

**Document Version**

Final published version

**Licence**

CC BY

**Citation (APA)**

van der Burg, V., de Boer, G., Benjamin, J. J., Halperin, B. A., Akdag Salah, A. A., Chandrasegaran, R. S. K., & Lloyd, P. A. (2026). Reflective AI: A Slow Technology Approach for Design Education . In *Proceedings of CHI '26, Barcelona, Spain* Association for Computing Machinery (ACM). <https://doi.org/10.1145/3772318.3791691>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.  
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

# Reflective AI: A Slow Technology Approach for Design Education

Vera van der Burg

Faculty of Industrial Design  
Engineering  
TU Delft  
Delft, Netherlands  
v.vandenburg@tudelft.nl

Gijs de Boer

MA Critical Inquiry Lab  
Design Academy Eindhoven  
Eindhoven, Netherlands  
gijs.deboer@designacademy.nl

Jesse Josua Benjamin

Industrial Design  
Eindhoven University of Technology  
Eindhoven, Netherlands  
j.j.benjamin@tue.nl

Brett A. Halperin

Human Centered Design &  
Engineering  
University of Washington  
Seattle, Washington, USA  
bhalp@uw.edu

Alkim Almila Akdag

Faculty of Science, Information and  
computing sciences  
Utrecht University  
Utrecht, Netherlands  
a.a.akdag@uu.nl

Senthil Chandrasegaran

Faculty of Industrial Design  
Engineering  
TU Delft  
Delft, Netherlands  
r.s.k.chandrasegaran@tudelft.nl

Peter Lloyd

Faculty of Industrial Design  
Engineering  
TU Delft  
Delft, Netherlands  
p.a.lloyd@tudelft.nl



**Figure 1: Scenes from our workshop study, showcasing participants engaging in data collection, annotation, Object Portrait creation, and reflective exploration of self-trained AI interpretations. These activities informed the development of implications for defining characteristics of a Reflective AI design practice.**

## Abstract

The proliferation of efficiency-focused AI tools in creative processes threatens to undermine critical, reflective practices foundational to design education. This approach can lead to creativity exhaustion and diminished agency among designers and students. As an antidote, we propose Reflective AI: an approach grounded in slow technology principles that reframes AI not as a production tool, but

as a medium for reflecting on the creative process itself. This paper presents the Objective Portrait Workshop where design students engaged in slowed data collection, annotation, and model finetuning. Our contribution is threefold: we (1) document a methodology for implementing Reflective AI in design education; (2) provide empirical evidence that slow engagement cultivates reflection on creative processes and technical understanding of AI; and (3) propose material and temporal disentanglement as core mechanisms for Reflective AI practice. This work offers a practical alternative to “fast” AI, providing methodology that cultivates critical capabilities essential to design.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2278-3/26/04  
<https://doi.org/10.1145/3772318.3791691>

## CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → *Artificial intelligence*.

## Keywords

Slow Technology, Artificial Intelligence, Design Practice, Reflection

### ACM Reference Format:

Vera van der Burg, Gijs de Boer, Jesse Josua Benjamin, Brett A. Halperin, Alkim Almila Akdag, Senthil Chandrasegaran, and Peter Lloyd. 2026. Reflective AI: A Slow Technology Approach for Design Education. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3772318.3791691>

## 1 Introduction

The proliferation of AI-empowered tools in creative processes promise to enhance creativity and accelerate design processes through rapid content generation and automated workflows. Yet, this efficiency-focused paradigm introduces fundamental tensions for design practice. Studies document how AI tools can lead to design fixation [22], creativity exhaustion [29], homogenization of style [27, 42], and diminished creative agency [13, 60]. For design education, these challenges are particularly acute: students need to develop critical, reflective practices fundamental to design [46, 48]. Such design education needs are further compromised by current AI tools being presented as polished, chat-based interfaces optimized for speed—these actively disrupt the iterative, reflective pedagogical goals essential to design education [15, 17]. While researchers have noted the urgent need for engaging with AI's impact on education [6, 7, 30], most focus on K-12 contexts rather than higher education and design education specifically.

We argue that the core problem lies not in AI technology itself, but in its prevailing presentation as polished, prepackaged solutions that prioritize a “move fast and break things” mentality [20]. This emphasis on efficiency sidelines the nuanced, iterative practices foundational to design [46, 48, 49] and inhibits the exploration of AI as a design material itself [3, 52]. When AI tools prioritize rapid output generation, as we observe with current generative AI tools like ChatGPT [38] and Midjourney [33], they obscure the underlying processes, decisions, and material conditions necessary for designers to critically understand and engage with the technology. This creates a pedagogical paradox: we risk educating designers who can use AI tools but cannot critically examine their implications—a concerning prospect given technology's increasingly central role in creative design practice.

To address this, we turn to slow technology principles [19], which create “intentional slowness” by prioritizing reflection and meaningful engagement over efficiency and high performance, transforming technologies from productivity tools into reflective media. This approach creates time and space for understanding how technology works, why it works that way, and its consequences. While slow technology has traditionally focused on the thoughtful (re)design of discrete, tangible artifacts [35–37], AI systems exist as complex assemblages of interconnected components, temporal processes, and

probabilistic mechanisms. This assemblage nature demands new approaches to engage with AI's distributed and temporal complexity while maintaining slow technology's emphasis on reflection.

Building on these principles, we introduce *Reflective AI*<sup>1</sup> specifically for design education contexts: a slow technology approach to AI in creative practice that treats AI systems as mediums for self-reflection rather than generators of design outputs. Where conventional AI tools hide complexity behind streamlined interfaces, Reflective AI deliberately exposes and engages with AI's constituent components—data collection, annotation, training, inference—as sites for critical reflection and as interpretations of slow technology. In an education context, this enables students to co-construct understanding through direct material engagement (cf. [39]), discovering how their subjective judgments shape and are shaped by technical systems.

In this paper, we present the Objective Portrait Workshop as a concrete implementation of Reflective AI in design education. Designed to cultivate reflective positioning and critical technical awareness, this structured three-day workshop involved 11 design students who created custom object detection models trained on personally meaningful datasets, then used these models to analyze their own creative work. Building on van der Burg and colleagues' Objective Portrait method [58], we adapt this approach specifically for educational contexts. While the foundational work explored AI as a reflective partner for professional designers, we repurpose the method to cultivate critical technical literacy in students, emphasizing prolonged engagement with each phase of the AI pipeline. Our contribution is threefold: (1) we document a concrete methodology for implementing Reflective AI in design education; (2) we provide empirical evidence of how slow engagement with AI components cultivates both reflection on creative processes and technical understanding of the technology; and (3) we propose a conceptual framework based on slow technology principles, identifying material and temporal disentanglement as the core mechanism for a Reflective AI practice. Through this work, we offer design educators and practitioners a practical alternative to fast, generative AI—one that cultivates the critical, reflective capabilities essential to design practice.

## 2 Related Work

This section examines research motivating our Reflective AI approach in design education. We establish how “fast” AI tools can undermine reflection essential to design learning, explore slow technology as an alternative framework, and identify the gap between conceptual approaches and concrete educational methods that our Objective Portrait Workshop addresses.

### 2.1 The Need for Reflective AI Methodologies in Design Education

Reflection on creative process is arguably the defining feature of design practice and research [16, 17, 46, 48]. Design education has traditionally cultivated this reflective capacity through iterative, slow processes encouraging students to examine their thinking and creative decisions over time. This pedagogical approach assumes

<sup>1</sup>Note that we are not proposing reflection *on part of* the AI technology (as in e.g. Lewis and Sarkadi [28]).

that designers learn not just by producing outputs, but by developing critical awareness of how and why they make creative choices—what Schön termed “reflection-in-action” that distinguishes expert practice from mere technical execution [46].

