"levels of interaction"*, that of artifacts and models.

5.2.3 Authenticity. The sub-theme on "The old, the new, and the in-between" describes contrasting views on AI, one doubtful of its ability to enable genuine expression as with traditional media, and another motivated to revisit traditional media or their previous body of work through AI. Similarly, artists in Rajcic et al.'s study [67] saw in personalized generative models a way to offer an "alternative interpretation of their visual aesthetic" and to capture a "small piece of their artistic identity". These contrasting views raise a question of how the "genuine" sense ascribed to traditional media can be preserved within the increasingly mediated practices of AI art [23].

5.2.4 Aesthetics. AI techniques are not aesthetically neutral once they are embedded in AI art tools (e.g., Autolume as a distinctive aesthetic choice in Section 4.2.2). AI-generated visual aesthetics can arise from (1) interaction, (2) training data, and (3) the training process itself. First, aesthetics can arise from interaction. Even with the same model, different ways of navigating a latent space can produce different visual and temporal qualities. For example, in response to the smooth latent spaces of GANs, some participants wanted non-linear interpolation controls, akin to an animation tool, while others wanted percussive jumps that break smoothness. These changes do not modify the model, only how its latent space is traversed. Second, aesthetics can come from training data. Generative models can learn, mix, and reinterpret an artist's visual aesthetic, especially when trained or fine-tuned on artist-curated datasets. Third, aesthetics can be shaped by the training process itself. For example, some generated visual features can be linked to training dynamics, such as GANs' morphing and visual indeterminacy [40], or LTGMs' tendency to favour representational over abstract imagery [88] due to their text-conditioning. Combining small-data training and interactivity, desired by some participants (Section 4.2.4), influences visual aesthetics on the three levels, and can enable things like morphing through familiar data points for what is in-between them (Section 4.1.5).

5.3 RQ3: Design Opportunities

In what follows, we examine artists' impressions of small data and model crafting (RQ1), and the creative practice considerations that surfaced during their engagement with Autolume (RQ2), to derive design opportunities for small-data AI art tools.

5.3.1 Combining Visual Direct Manipulation, Low and High Level Controls, with Small-data Training. Artists in our study expressed frustration with prompting systems (Section 4.1.6), which was similarly voiced in other studies of text-to-image art practices [52]. Instead, they preferred direct, visual manipulation over text or code, favouring embodied and tangible navigation of latent spaces; a shift in practice also advocated by McCormack et al. [55] and explored in the design research of Lindley and Whitham [50]. Furthermore, our participants appreciated low-level control over visual characteristics such as colour palettes, textures, or noisiness, similar to the granular controls offered in traditional software like Photoshop or generative art frameworks like Processing or TouchDesigner, alongside the high-level intentional authoring controls afforded by prompting systems or Autolume's Feature Extraction module. Furthermore, many participants shared project ideas and expressed excitement over the potential of using Autolume and latent space navigation in interactive installations (Section 4.2.5). The combination of training models on small personal datasets with real-time interactivity was seen by some as particularly promising (Section 4.2.4). The combination of the above ideas may lead to personal, meaningfully controllable generative AI systems for AI art.

5.3.2 Surfacing the Generative Process. Our findings in Theme 3 revealed that the "black box" nature of AI remains a barrier, even with no-code tools. Structuring the tool's interface to mirror the ML process —data preparation, training, and exploration— resonated with participants and made the interface more legible to them (Section 4.3.2). This exposes a core design trade-off between simplicity for novices and control for experts. Overly simplified interfaces may lower the entry barrier and reduce the chance of failure, but might also reduce output diversity (Section 4.2.3) or hide the underlying materiality essential for sustained artistic engagement [81].

This tension highlights the relevance of the nascent field of explainable AI for the Arts (XAIxArts) [31], which advocates for artist-centred rather than technocentric approaches to XAI [17]. Within this framework, artistic practice itself becomes a pathway to explainability [17]. By treating generative AI as an "artistic material" to be hacked, broken, and explored, artists can reveal the inner workings and biases of these systems in tangible ways [12], not just for the artist but for their audience as well [39].