There is considerable enthusiasm for incorporating AI in design practice and design education, with AI tools promising enhanced creativity and faster design iterations. We identify this prevailing efficiency-driven approach as “fast AI”. In this mode, students use tools such as Midjourney or ChatGPT primarily to accelerate ideation, for instance by generating instant moodboards or visual variations [54], or to automate production tasks like copywriting and UI prototyping [29, 60]. However, incorporating “fast” AI in design practice introduces significant risks that could undermine this essential reflection. Research has demonstrated that contemporary AI tools often lead to “design fixation” [22, 59], where the speed and perceived finish of AI outputs cut short the exploration of design spaces. The technical opacity of AI systems—arising from their layered structure and complex probabilistic mechanisms, further obscured by proprietary text- or chat-based interfaces—limits designers’ sense of control and agency over the creative process [21, 61]. Moreover, Li and colleagues [29] identified the phenomenon of “creativity exhaustion”, where designers feel pressured to match the pace of generative AI outputs. While generative AI can enhance individual creativity to an extent, its collective creative outputs often exhibit limited novelty [12], and these tools do not inherently inspire more novel or innovative ideas [54].

Then, in design education specifically, these challenges are particularly acute. Sandhaus and colleagues’ study of the unprompted and uncoordinated use of generative AI by all student groups in an HCI design course found that despite some benefits in “enhanced creativity and faster design iterations”, the lack of structured engagement led to “shallow learning and reflection” rather than deep understanding [45]. This is particularly concerning for HCI design education, where students must develop not only technical skills but also critical sensitivities toward the social, ethical, and experiential dimensions of technology. While research on AI technologies within educational contexts has primarily focused on K-12 education (e.g. [6, 7, 30]), HCI design education itself has received comparatively little attention despite its unique requirements for cultivating both technical understanding and critical reflection. The general technical opacity, and even more so the “opacity at runtime” [9] of automatic generation, hinder the in-depth, reflective integration that design education seeks to cultivate—particularly problematic when design students need to develop nuanced understanding of how their creative choices interact with technological capabilities and constraints.

Some existing work has begun addressing these challenges in design education. Murray-Rust and colleagues’ “Grasping AI” work [34] developed nine experiential exercises for design students to engage with AI as a socio-technical system, using enactment, metaphor work, and role-playing to help students develop critical perspectives. This aligns with broader efforts emphasizing tangible [7, 30] and constructionist [6] educational approaches to AI and machine learning technologies. However, existing approaches primarily address the technical opacity and social implications of AI systems, rather than the presentation of AI as ready-made, productivity-oriented “fast” tools. While these efforts successfully create reflective spaces

around AI’s socio-technical dimensions, they do not specifically address how the efficiency-focused paradigm of current AI tools undermines the slow, reflective practices foundational to design education.

We argue that this constitutes an opportunity to rethink engagement with these technologies through slower, more reflective interactions. In the following, we therefore outline how principles from “slow technology” [19] could be utilized, before presenting our methodological considerations for a study in which we use design education as an exploratory and exemplary context for a potential ‘Reflective AI design’ approach.

## 2.2 Slow Technology and Reflective AI as Alternative Approaches

Slow technology emerges from a fundamental critique of efficiency-driven design paradigms that prioritize speed and productivity over human understanding and reflection. Hallnäs and Redström [19] established slow technology as “a design agenda for technology aimed at reflection and moments of mental rest rather than efficiency in performance”. Rather than simply introducing temporal delays, slowness in this context offers users time to: 1) learn how the technology functions, 2) understand its design rationale, 3) apply it effectively, 4) observe its outcomes, and 5) critically evaluate its consequences [19]. This type of slowness concerns not time perception but rather “time presence”: when we use technology as an efficient tool, time disappears as we accomplish tasks, but accepting an invitation for reflection means that “time will appear”, creating space for deliberate engagement.

However, applying slow technology principles to AI in design education contexts requires adaptation. Traditional slow technology work has focused on creating discrete artifacts that embody slowness for others to experience [36, 37, 50]. To illustrate these principles in practice, consider Odom and colleagues’ Olly [37]; a domestic music player that deliberately constrains access by only playing music from the same day in previous years, transforming music consumption from efficient task completion into contemplative encounter with one’s musical history. AI technologies, by contrast, exist as composite assemblages of interconnected components—data collection, annotation, training, inference—rather than explicitly tangible, single interactive objects. Moreover, the educational goal is not to create slow artifacts for others to use, but to develop a reflective understanding of their own creative process—through a deliberate, slow, intentional engagement of AI technology.

Some work has explored slow technology principles in AI contexts, though with limited scope. Park et al. [40] showed that intentionally slowing algorithmic response times prompted users to reflect on processes behind algorithmic decisions in high-stakes contexts like crime prediction. Similarly, Cremaschi et al. [11] introduced temporal constraints in generative AI through a modified typewriter for “slow tweets,” critiquing accelerated content creation. However, these approaches focus primarily on temporal delays—artificially slowing down interactions—rather than creating comprehensive reflective spaces that engage with AI’s material and procedural complexity. They represent one interpretation of slowness but do not address the deeper challenge of transforming AI from opaque tool to reflective medium.

The proposal for Reflective AI presents an alternative way to incorporate AI in the creative process, and is a specific application of slow technology principles to AI in creative practice. Instead of using AI as a tool to produce or design *with*, the key is to reframe AI technologies as reflective media that enable critical examination of one’s own creative practice. The reflective use of AI has been explored in previous work, notably van der Burg and colleagues’ “Objective Portrait” method [58], which demonstrates how AI can serve as a mirror for design thinking rather than a generator of design outputs. Similarly, Arzberger and colleagues’ approach to “reflexive data curation” [2] shows how deliberately engaging with AI’s biases and uncertainties can become a site for self-confrontation and critical reflection on one’s own worldview. These approaches treat AI technologies as sociotechnical assemblages that can be disentangled and engaged with reflectively. However, while implicitly embedding slow technology principles within the creation and interaction process, they have not yet been explicitly established within the larger body of slow technology literature. We consider these early, implicit examples as foundational for what we term Reflective AI, aiming to formally establish and apply these slow technology principles within design education.

Building on these foundational examples, this paper formalizes the Reflective AI approach as a direct antidote to the logic of opaque efficiency prevalent in creative AI tools. We translate this approach into a design education practice through the “Objective Portrait” Workshop, a concrete educational methodology designed to use AI to cultivate, rather than diminish, creative self-reflection and creative agency within design education.

### 3 Methodology

This section presents the Objective Portrait Workshop as a structured educational method for integrating a Reflective AI position into design education. Generally, we position our research methodology as Research-through-Design (RtD) in that we design and deploy an educational intervention (cf. [6]) with the goal of sourcing “intermediate-level knowledge” [23] on the potential of Reflective AI in design education for further research. We adapted van der Burg et al.’s Objective Portrait method [58] into a pedagogical framework grounded in slow technology principles, transforming individual design reflection into a transferable educational methodology. To do this, we first establish the theoretical rationale for our approach, then describe four transferable methodological components, followed by details of the three-day workshop process, our participant group, and finally our data collection and analytical procedures.

#### 3.1 Rationale and Principles of the Objective Portrait Workshop

The Objective Portrait Workshop builds on van der Burg et al.’s Objective Portrait method [58] as its foundational approach because this method inherently embodies the reflective engagement we seek to cultivate in design education. Unlike engaging with AI as an opaque tool merely for rapid output generation, the Objective Portrait method requires users to engage with the entire AI pipeline—from data collection and annotation to model training and output interpretation. Whereas the previous work was done with three professional designers, we translate the method into a



**Figure 2: An example of an Object Portrait from van der Burg et al. [58]. On the left, the original Object Portrait is shown, while on the right, the same portrait is analyzed by a fine-tuned YOLOv5 model trained on the creator’s personal fascinations.**

structured engagement over three days, specifically aimed at design students. Our educational implementation diverges from the original work. The prior study was a first-person, practice-based inquiry where professional designers explored AI as a reflective partner. In contrast, our workshop translates this open-ended exploration into a structured pedagogical exercise. The conceptual evolution lies in repurposing the method from a research probe for professionals into a guided learning instrument for students. We shift the focus from an open-ended inquiry into designer-AI relations toward a guided workshop designed to cultivate critical technical literacy and reflective practice.

The Objective Portrait Workshop aligns with slow technology principles [19] in several key ways. First, it deliberately prioritizes reflection over efficiency by requiring extended engagement with AI components rather than optimizing for quick results. Participants spend considerable time curating personal datasets, developing annotation strategies, and interpreting model outputs—activities that would be considered inefficient in efficiency-driven practices but are essential for educational growth in designerly reflection. Second, the method transforms AI from a productivity tool into a medium for self-examination, allowing participants to observe their own design process reflected back through the model’s interpretations of their work. This comprehensive engagement creates multiple opportunities for critical reflection as participants must make explicit their subjective perspectives and observe how these perspectives shape technological outcomes.

The workshop aims to achieve two specific learning objectives. First, we support reflective positioning, encouraging students to deepen their relationship with their research topic through the

external lens of data curation. Second, we focus on critical technical awareness. By engaging hands-on with the mechanics of training, we motivated the students to move beyond passive use and develop a sense of agency, recognizing model behaviors as reflections of their own curatorial decisions rather than objective technical realities.


### 3.2 Methodological Components

By deploying the Objective Portrait as a workshop for design students, we sought to explore whether a Reflective AI approach could offer an alternative to current AI integration in design education. We hypothesized that shifting from “fast AI” integration, where students might simply prompt a generative model for quick outputs, to “slow AI” engagement could reveal the subjective choices embedded within seemingly objective technological systems, similar to how van der Burg et al. stated: “Reflection is in the annotation rather than only in using the AI” [58]. Building on this insight, we further hypothesized that requiring students to train their own models with personally meaningful data would make visible how individual perspectives and biases shape AI behavior, potentially creating concrete anchor points for reflection throughout the AI training process. We anticipated that extending engagement across data collection, annotation, training, and interpretation might ensure that reflection becomes integral to the process rather than an afterthought. Based on these hypotheses, we developed four transferable components that could potentially be adapted across different educational contexts.

**3.2.1 Object Detection as Reflective AI Tool.** Our technical decisions prioritized pedagogical opportunities for reflection and critical engagement over computational efficiency or state-of-the-art capabilities. Therefore, our choice of an object detection model over a generative model was intentional, informed by previous work from [57, 58] and [31], which recognize object detection results as a foundation for interpretation, speculation, and reflection.

Object detection technology is similar to that used in self-driving cars, facial recognition software, or traffic monitoring systems, and is typically trained on large datasets of images annotated with objective or descriptive labels such as street signs, zebra crossings, and traffic lights. However, in our workshop, students customized models to analyze images *they* produced following the adapted Objective Portrait approach (see section 3). Due to this setup, dataset curation (i.e. collecting and annotating data) becomes a reflective process involving human annotation of specific sections of an image (i.e. creating multiple annotations with bounding boxes within one image) rather than providing descriptions of the entire image, as seen in generative diffusion models. This approach offered greater potential for creating a reflective space, as the annotation process for object detection models is inherently slow and deliberate, i.e. pointing out multiple labels within one image, in contrast to generating a single sentence per image.

We used YOLOv5 [55], an open-source object detection model pre-trained on a large dataset. It is easily re-trained with minimal computational resources. Unlike generative AI technologies, which for instance create new images from textual prompts, YOLOv5 interprets existing images using a suggestive visual language, a bounding box with a confidence percentage, that invites reflection by offering



**Figure 3: Quadrant tool for determining labels with which to annotate their data, shown here as filled by Pepa, a participant interested in the topic of age limitations on digital media. The tool, based on van der Burg et al. [58], helped students think of labels as occupying opposing ends of two orthogonal axes, which helped them come up with minimally-overlapping labels. Using this tool, they were able to use more precise language for the feelings associated with the “hot/cold” and “push/pull” dimensions of the quadrant in order to crystallise into words their subjective judgement on aspects of their topics of interest.**

an interpretation of something already present. In terms of both workshop design as well as the actual practices of slow engagement YOLOv5 is ready-to-hand due to being (a) openly available for customization and (b) due to its comparatively small computational requirements processing and training can occur more accessible for hands-on use, making the workshop method feasible. More importantly, the limitations of YOLOv5 in terms of detection and classification are welcome since they require students to actively interpret the model’s outputs rather than simply accepting them. We trained each of the students’ models on an external server to address the challenge of students having to find significant computing power on their own.

**3.2.2 The Quadrant Tool.** The Quadrant Tool (explicitly taken from [58]) creates the reflective foundation for AI training within Slow Technology principles. It serves as a vehicle to flesh out suitable labels for annotation, preceding the data annotation task. The tool features intentionally undefined ‘push/pull’ (x-axis) and ‘warm/cold’ (y-axis) axes. We asked students to define these dimensions in the context of their topic and physically place potential emotional keywords onto the map. This required them to articulate why a specific feeling was “warmer” or “more pushing” than another. This spatial

**Table 1: Overview of workshop students and their chosen topics.**

Name	Topic/Question	Image Dataset	Labels
Eva	Humor in memes	Memes	ridiculous, amusing, appalling, facetious
Aisha	How would AI react to humanistic beliefs, cultural oppression, and social etiquettes?	Religious events	purity, salient, dominant, synchronous
Carlo	Queer identity expression	People in extravagant outfits or drag fashion	enticing, flamboyant, post-human, cloying
Sascha	Vibes at a festival	Festival stages & concert stages	commercial saturation, convivial, lifeless, otherworldly
Yann	What can the refusal of advancement actually look like?	Abandoned buildings, artworks of Gordon Matta-Clark and Alvin Baltra	devotion, inconsistent, pretentious, rebelling
Borisz	Rubens' paintings and compositions	Rubens' paintings	majestic, repulsive, sublime, terrifying
Pepa	Age limitations on digital mediums	Children's cartoons from age 4–6	comforting, disruptive, frightening, overwhelming
Jassir	"Morphed by force"	Faces morphed by force	forged, mythical, reject, thrill
Rey	AI comedian	Wildlife photography	compromise, drastic, gloomy, mysterious
Yuky	Judgment in fashion design	Models on runways wearing biomimicry fashion	charismatic, generic, premature, tasteless
Agustina	Birthdays	Stock photography of birthday celebrations	isolated, neglect, safe, vain

mapping forced a commitment to a vocabulary that would later serve as the classification labels for the object detection model.

The intentional ambiguity regarding word organization facilitated the articulation of opposing feelings or judgments. This structured the label selection task and minimized meaning overlap. Students associated descriptive words with these feelings and placed them on the axes, organizing emotions into distinct sets of opposites (e.g., differentiating 'happy' from 'content').

Students drew and populated the quadrant tool by hand, making it an accessible resource for frequent annotations, allowing them to strike through prior choices and weight labels on the axes. This manual labeling process aimed to open a wider space for reflection and shifting interpretation. An example of a filled in quadrant tool appears in Figure 3. In contrast to efficiency-driven annotation pipelines aimed at producing objectively correct labels, our goal was to explore how the comparatively slow setup and integration of an AI model into a design process could become reflectively meaningful. We adapted the tool to slow down annotation intentionally. This encouraged students to pay closer attention to the task requirements and to customize it to reflect their design practice interests.