5.3.3 Integrating Multiple Forms of Creative Control. Our findings frame authorship in AI art as a spectrum. Participants reported a stronger sense of authorship through data curation and model training, noting that *"using your own corpus gives one a sense of authorship"*. At the same time, some observed that prompt-based systems can offer finer control over specific outputs. This suggests a distinction between authorship at the model level (small-data training and model crafting) and at the instance level (e.g., prompting, latent-space navigation), leading one participant to ask, *"on which level of interaction does ownership come in?"*. We therefore suggest that designers should support fluid movement between these forms of control, for example, by integrating model crafting and small-data personalization into new or existing creative tools [2]. Both our study of artists' impressions of small-data workflows and Rajcic et al.'s analysis of artists' practices with personalized LTGMs trained on small data [67] suggest that artists prefer personalized models over generic ones. However, difficulties in successfully training these personalized models, as described in Theme 3, may confound artists' sense of authorship (Section 4.1.4).

Integrating multiple levels of control –specifically at the instance and model levels– allows artists to navigate, adapt, and craft generative spaces [1]. This heightened engagement may enhance an artist's perception of their creative labour, preserving a sense of authorship that might otherwise diminish if the output feels too effortless [67]. Furthermore, by making these controls reliable (ensuring artists can successfully create or adapt models (cf. Section 4.1.4), without getting lost in the models' latent space or being frustrated by unpredictable prompting [52]), AI art tools may be able to foster a more robust sense of ownership in creators.

5.3.4 Catering to the Broader AI Art Ecosystem. Our findings highlight the need to situate AI tools within the broader artistic ecosystem. Artists sought support across the small-data workflow—from dataset curation and training guidance (Section 4.3.3) to embodied latent-space navigation (Section 4.1.6) and communal sharing of models, data, and training experiments (Section 5.1.5). A recent study [6] with visual artists (including AI/data artists) similarly frames creative practice within a creativity supportive ecosystem, emphasizing how sociotechnical systems shape not only creation but also sharing, provenance, and collective negotiation around disruptions such as generative AI. Designing for artists requires balancing needs-driven and affordances-driven approaches: while no single tool can serve every practice, public models, open-sourced tools, and shared community resources can enable exploration and appropriation, as observed in other studies [19, 66] with AI artists.

5.3.5 Small-data AI Tools for Timely Learning. Educators in our study saw a potential for using Autolume and small-data approaches in their teaching practice. Beyond being creative instruments, tools like Autolume might also foster critical AI literacy. Whereas public understanding of earlier ML systems grew gradually through community initiatives and accessible tools (e.g., Wekinator [30], AI literacy competencies [51]), generative AI emerged rapidly to the public as opaque black boxes. By engaging users directly in dataset curation and model training, small-data and model crafting tools may turn generative AI from a commodified black box into a tangible, malleable material [3].

6 Limitations and Future Directions

The participants' recruitment call was open to artists from all visual disciplines and technical backgrounds who are interested in exploring generative AI or deepening their engagement with it. However, as shown in Figure 4, most participants already used generative AI occasionally or were familiar with specific AI tools. Consequently, this group is not representative of the broader artist population, as many artists abstain from AI engagement on principle due to concerns such as sustainability and exploitative data scraping. On the other hand, we may expect similar responses from artists whose existing practices already align with algorithmic or data-driven methodologies, regardless of their specific creative discipline.

We used focus groups to elicit diverse, rich conversations with visual artists from varied backgrounds. This approach helped us understand their impressions of small-data and model-crafting practices within the complex, entangled AI art space. However, focus groups can make it practically difficult to attribute viewpoints to specific individuals (especially when consensus was expressed through non-verbal cues) or to systematically compare responses across workshops (given variability in the questions asked). Several participants also refrained from judging Autolume's aesthetics, for example, noting that they would need more time with the system before forming an opinion. This underscores the value of longer-term collaborations with artists. Such engagements, however, require substantial artistic investment, especially when participants are expected to curate data and train models independently, which can make recruitment challenging.

Training low-resource models on small datasets and exploring their latent spaces is one mode of creating visual artifacts. While this mode foregrounds data curation and model training in ways that align with our study, it offers limited control and visual fidelity. Future work could examine small-data and model-crafting practices in other creative modalities, such as training LoRA adaptations of LTGMs and building end-to-end generative AI workflows.

This study focuses on a specific small-data approach enabled by Autolume, which foregrounds model training and data curation coupled with latent space exploration. While suitable for interactive arts and abstract or ambient visuals, this approach lacks, on its own, the granular control afforded by large-scale generative models (e.g., regional editing, style references, or prompting). Future research, following Rajcic et al. [67], could examine how artists engage with small-data practices when integrated into large-scale AI workflows via methods such as fine-tuning or model conditioning.