**3.2.3 The Object Portrait.** The Object Portrait serves as both the workshop's culminating activity and its conceptual anchor, representing a fundamental shift in how designers might engage with AI systems. An Object Portrait is a carefully composed image that students create to be analyzed by their trained model. The process of creating an Object Portrait involves five main steps: (1) identify a particular topic or fascination (e.g., within one's creative practice); prepare for training an AI model by (2) collecting relevant data and (3) annotating it with one's own subjective labels; (4) create a portrait (i.e., scene, still life, etc.) that subjectively reflects one's topic as defined in (1); and (5) have the portrait interpreted by the

AI model trained on the topic [58]. An example of an original Object Portrait is shown in Figure 2.

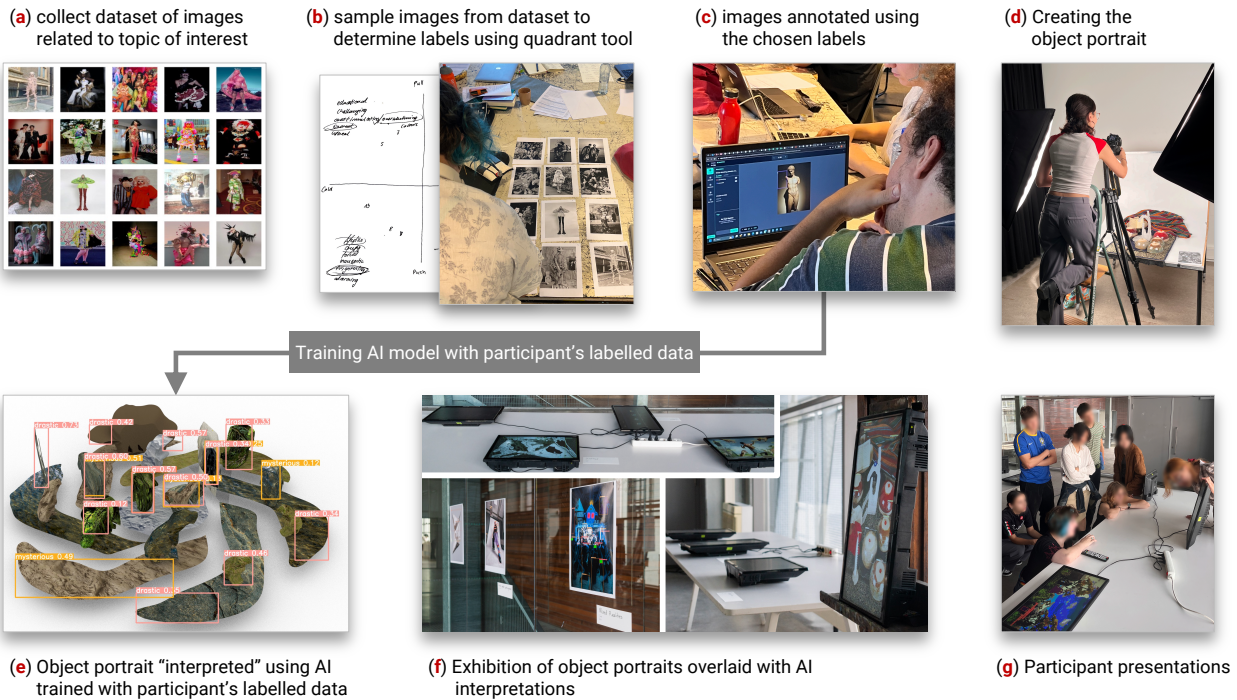
Students were asked to frame and articulate a central topic or research question which motivates their design practice leading to the Object Portrait. This formed the basis for data collection, labeling and training their object detection model before creating the Object Portrait which would represent their topic and deploying the model to analyze it. Through this approach, we anticipated that we would gain insights on (i) how qualitatively and subjectively rich categories such as design practice characteristics are made meaningful in the labeling and training processes and (ii) what kind of effect the customization of an AI model would have on the overall design process. *Note on terminology: We refer to the broader methodology as the Objective Portrait (following [58]), while the specific visual artifact created by the student is referred to as the Object Portrait*

## 4 The Object Portrait Workshop

This section details our implementation of the Objective Portrait Workshop, structured across three days (see Table 2 and Figure 4 for visual overviews), and outlines our student participants, data collection, and analytic approach. We adapted the Objective Portrait method into a workshop format where each day emphasized a distinct aspect of Reflective AI engagement. The section proceeds in four parts: (1) a description of student participants, (2) a day-by-day account of the workshop activities, (3) our approach to data collection, and (4) the analytic procedure. For a detailed breakdown of the specific prompts and instructions provided to students during the workshop, see the Appendix A.

### 4.1 Student Participants

We recruited 11 students from a bachelor's program in digital systems and new media at a recognized design school in Europe who



**Figure 4: Steps in the workshop over the three days. Steps (a) through (c) happen on day 1, step (d) on day 2, and steps (d) through (f) on day 3. Images shown in each step are not necessarily of any one participant’s dataset, labels, or object portrait, but the steps represent those of each participant.**

could benefit from developing critical AI literacy within their design education. This program’s focus on technology-centered design education aligned well with our workshop’s exploration of Reflective AI. All students reported common interest in technology and came from diverse cultural and professional backgrounds, though we chose not to collect detailed demographic data due to its irrelevance for our methodological approach. The workshop timing strategically coincided with the early phase of students’ academic year, when students were tasked with exploring topics representing their individual research interests. This alignment meant students could integrate their workshop engagement with ongoing coursework, selecting topics and corresponding datasets that reflected their personal areas of exploration (see Table 1). Given that students retain authorship of their Object Portraits, first names have not been anonymized at their request. We received informed consent from all students, and this research was approved by the [removed for review] Ethics Committee.

### 4.2 Day 1: Dataset Creation and Annotation

To foreground students’ individual practices, we asked them to complete topic selection and initial data collection prior to the workshop. This ensured personal investment in their chosen themes and allowed for organic meaning-making during the three-day process. To personalize the model, the images had to be thematically aligned with each participant’s research topic or question (see Table 1).

To ensure the effective functioning of the object detection model, the dataset needed to contain at least 100 images with consistent styles and sizes (e.g., no mixing of photos with paintings or screenshots). This emphasized the “small dataset mindset” described in [1], making annotation manageable while giving students greater control over model training data. The day began with a group reflection where students discussed their chosen datasets and research interests. During the first part of the workshop, students created four subjective labels (Task 1.2) using the Quadrant Tool (see 3.2.2), which helped them articulate emotional and judgmental categories through close examination of 10 images. They refined these labels iteratively, identifying overlaps or gaps in their initial choices (see example in Figure 3). In the afternoon, students annotated their datasets using Roboflow<sup>2</sup> (Task 1.3). This two-hour activity emphasized slow, deliberate annotation, with students encouraged to revisit and adjust labels as needed. Afterward, they reflected individually by annotating their quadrant tools with notes on how the labeling process influenced their perception of the labels and their research topics (Task 1.4 in Table 2). The steps from Day 1 are illustrated in Figure 4(a)–(c).

### 4.3 Day 2: Object Portrait Making

We introduced the concept of the Object Portrait [58] (see also Section 3.2.3) as an artwork explicitly created for AI interpretation.