Some results of the study can be understood as specific reactions to GAN models and their distinctive morphing aesthetic (Section 4.2.2), while others are tied to the Autolume interface (Section 4.2.5; Section 4.3.2) or to a combination of both (Section 4.3.4). Impressions of small-data and model crafting, however, seem to point to broader conceptual shifts in AI-driven creative practice and may therefore remain relevant as artists transition to other generative models and creative AI tools.

7 Conclusion

We presented a qualitative thematic analysis of five focus groups conducted with 25 artists to elicit their impressions of counter-initiatives from the art world to large-scale generative AI, namely small-data and model crafting, through a probe that embodies these concepts (Autolume). The results show an excitement for small data, including its potential for creating specific and personal results and its alignment with artists' ethics. Training was seen as a barrier to small data, requiring better explainability and support. Considerations around control, authenticity, and aesthetics surfaced in relation to small data and AI tools in general. We conclude by providing design opportunities for AI art interfaces: support for direct and low-level visual control, mapping UI to ML processes, integrating small data and catering to its broader ecosystem, and using small data as a pedagogical tool.

Small data is not a new concept, but a general approach to creative work with AI. Artists have been creating, training and adapting generative models long before the rise of LTGMs. However, the current landscape of large-scale AI (discussed in the Critiques section) makes this approach particularly compelling today and helps explain the enthusiasm expressed by artists in our study. Due to this larger context, small data may shift from a fringe experimental practice used by early AI art pioneers to a core method in contemporary, artist-centred creative work with AI.

Acknowledgments

We would like to thank all the study participants for their time, as well as the reviewers for their constructive feedback on this paper.

References

[1] Ahmed M. Abuzuraiq and Philippe Pasquier. 2024. Seizing the Means of Production: Exploring the Landscape of Crafting, Adapting and Navigating Generative AI Models. In *the 3rd Generative AI and HCI Workshop*. Hawaii, USA.

[2] Ahmed M. Abuzuraiq and Philippe Pasquier. 2024. Towards Personalizing Generative AI with Small Data for Co-Creation in the Visual Arts. In *HAI-GEN 2024: 4th Workshop on Human-AI Co-Creation* (2024). ACM, South Carolina, USA.

[3] Ahmed M Abuzuraiq and Philippe Pasquier. 2025. Explainability-in-Action: Enabling Expressive Manipulation and Tacit Understanding by Bending Diffusion Models in ComfyUI. In *In Proceedings of Explainable AI for the Arts Workshop 2025 (XAIxArts 2025)*.

[4] Adobe. 2026. Adobe Firefly. https://www.adobe.com/products/firefly/features/text-to-image.html.

[5] Memo Akten, Rebecca Fiebrink, and Mick Grierson. 2019. Learning to see: you are what you see. In *ACM SIGGRAPH 2019 Art Gallery* (Los Angeles, California) *(SIGGRAPH '19)*. Association for Computing Machinery, New York, NY, USA, Article 13, 6 pages. doi:10.1145/3306211.3320143

[6] Shm Garanganao Almeda, Joy O Kim, and Bjoern Hartmann. 2025. Creativity Supportive Ecosystems: A Framework for Understanding Function and Disruption in Online Art Worlds. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 267, 17 pages. doi:10.1145/3706598.3713734

[7] AUTOMATIC1111. 2026. Stable Diffusion Web UI. https://github.com/AUTOMATIC1111/stable-diffusion-webui.

[8] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) *(FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 610–623. doi:10.1145/3442188.3445922

[9] Charlotte Bird. 2024. Artists and AI: Creative Interactions and Tensions. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–6. doi:10.1145/3613905.3651041

[10] Black Forest Labs. 2026. Flux.2. https://bfl.ai/models/flux-2i.

[11] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.

[12] Terence Broad. 2024. Using Generative AI as an Artistic Material: A Hacker's Guide. In *In Proceedings of Explainable AI for the Arts Workshop 2024 (XAIxArts 2024)*.

[13] Terence Broad, Sebastian Berns, Simon Colton, and Mick Grierson. 2021. Active Divergence with Generative Deep Learning–A Survey and Taxonomy. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)*.

[14] Terence Broad, Frederic Fol Leymarie, and Mick Grierson. 2021. Network Bending: Expressive Manipulation of Deep Generative Models. In *Artificial Intelligence in Music, Sound, Art and Design: 10th International Conference, EvoMUSART 2021, Held as Part of EvoStar 2021, Virtual Event, April 7–9, 2021, Proceedings 10* (2021). Springer, 20–36.