<sup>2</sup>Roboflow Annotate: <https://roboflow.com/annotate>

**Table 2: Overview of the tasks, estimated duration, and underlying purpose over the 3 days of the workshop.**

Day	Agenda	Task	Duration	Purpose
1	Training the Object Detection Model Reflectively	1.1 Collect image dataset	Prior	Obtain thematically-grounded training data for the AI model
		1.2 Label Selection	1h	Crystallise emotions and subjective judgements about the dataset
		1.3 Dataset annotation	3h	Assign above labels to training data to inform AI model “interpretation”
		1.4 Individual reflection	30m	Reflect on (change in) perception of chosen labels
2	Creating the Object Portrait	2.1 Object Selection	2h	Reflect on own practice and material
		2.2 Object Portrait creation	3h	Reflect on own practice in relation to the dataset and labels
		2.3 Collective reflection	30m	Gather insights on the process of creating an image for AI interpretation
3	Integrating the AI predictions	3.1 Editing Object Portrait based on AI output	1h	Reflect on expected and unexpected aspects of AI model interpretation
		3.2 Exhibiting the work	1h	Adapting to changes in perspective based on AI model output
		3.3 Collective Interpretation	1h	Shared reflection on AI “interpretation” of their work
		3.4 Presenting the work	1h	Articulating the essence of the work and what they learned
		3.5 Individual reflection	30m	Reflecting on change in perspective on AI

Students followed guidelines requiring that portraits (1) feature objects related to their practice (Task 2.1), (2) mimic visual or aesthetic elements of their dataset to ensure model recognition (Task 2.2), and (3) take the form of either a single image or a 10-second video. Importantly, students were asked to produce these portraits during the workshop, encouraging a making-oriented process that deliberately slowed the transition to model testing. Once completed, portraits were shared in a group reflection session, where students discussed the experience of “creating an image for the AI model.”

#### 4.4 Day 3: AI Analysis and Reflection

Prior to Day 3, we trained object detection models on each participant’s annotated dataset. This process took approximately eight hours of computational processing time to complete the training for the full cohort of detection models. On the final day, students tested their Object Portraits with these models, interpreting the results as the AI’s “reading” of their work. They were given the option to modify their portraits based on these insights (Task 3.1 in Table 2). The day concluded with a collective exhibition of portraits (Task 3.2), followed by group discussion (Task 3.3) where students reflected on what the AI interpretations revealed, especially overlooked or surprising elements. After presenting to tutors (Task 3.4), students completed structured reflection cards (Task 3.5), documenting their learnings, shifts in perspective, and evolving relationships with AI. Steps from Day 3 are illustrated in Figure 4(e)–(g). See Figure 8 for a collection of the final Object Portraits.

#### 4.5 Data Collection

In designing the workshop, we anticipated large amounts of qualitative data and prepared for the emergent findings typical of RtD methodologies (cf. [16]). Since we could not predict exactly when or

how students would reflect, we identified key moments to explicitly prompt written or group reflections (see Appendix A for when and how these occurred). These moments were captured using a voice recorder. We also maintained tutor notes throughout the three days to document our guidance and observations.

Data collected included students’ Object Portraits, image datasets, annotated Quadrant Tools, audio-recorded reflections, written reflection cards, and tutor notes. Two authors began by segmenting the collected data according to each stage of the workshop using a collaborative online whiteboard. They also contributed their own reflections as tutors to contextualize the students’ processes. Transcribed reflections were mapped onto the corresponding workshop steps and linked with students’ images and labels.

#### 4.6 Analytic Procedure

We employed thematic analysis [41] to identify patterns across the collected data. First and second author, who also served as workshop tutors, collaboratively mapped transcribed and written reflections onto a visual representation of the workshop stages, linking them to the corresponding steps, labels, and images. To ensure rigor, we adopted a collaborative approach inspired by reflexive thematic analysis [8]. This involved third author, with a background in philosophy of technology, who reviewed the initial thematic mapping to challenge our interpretations from a theoretical standpoint. Rather than applying a rigid coding scheme or calculating inter-rater reliability scores, we prioritized developing a rich interpretation of the data through investigator triangulation. We resolved interpretive discrepancies through discussions until we reached consensus. The final thematic structure resulting from this analysis is detailed in the Appendix A.

Our analytic approach addressed two areas of inquiry: (1) how students reflected on their own creative practice and design thinking through Reflective AI, and (2) how engagement with AI systems facilitated insights, particularly across extended temporal processes of data collection, annotation, training, and interpretation. We paid close attention to moments where students’ understanding shifted, whether in their research focus, annotation strategies, or relationship to AI technology. This analysis allowed us to identify both immediate educational outcomes and broader implications for implementing Reflective AI in design education.

## 5 Findings

Our analysis shows that the slow, hands-on process of dataset curation, manual annotation, model training, and Object Portrait creation facilitated two distinct yet interconnected forms of learning. First, students engaged in learning *through* the AI: reflection on their own creative practices, with AI serving as a reflective medium—the workshop’s primary pedagogical goal. Second, students experienced learning *about* the AI: developing a deeper, practice-based understanding of AI systems and their subjectivities. This section presents these outcomes in turn. To ground these findings in the students’ lived experiences, our analysis draws extensively on direct quotations from their reflections.

### 5.1 Learning Through AI as a Reflective Medium

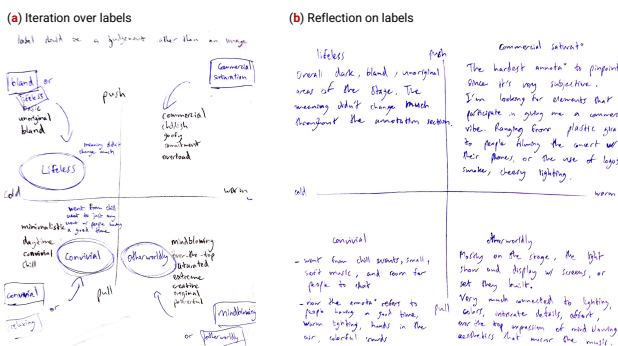
Learning *through* the AI captures insights where the AI workshop structure became a vehicle for self-reflection on one’s own creative process, enabling students to discover new understandings, aesthetic judgments, and ways of seeing. Here, the AI served as a reflective medium rather than merely a technical tool, with learning outcomes extending beyond AI literacy itself.

that represented potentially conflicting emotions and perspectives. This task revealed a significant shift in how labels are typically conceptualized, transitioning from static, object-based categories to temporal and fluid ones in a dynamic process.

Carlo’s experience exemplifies this process. Working with a dataset focused on queer identity expression and featuring photographs of people in queer and drag fashion, Carlo initially felt overwhelmed by their “*complex and overwhelming admiration*” for the individuals in their dataset. Through multiple iterations with the quadrant tool, Carlo arrived at two contrasting labels: “*enticing*” and “*cloying*.” The breakthrough came when Carlo began “*personifying*” the labels: “*I’ve sort of thought like if these [labels] were people, how would I describe them? For some reason that really helped me. So for example for [the label] ‘enticing,’ I say that it’s someone who’s intriguing and calm. And then, cloying was when it’s like a sweet, you know, like too much. It’s so much that it overpowers your senses and you’re just drowning in the amazingness.*” This process helped Carlo recognize that while their admiration was consistently positive, it contained an important distinction between measured appreciation and overwhelming intensity.

Sascha’s exploration of festival “*vibes*” demonstrates how the structured approach revealed nuanced emotional landscapes. Working with festival stage imagery, Sascha used the label selection process to map different feelings he encountered at festivals, ranging from “*otherworldly*” magical atmospheres to moments of “*commercial saturation*” that lacked depth. As he explained: “*So ‘lifeless’ and ‘commercial saturation’ are two kinds of themes in the festival stages that I thought didn’t really draw me forward and were not really appealing to me. And then ‘otherworldly’ and ‘convivial’ are two words that I thought were warmer and kind of intriguing, kind of as a reflection of the music and the artist.*” By organizing these feelings on contrasting axes, Sascha created a structured understanding of his aesthetic preferences that he had not previously articulated. Sascha’s iterations over the labels, followed by his reflections on the selected labels are shown in Figure 5(a) and (b) respectively, capturing how he tried to search for the meaning of these labels in the quadrant tool.

The process of articulating complex judgments into single- or two-word terms did not simplify students’ positions; instead, it distilled them. Positioning labels on the quadrant revealed connections, whether aligned or opposing, and brought latent feelings and judgments to consciousness, enabling students to develop more nuanced self-awareness about their aesthetic and emotional responses.



**Figure 5: Labels chosen by Sascha relating to his topic “Vibes at a festival”, with photographs of festival and concert stages as his dataset. The figure shows (a) his iteration over the labels using his quadrant tool, and (b) his reflection on the final choice of labels.**

**5.1.1 Articulating Subjective Judgments Through Structured Label Selection.** The quadrant tool and label selection process helped students articulate and understand their own complex judgments about their topics of interest by encouraging them to consider labels

**5.1.2 Anticipating AI Analysis and Rethinking Design Decisions.** The process of preparing data for the AI model influenced students’ approach to creating their object portraits. The anticipation of their work being “*interpreted*” by the AI triggered a form of self-critique, as students began to view the creative process through a lens shaped by the annotation experience.

Sascha’s reflection encapsulates how participating in data labeling reshaped his creative orientation: “*It’s so funny how you’re not just training an algorithm. You’re also training your own gaze. You might start to see whatever you’re looking for in the images also in the real world.*” This insight reveals how the labeling process extends beyond technical annotation to become a form of perceptual training for the students themselves. When creating his object portrait,

Sascha deliberately incorporated elements he had labeled “*other-worldly*,” using altered blue lighting, scaled-up toys and a “*Mark Zuckerberg in the style of a cyborg*” figure, anticipating how his trained model might interpret these surreal elements. Similarly, Yann noticed his use of what he called “*DIY aesthetics*” when laying out objects from his previous design work, and started to wonder about how these might embody “*rebellion*,” one of his chosen labels. The knowledge that their works would be “seen” by their own AI model seemed to instigate students to anticipate interpretations, which guided how they would see and create their object portraits.

**5.1.3 Symbolic and Metaphorical Meaning Discovery Through AI Interpretation.** Many students reported that the model predictions—regardless of whether they were anticipated or unexpected—served as starting points for speculation about the artistic and symbolic meaning of their work. Even though they had prepared the training data themselves, they still perceived the AI model as offering new, unexpected perspective on their creative work. These ideas were not based on the technical accuracy of the model, but rather abstracted insights about the meaning of their work.

For example, Aisha’s reflection on her object portrait analysis demonstrates this symbolic discovery process. When her carefully considered bluish white lantern was not detected but a candle received all four of her labels unexpectedly, this made her reflect on its metaphorical meaning: “*in a metaphorical way, I liked this part. When I light up the candle, it’s turning into ‘dominant’, and when I blow it out it disappears. Cultural or religious activities see fear and hope at the same time. So here, I can really see what I want to say in this work.*”

Similarly, Pepa’s initial disappointment with her model’s interpretation led to broader insights about cartoon aesthetics and age-appropriate media. Her model detected only a small area in her portrait using just one label (“*comforting*”), despite her having labeled most faces in her dataset as “*frightening*.” “*I think this shows that the type of media we would call age-appropriate for children is very far from real life.*” The AI’s analysis became a catalyst for reflecting on the broader cultural implications of her chosen visual domain.

The deliberate evaluation of the model’s interpretations by the students—whether anticipated or unanticipated, and whether fully understood based on their annotator behavior or not—seemed to create a space for speculating on multiple levels of the creative process they had undergone. This ranged from attempting to grasp their annotator behavior through the model’s detections, to engaging in higher-level speculation on the potential meanings of their work, where the model’s detections served as inspiration or departure points for further reflection. Thus by training the model and anticipating on the model analyzing their own work, the students had also taught themselves new ways of looking at their work.

## 5.2 Learning About AI Through Hands-On Practice

Learning *about* the AI encompasses findings where students gained explicit knowledge about AI systems themselves: understanding how machine learning models function, discovering the challenges of data annotation, and attaining a form of AI literacy through hands-on engagement.

**5.2.1 Discovering Machine Learning Principles Through Model Performance Analysis.** Students gained insights into fundamental machine learning principles by analyzing instances where their models performed in unexpected ways, leading to direct understanding of how these systems process and interpret information. These moments of surprise created opportunities to grasp core concepts about machine learning that typically remain abstract in theoretical explanations.

Eva’s observation of her meme dataset analysis exemplifies this learning. She noted that her model recognized “*a whole text paragraph, even though I didn’t always annotate the full piece of text. It was sometimes just a singular word...*” This unexpected behavior taught her about how AI systems can generalize beyond exact training examples, learning to detect broader visual patterns rather than only the specific elements that were annotated. Her reflection that the AI “*did put the labels in almost the right place every single time*” despite not being able to “*read the text because it’s only analyzing the pixels*” revealed core principles about how computer vision systems process information—focusing on visual patterns rather than semantic content.

Similarly, students learned about pattern recognition principles when their models detected similarities they hadn’t explicitly taught. Borisz’s realization that his model labeled the same visual element as both “*majestic*” and “*repulsive*” led him to understand: “*I labeled animal legs majestic, and satyr legs repulsive, but they are actually the same if you just look at the legs. Now that I think about it, I understand why.*” This moment illuminated how machine learning systems focus on visual features rather than contextual meaning—a fundamental principle about AI processing that emerged through hands-on analysis rather than theoretical instruction.

**5.2.2 Discovering the Annotation Consistency Challenge Through Practice.** A central learning outcome concerned the fundamental tension between subjective human judgment and the requirements of machine learning training. Students discovered firsthand that while personal judgment is inherently fluid, consistent annotation remains crucial for reliable model performance. This challenge manifested as students found their understanding of their chosen labels shifting over the course of annotating a large amount of images. The repetitive and enduring nature of the labeling task led to a gradual evolution in their interpretations, creating a practical dilemma that illuminated core principles of AI training. As Yann reflected: “*for me all the words change their meaning a bit*” during the annotation process. Students developed three distinct strategies to navigate this tension:

**Revision strategy:** Some students chose to adapt by revising their labels to better capture their evolving judgments. Yann, working with a dataset of abandoned buildings, experienced a significant shift in his perception of labels during the annotation process. His understanding of “*pretentious*” evolved: “*I would now understand [it] as ‘hiding’ and not really the ‘pretentious’ part. Hiding from something that it is not,*” he explained, suggesting that modern buildings began to feel like they were hiding from a certain “*reality*” that abandoned buildings conveyed. Similarly, his label “*inconsistent*,” which initially had negative connotations, “*became way more positive*” over time as he began to perceive those elements in a positive



Figure 6: Samples from annotated training data and object portraits by Aisha and Pepa, discussed in Section 5.1.3.

light, almost discerning what he felt was the original artist’s intended meaning. In response to these shifts, Yann chose to revise his entire labeling scheme—replacing “restricting” with “rebellious,” “reductive” with “inconsistent,” and “lying” with “devotion”—then re-annotated his entire dataset. This approach prioritized authentic representation of his current understanding over maintaining consistency with his initial judgments.

**Acceptance strategy:** Others continued labeling with their original labels while accepting the inevitable shift in judgment, adopting a more flexible approach to the annotation task. Sascha, who worked with festival stage imagery, acknowledged: “I made those choices at one point at the beginning and then I realized that sometimes it didn’t apply. I was like, ‘okay, here’s this element so I’m going to annotate it’, and then when I look back at the whole image, I’m like, ‘actually, it should be something different, something else.’” Rather than revising his labels or forcing strict consistency, Sascha adapted his interpretation approach while keeping the original label structure. He explained how his understanding of “convivial” evolved: “I had the label ‘convivial,’ and I think in the beginning I interpreted that more in the actions, like hugging or sitting together, which was really hard to identify. Then I switched to things like light sources, for example, if it’s soft light compared to the harsh colors in all animations.” This strategy represented an acceptance of the inherent subjectivity in

the annotation process, allowing for interpretive flexibility while maintaining the original conceptual framework.