[15] Nick Bryan-Kinns, Phoenix Perry, Elizabeth Wilson, and Anna Wszeborowska. 2025. Human-Centered AI in Arts. In *Handbook of Human-Centered Artificial Intelligence*. Springer, Singapore, 1–32. doi:10.1007/978-981-97-8440-0_86-1

[16] Nick Bryan-Kinns, Anna Wszeborowska, Olga Sutskova, Elizabeth Wilson, Phoenix Perry, Rebecca Fiebrink, Gabriel Vigliensoni, Rikard Lindell, Andrei Coronel, and Nuno N. Correia. 2025. Leveraging Small Datasets for Ethical and Responsible AI Music Making. In *Audio Mostly 2025* (2025). University of Coimbra, Portugal. https://audiomostly.com/

[17] Nick Bryan-Kinns, Shuoyang Zheng, Francisco Castro, Makayla Lewis, Jia-Rey Chang, Gabriel Vigliensoni, Terence Broad, Michael Paul Clemens, and Elizabeth Wilson. 2025. XAIxArts Manifesto: Explainable AI for the Arts. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 650, 8 pages. doi:10.1145/3706599.3716227

[18] Daniel Buschek, Lukas Mecke, Florian Lehmann, and Hai Dang. 2021. Nine Potential Pitfalls When Designing Human-AI Co-Creative Systems. In *HAI-GEN 2021: 4th Workshop on Human-AI Co-Creation* (2021). College Station, USA.

[19] Baptiste Caramiaux and Sarah Fdili Alaoui. 2022. "Explorers of Unknown Planets": Practices and Politics of Artificial Intelligence in Visual Arts. *Proceedings of the ACM on Human-Computer Interaction* 6 (2022), 1–24. Issue CSCW2. doi:10.1145/3555578

[20] Eva Cetinic and James She. 2022. Understanding and Creating Art with AI: Review and Outlook. *ACM Trans. Multimedia Comput. Commun. Appl.* 18, 2, Article 66 (2022). doi:10.1145/3475799

[21] Keshigeyan Chandrasegaran, Michael Poli, Daniel Y. Fu, Dongjun Kim, Lea M. Hadzic, Manling Li, Agrim Gupta, Stefano Massaroli, Azalia Mirhoseini, Juan Carlos Niebles, Stefano Ermon, and Fei-Fei Li. 2025. Exploring Diffusion Transformer Designs via Grafting. In *Advances in Neural Information Processing Systems*, Vol. 38. https://arxiv.org/abs/2506.05340

[22] Minsuk Chang, Stefania Druga, Alexander J. Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie J Cai, and Michael Terry. 2023. The Prompt Artists. In *Proceedings of the 15th Conference on Creativity and Cognition* (2023-06-19) *(C&C '23)*. Association for Computing Machinery, New York, NY, USA, 75–87. doi:10.1145/3591196.3593515

[23] Suk Kyoung Choi, Steve DiPaola, and Hannu Töyrylä. 2021. Artistic Style Meets Artificial Intelligence. *Journal of Perceptual Imaging* 4 (09 2021). doi:10.2352/J.Percept.Imaging.2021.4.2.020501

[24] Civitai. 2026. Civitai: The Home of Open-Source Generative AI. https://civitai.green/.

[25] ComfyUI. 2026. ComfyUI. https://www.comfy.org/.

[26] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion Models beat GANs on Image Synthesis. *Advances in neural information processing systems* 34 (2021), 8780–8794.

[27] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017-05-02) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 278–288. doi:10.1145/3025453.3025739

[28] Ahmed Elgammal. 2023. *Text-to-Image Generators Have Altered the Digital Art Landscape—But Killed Creativity. Here's Why an Era of A.I. Art Is Over.* Artnet News. https://news.artnet.com/art-world-archives/ahmed-elgammal-op-ed-ai-art-is-over-2304028 Accessed April 2026.

[29] Ziv Epstein, Aaron Hertzmann, the Investigators of Human Creativity, Memo Akten, Hany Farid, Jessica Fjeld, Morgan R. Frank, Matthew Groh, Laura Herman, Neil Leach, Robert Mahari, Alex "Sandy" Pentland, Olga Russakovsky, Hope Schroeder, and Amy Smith. 2023. Art and the Science of Generative AI. *Science (New York, N.Y.)* 380, 6650 (2023), 1110–1111. doi:10.1126/science.adh4451

[30] Rebecca Fiebrink and Perry R Cook. 2010. The Wekinator: a system for real-time, interactive machine learning in music. In *In Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)*.

[31] Corey Ford, Elizabeth Wilson, Shuoyang Zheng, Gabriel Vigliensoni, Jeba Rezwana, Lanxi Xiao, Michael Clemens, Makayla Lewis, Drew Hemment, Alan Chamberlain, et al. 2025. Proceedings of the third international workshop on eXplainable AI for the Arts (XAIxArts). *arXiv preprint arXiv:2511.10482* (2025).

[32] Monika Fratczak, Erinma Ochu, Jo Bates, Itzelle Medina-Perea, and Helen Kennedy. 2023. Cultures of AI Practice in the Arts. Patterns in Practice: cultures of data mining in science, education and the art.

[33] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. *Advances in Neural Information Processing Systems* 27 (2014).

[34] Google. 2026. Gemini Image Generation: Nano Banana Pro. https://gemini.google/overview/image-generation/.

[35] Imke Grabe, Miguel González Duque, Sebastian Risi, and Jichen Zhu. 2022. Towards a Framework for Human-AI Interaction Patterns in Co-Creative GAN Applications. In *Joint Proceedings of Intelligent User Interfaces (ACM IUI) Workshops* (2022). 92–102.

[36] Imke Grabe and Tom Jenkins. 2025. Hidden Layer Interaction: A Technique to Explore the Material of Generative AI. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference* (2025-07-04) *(DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1913–1927. doi:10.1145/3715336.3735437

[37] Anuj Gupta, Yasser Atef, Anna Mills, and Maha Bali. 2024. Assistant, parrot, or colonizing loudspeaker? ChatGPT metaphors for developing critical AI literacies. *Open Praxis* 16, 1 (2024), 37–53.

[38] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. 2020. Ganspace: Discovering interpretable gan controls. *Advances in neural information processing systems* 33 (2020), 9841–9850.

[39] Drew Hemment, Dave Murray-Rust, Vaishak Belle, Ruth Aylett, Matjaz Vidmar, and Frank Broz. 2024. Experiential AI: Between arts and explainable AI. *Leonardo* 57, 3 (2024), 298–306.

[40] Aaron Hertzmann. 2020. Visual indeterminacy in GAN art. In *ACM SIGGRAPH 2020 Art Gallery*. 424–428.

[41] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. https://openreview.net/forum?id=nZeVKeeFYf9

[42] Xi Hu, Yiwen Xing, Xudong Cai, Yihang Zhao, Michael Cook, Rita Borgo, and Timothy Neate. 2025. Designing Interactions with Generative AI for Art and Creativity: A Systematic Review and Taxonomy. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1126–1155. doi:10.1145/3715336.3735843

[43] Hugging Face. 2026. Hugging Face Models. https://huggingface.co/models.

[44] Minguk Kang, Joonghyuk Shin, and Jaesik Park. 2022. StudioGAN: A Taxonomy and Benchmark of GANs for Image Synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45 (2022), 15725–15742.

[45] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. 2023. Scaling up gans for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10124–10134.

[46] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8110–8119.

[47] Hyung-Kwon Ko, Gwanmo Park, Hyeon Jeon, Jaemin Jo, Juho Kim, and Jinwook Seo. 2023. Large-Scale Text-to-Image Generation Models for Visual Artists' Creative Works. In *Proceedings of the 28th International Conference on Intelligent User Interfaces* (2023) *(IUI '23)*. Association for Computing Machinery, New York, NY, USA, 919–933. doi:10.1145/3581641.3584078

[48] Akio Kodaira, Chenfeng Xu, Toshiki Hazama, Takanori Yoshimoto, Kohei Ohno, Shogo Mitsuhori, Soichi Sugano, Hanying Cho, Zhijian Liu, Masayoshi Tomizuka, and Kurt Keutzer. 2025. StreamDiffusion: A Pipeline-level Solution for Real-Time Interactive Generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 12371–12380.

[49] Stacey Kuznetsov and Eric Paulos. 2010. Rise of the expert amateur: DIY projects, communities, and cultures. In *Proceedings of the 6th Nordic conference on human-computer interaction: extending boundaries*. 295–304.