**Consistency strategy:** A third approach involved maintaining strict consistency with model performance explicitly in mind, prioritizing technical reliability over evolving personal interpretations. Borisz, working with Rubens paintings, developed a systematic approach that helped him maintain consistency even as his internal understanding shifted. He created what he described as a specific labeling code, focusing on elements that were easy to categorize: “Yeah, first I was labeling hands as majestic. But then I realized they are more sublime to me, especially since they always had some symbolic meaning in the paintings. So I went back and relabeled them. I think that helped with the consistency.” Once he made these initial adjustments, Borisz remained deliberately consistent: “I was super consistent in labeling... even when I started changing things in my head, like the meaning changed... I let go of some things, but I didn’t really label them differently.” This approach reflected a strategic decision to prioritize the model’s ability to learn clear patterns over authentic representation of evolving personal judgment, demonstrating an understanding that consistent input would lead to more predictable AI behavior.

Through grappling with this practical challenge, students gained direct insight into one of the fundamental tensions in machine learning: the need to transform fluid human judgment into consistent

(a) Sample images from Borisz's annotated training data



(b) Sample images from Yuky's annotated training data



(d) Yuky's object portrait interpreted by her AI model

(c) Borisz's object portrait interpreted by his AI model

Figure 7: Samples from annotated training data and object portraits by Yuky and Borisz, discussed in Section 5.2.1.

training data. Regardless of which strategy they adopted, students discovered that the annotation process functioned as more than a technical requirement—it became a form of meta-reflection where “teaching” the AI about their chosen topics simultaneously revealed the evolving nature of their own perceptions and judgments. This dual learning process illuminates how Reflective AI engagement can make visible the subjective choices typically hidden within seemingly objective technical tasks, transforming annotation from a means to an end into a reflective practice that generates insights about both AI systems and one’s own ways of seeing.

## 6 Discussion

Our Object Portrait Workshop explores how Reflective AI—what we propose as a slow technology approach to AI in creative practice—may present an alternative approach to efficiency-focused AI integration in design education. Through structured, slow and hands-on engagement with AI components, we observed how students developed both deep reflection on their creative practice and practical understanding of AI systems. In this section, we first consider the implications of our workshop for integrating AI technologies into design education, examining how it cultivated reflection and agency among students. We then explore what these educational insights might suggest for Reflective AI in design practice more broadly, considering how the slow technology principles adapted to AI could extend beyond discrete artifacts and instead engage with the complex temporal assemblages that constitute AI technologies.

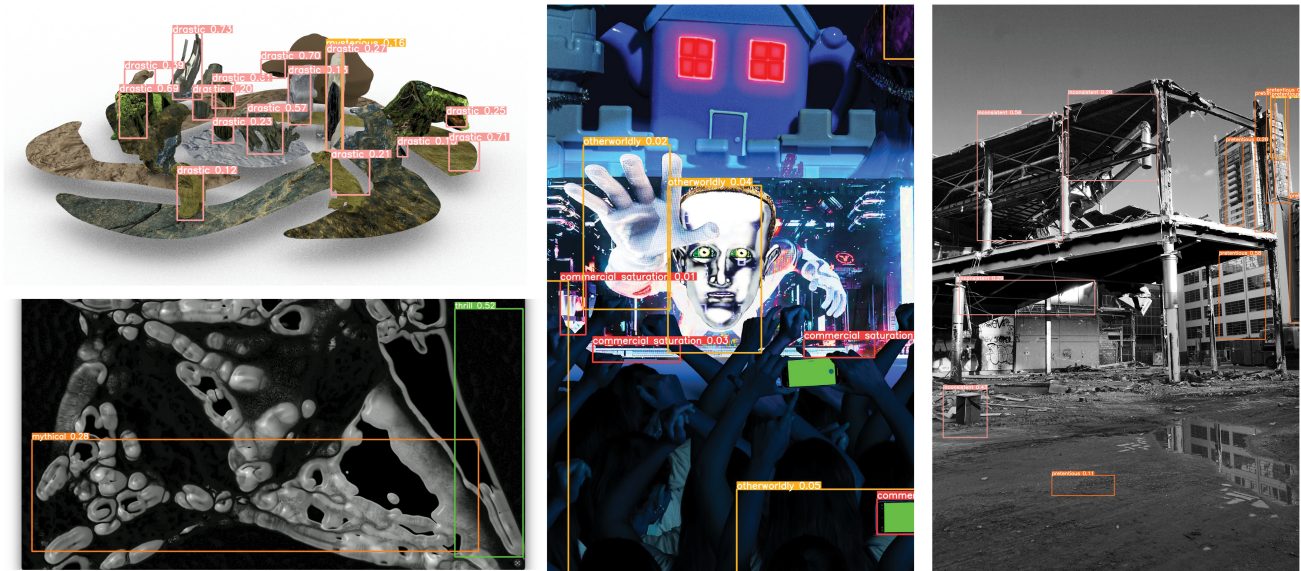


Figure 8: An overview of the Object Portraits created in the workshop by the students, illustrating diverse artistic interpretations and the visible application of AI object detection and labeling as a reflective medium within their creative process.

### 6.1 Integrating Reflective AI into Design Education: Pedagogical Implications

To synthesize the educational insights from our workshop, we identify two core pedagogical principles for implementing Reflective AI: *Intentional Slowness* and *Material Disentanglement*. These principles serve as actionable mechanisms for educators to cultivate reflection. By prioritizing slowness (Section 6.1.1), we shift the focus from output generation to recursive inquiry; by disentangling the material pipeline (Section 6.1.2), we foster agency and critical literacy. We detail these implications below.

#### 6.1.1 Shifting Focus from AI Output to Reflective Practice.

AI-empowered tools in creative processes—often presented as polished, prepackaged chat- or text-based interfaces—tend to sideline the reflective practices foundational to design education. Our findings show that, in contrast, the Object Portrait Workshop’s “intentional slowness” [19] in label selection, annotation, and model usage created a reflective space, enabling participants to engage deeply with both their own creative practices and the technology itself. By extension, a Reflective AI approach seeks to reframe AI integration in design education—not as an efficiency tool but as a mirror for reflecting on the creative process.

The evidence from design students reveals multiple forms of reflection that emerged through this structured approach. The label selection process enabled design students to distill complex emotional responses into nuanced distinctions, revealing deeper self-understanding about their aesthetic judgments and creative preferences. Through systematic engagement with the quadrant

tool, design students developed structured understandings of previously unarticulated feelings and preferences within their chosen domains. Recall Carlo’s process of “personifying” labels to distinguish between “enticing” and “cloying” forms of admiration, and Sascha’s mapping of festival “vibes” onto axes of “commercial saturation” versus “otherworldly,” exemplify how this structured process surfaced nuanced aesthetic judgments (Section 5.1.1). This aligns with our goal of treating AI as reflective media in the constructionist sense [39], where design students co-construct understanding through direct engagement rather than through streamlined interfaces.

The workshop’s structure played a crucial role in facilitating this reflective practice. When design students “taught” an AI about themselves, they simultaneously learned how they perceived their own work, creating a recursive learning process. This exemplifies the reflective “recursive effect” [24] operating not just between maker and making, but among the various components in the making process. Crucially, the outcome of the workshop process—the prediction of labels on participant-created Object Portraits—was not AI-generated content to be consumed, but rather served as a reflective vehicle for students to further consider the effect of inevitably reductive labels and the unpredictability of the model’s assignment of these labels. Ultimately, the workshop’s success is defined not by the quality of the AI’s output, but by the depth of the student’s reflective engagement with their own creative process. This was evident when students found metaphorical meaning even in the AI’s “failures” or unexpected interpretations. As Aisha reflected on her model’s analysis of a candle, the output became a catalyst to

“*really see what I want to say in this work*,” transforming the AI from an analysis tool into a reflective medium (see Section 5.1.3).