[50] Joseph Lindley and Roger Whitham. 2025. From Prompt Engineering to Prompt Craft. In *Proceedings of the Nineteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '25)*. Association for Computing Machinery, New York, NY, USA, Article 54, 12 pages. doi:10.1145/3689050.3704424

[51] Duri Long and Brian Magerko. 2020. What Is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020-04-23) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3313831.3376727

[52] Atefeh Mahdavi Goloujeh, Anne Sullivan, and Brian Magerko. 2024. Is It AI or Is It Me? Understanding Users' Prompt Journey with Text-to-Image Generative AI Tools. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (2024-05-11) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3613904.3642861

[53] Moshe Mahler, Yuchen Liu, and Jodi Forlizzi. 2025. *The Evolution of Artists' Toolsets: A Call for Artist-Driven Design.* Association for Computing Machinery, New York, NY, USA, 335–339. https://doi.org/10.1145/3715668.3736335

[54] Jon McCormack, Camilo Cruz Gambardella, Nina Rajcic, Stephen James Krol, Maria Teresa Llano, and Meng Yang. 2023. Is writing prompts really making art?. In *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*. Springer, 196–211.

[55] Jon McCormack, Camilo Cruz Gambardella, Nina Rajcic, Stephen James Krol, Maria Teresa Llano, and Meng Yang. 2023. Is Writing Prompts Really Making Art?. In *Artificial Intelligence in Music, Sound, Art and Design: 12th International Conference, EvoMUSART 2023, Held as Part of EvoStar 2023* (2023-04). Brno, Czech Republic. arXiv:2301.13049 [cs] doi:10.48550/arXiv.2301.13049

[56] Merzmensch. 2024. Good Bye, RunwayML: Not a fairy tale. Based on real events. Medium (Merzazine). https://medium.com/merzazine/good-bye-runwayml-35f5516aef1e

[57] Uwe Messer. 2024. Co-creating art with generative artificial intelligence: Implications for artworks and artists. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100056. doi:10.1016/j.chbah.2024.100056

[58] Metacreation Lab. 2026. Autolume: A Neural-network Based Visual Synthesizer. https://www.metacreation.net/autolume. Accessed April, 2026.

[59] Midjourney. 2026. Midjourney. https://www.midjourney.com/.

[60] OpenAI. 2026. ChatGPT Images: Create and edit images. https://chatgpt.com/features/images/.

[61] Jonas Oppenlaender, Rhema Linder, and Johanna Silvennoinen. 2025. Prompting AI Art: An Investigation into the Creative Skill of Prompt Engineering. *International Journal of Human–Computer Interaction* 41, 16 (2025), 10207–10229. doi:10.1080/10447318.2024.2431761

[62] Avik Pal and Aniket Das. 2019. TorchGAN: A Flexible Framework for GAN Training and Evaluation. *J. Open Source Softw.* 6 (2019), 2606. https://api.semanticscholar.org/CorpusID:202540048

[63] Justin N. M. Pinkney and Doron Adler. 2020. Resolution Dependent Gan Interpolation for Controllable Image Synthesis between Domains. In *Machine Learning for Creativity and Design Workshop (34th Conference on Neural Information Processing Systems (NeurIPS 2020).)*.

[64] Playform. 2024. Playform: No-Code AI for Creative People. https://www.playform.io/train.

[65] Anne Ploin, Rebecca Eynon, Isis Hjorth, and Michael A Osborne. 2022. AI and the Arts: How Machine Learning Is Changing Artistic Work. *Report from the Creative Algorithmic Intelligence Research Project* (2022).

[66] Rida Qadri, Piotr Mirowski, and Remi Denton. 2025. AI and Non-Western Art Worlds: Reimagining Critical AI Futures through Artistic Inquiry and Situated Dialogue. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (2025-04-25) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3706598.3714049

[67] Nina Rajcic, Maria Teresa Llano Rodriguez, and Jon McCormack. 2024. Towards a Diffractive Analysis of Prompt-Based Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (2024-05-11) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3613904.3641971

[68] Jeba Rezwana and Mary Lou Maher. 2022. Designing Creative AI Partners with COFI: A Framework for Modeling Interaction in Human-AI Co-Creative Systems. *ACM Transactions on Computer-Human Interaction* (2022).

[69] Chantal Rodier, Jason Millar, Willem Deisinger, and Sarah Jasmine Hodgson. 2023. Art Critically Examining Generative AI. In *2023 World Engineering Education Forum-Global Engineering Deans Council (WEEF-GEDC)*. IEEE, 1–9.

[70] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 10684–10695.

[71] Sonja Rozental, Michel Van Dartel, and Alwin De Rooij. 2025. How Artists Use AI as a Responsive Material for Art Creation. In *ISEA 2025: 30th International Symposium on Electronic/Emerging Art* (2025). Korea. doi:10.31234/osf.io/gjdnw

[72] Runway. 2026. Runway. https://runwayml.com/.

[73] Anush Sankaran, Raunak Sinha, and Naveen Panwar. 2018. *IBM GAN Toolkit.* https://https://github.com/IBM/gan-toolki

[74] Ben Shneiderman. 2007. Creativity Support Tools: Accelerating Discovery and Innovation. *Commun. ACM* 50, 12 (2007), 20–32. doi:10.1145/1323688.1323689

[75] Arshia Sobhan, Ahmed M Abuzuraiq, Philippe Pasquier, Lionel Ringenbach, and Kris Krüg. 2025. Autolume: A GAN-based No-Coding Small Data and Model Crafting Visual Synthesizer for Artistic Creation. In *Proceedings of the Conference on Animation and Interactive Art.* 165–172.

[76] Arshia Sobhan, Philippe Pasquier, and Gabriela Aceves Sepulveda. 2025. Broken Letters, Broken Narratives: A Case Study on Arabic Script in DALL-e 3. In *International Symposium on Electronic/Emerging Art (ISEA)*.

[77] Arshia Sobhan, Philippe Pasquier, and Adam Tindale. 2024. Unveiling New Artistic Dimensions in Calligraphic Arabic Script with Generative Adversarial Networks. *Proc. ACM Comput. Graph. Interact. Tech.* 7, 4, Article 61 (July 2024), 12 pages. doi:10.1145/3664210

[78] Mar Canet Sola and Varvara Guljajeva. 2025. Why Open Small AI Models Matter for Interactive Art. In *To appear in proceedings of ARTECH International Conference 2025*. ACM. arXiv preprint arXiv:2511.09788.

[79] StabilityAI. 2024. Stable Diffusion. https://stability.ai/stable-image.

[80] StabilityAI. 2026. SDXL-Turbo Model. https://huggingface.co/stabilityai/sdxl-turbo.

[81] Austin Tecks, Thomas Peschlow, and Gabriel Vigliensoni. 2024. Explainability Paths for Sustained Artistic Practice with AI. In *In Proceedings of Explainable AI for the Arts Workshop 2024 (XAIxArts 2024)* (2024-07-21). arXiv. arXiv:2407.15216 [cs] doi:10.48550/arXiv.2407.15216

[82] Hoang Thanh-Tung and Truyen Tran. 2020. Catastrophic forgetting and mode collapse in GANs. In *2020 International Joint Conference on Neural Networks (IJCNN)*. 1–10. doi:10.1109/IJCNN48605.2020.9207181

[83] Gabriel Vigliensoni, Louis McCallum, and Rebecca Fiebrink. 2020. Creating Latent Spaces for Modern Music Genre Rhythms Using Minimal Training Data. In *Proceedings of the 11th International Conference on Computational Creativity (ICCC'20)* (2020). Goldsmiths, University of London, Coimbra, Portugal.

[84] Gabriel Vigliensoni, Phoenix Perry, and Rebecca Fiebrink. 2022. A Small-Data Mindset for Generative AI Creative Work. In *In Generative AI in HCI Workshop, CHI '22*. New York, NY, USA, 5.

[85] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. 2021. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1905–1914.

[86] Justin D. Weisz, Michael Muller, Jessica He, and Stephanie Houde. 2023. Toward General Design Principles for Generative AI Applications. In *4th Workshop on Human-AI Co-Creation with Generative Models (HAI-GEN) at IUI 2023* (2023-01-13). arXiv:2301.05578 [cs]

[87] 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* (2023). arXiv:2308.06721

[88] Sean Zhai and Rebecca Ruige Xu. 2025. Pareidolia in the Machine: Unintended Figuration in AI-Generated Video from Abstract Visual Music. In *Proceedings of the 2025 Conference on Creativity and Cognition* (2025-06-22) *(C&C '25)*. Association for Computing Machinery, New York, NY, USA, 964–974. doi:10.1145/3698061.3726957

[89] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023). 3836–3847.