**6.1.2 Developing Agency and Critical Literacy Through Material Engagement.** We identified a persistent “gap in power dynamics” [21] and “diminished sense of agency” [13, 60] as key challenges when designers integrate AI into their practice. AI technologies are reported to trigger “creativity exhaustion” [29] and “design fixation” [22, 25]. We also argued that design education needs to help students examine AI technologies not as opaque market solutions but as sociotechnical assemblages of components that can be disentangled and engaged with critically to potentially overcome those challenges. The workshop addresses this through its multi-phase structure, where students engage with data collection, annotation, training, and inference as distinct yet interconnected processes.

Employing Reflective AI according to the slow technology principles, in contrast, is a means to develop designerly agency over AI technologies. With designerly agency we mean the capacity to act on and via tools and materials in a reflective manner (cf. [46, 48]), rather than taking them as given. The evolving annotation strategies we observed demonstrate this agency development in practice. Some students adopted revision strategies, fundamentally reconceptualizing their labels and re-annotating their datasets to authentically represent their evolving understanding. Others developed consistency strategies, maintaining systematic approaches while remaining aware of their shifting internal interpretations, as we observed with Borisz and Yann (see Section 5.2.2). These strategies represent different approaches to negotiating the core tension between subjective judgment and the conventional goal of ‘objective’ technical efficiency—a fundamental challenge in AI and specifically machine learning that students discovered through practice rather than theoretical instruction.

The shifting meanings of subjective labels that we observed unfolded from a starting point associated with a particular artifact (our quadrant tool), but reverberated across all steps in the entire pipeline, calling for reflective recalibration: questioning and reorienting labels, discarding and developing associations. The evolving strategies of the students in annotation took shape over time in correspondence with the reflective application of labels. Both shifting meanings and evolving strategies were subject to the anticipation of an eventual output by the trained model, which itself lay on the other side of the scheduled interval of model training—a point of no return for students’ model customization within the context of our workshop.

This development of reflective strategies was afforded by the workshop’s temporal structure, which teaches students to understand AI not as an immediate response system but as a process with distinct phases, each requiring different forms of engagement and reflection. Through disentanglement, the components of the step-by-step technical pipeline of an AI technology actually dissolved into a rich and temporally intricate design space. In that space, our students navigated slowly the different facets of temporal interaction between components (e.g., anticipation, reorientation, speculation). Regardless of which annotation strategy they adopted, students discovered that the process functioned as more than a technical requirement—it became a form of meta-reflection where

“teaching” the AI about their chosen topics simultaneously revealed the evolving nature of their own perceptions and judgments.

Our Reflective AI design practice echoes work on artistic practices with AI technologies, particularly the notion of “model crafting” advocated by Steinbrück and colleagues [52], who argue that artists gain creative control with the ability to steer the AI system toward one’s creative intentions. Caramiaux and colleagues [10], based on their interview study with AI artists, describe how crafting best encapsulates how these artists interact with their tools, implying a “direct relationship with the algorithmic and data material”. Abuzurairi and colleagues [1] further highlight that this craftsmanship—experimenting with training sets, model architectures, and manipulating pre-trained models—offers a sense of ownership. However, while valuable, this approach does not always directly map to the methodological needs of design research as practiced in the HCI community, emphasizing why structured educational approaches like our workshop are essential for fostering designerly agency in this context.

This structured approach is particularly crucial for students still developing confidence in their designerly intuition and reflective abilities. While professional artists can rely on established creative practices, students require scaffolding to engage with AI’s complexity without becoming overwhelmed or accepting outputs uncritically. The annotation strategies that emerged represent not just technical approaches but developing design philosophies about maintaining creative integrity when working with probabilistic systems, which is essential preparation for professional contexts where such negotiations occur without support.

## 6.2 Toward a Broader Reflective AI Design Practice

While our findings demonstrate the educational value of integrating Reflective AI approaches, they also suggest tentative implications for a broader design practice beyond educational contexts. The reflective processes we observed—label meaning evolution, annotation strategy development, anticipatory creative decision-making—point toward a different paradigm for engaging with AI technologies in professional design practice that further extends the slow technology principles we advocated for.

**6.2.1 Material and Temporal Disentanglement as Core Practice.** Our introduction positioned slow technology as offering intentional slowness that prioritizes reflection over efficiency goals, transforming technologies from productivity tools into reflective media. In practice, we argue that the core characteristic of a *Reflective AI design practice* is a parallel material and temporal disentanglement of AI technologies. With this, we pick up entanglement as an increasingly prevalent ontological and epistemological standpoint for design research in HCI as argued by Frauenberger [14], and position *disentanglement* as an actionable principle for Reflective AI design practice in response.

Material and temporal here are seen as two sides of the same coin, needed to overcome AI technologies’ conventional appearance as an output-supplying artifact. The latter has the risk to play a superficial, fixed role in design practice: that of producing content. Similarly, from the perspective of slow technology, simply delaying

or artificially slowing down this role would not change the conventional logic of an “efficient” readymade tool. Instead, building on our approach to and findings from the Object Portrait workshop, we suggest that AI technologies in a Reflective AI design practice should be framed as a technical assemblage of specific components (e.g., data preprocessing, labeling, training, prediction) that can be disentangled, with each possessing diverse yet interwoven material and temporal characteristics.

Accordingly, to overcome the “artifact bias” of slow technology, a Reflective AI design practice does not only consider the level of temporal interaction with a specific function (e.g., delay or constraints thereof), but rather the various ways in which components act on each other in and across various facets of slowness. This connects our argument to work in design research which is based in considering the composite aspect of computational artifacts not as a given, but something that continuously takes shape (cf. [43, 56]).

**6.2.2 Pairing Technical Uncertainty with Designerly Ambiguity.** Our findings on how students negotiated unpredictable shifts in meaning suggest that pairing technical uncertainty with designerly ambiguity can be a powerful resource, both as a pedagogical strategy in design education and as a generative approach in professional practice. Ambiguity as a designerly resource is well established, both from an epistemological position as well as a practical aspect of engaging with materials (see e.g. [17, 46]). Understanding the design process in this way, outcomes are not predetermined as answers to positivistic hypotheses, but rather assume shape, value and qualities underway.

AI technologies may require more specific considerations regarding what ambiguity means in the design process. To this end, and drawing on the unpredictability of label prediction in the Objective Portrait workshop, we propose to pair designerly ambiguity with technical uncertainty. The latter refers to the probabilistic aspect of the techniques underlying AI technologies (i.e., certainty of prediction due to noise or model variance), and has been proposed by Benjamin and colleagues as a qualitatively distinct design material [3], which triggered explorations of metaphors and interaction modes (cf. [4, 51]).

In the educational context of our workshop, we observed the anticipatory effect of uncertainty on designerly ambiguity. For example, students’ awareness of the model’s eventual, unpredictable reading of their work shaped their creative choices when composing their Object Portraits. Students directly navigated the ambiguity of creating an Object Portrait through specific steps, while the impact of technical uncertainty was experienced afterward—through (non-)detections with particular labels, overlaps in detections, and the confidence scores for each detection. This suggests a pedagogical opportunity for future Reflective AI design practices: framing technical uncertainty as an explicit design material. For example, designers may consider to inject noise and variance at distinct points. In the context of creating an Object Portrait, a designer may add deliberate outlier images to the small dataset, use highly contrasting labels, add random annotation aberrations, and so forth.

In this sense, we see the pairing of ambiguity and uncertainty in Reflective AI design practices as an alignment in terms of its processes (i.e., methodical craft decision-making) as well as materials (i.e., AI technologies as technical assemblages). In sum, we suggest

that the essential methodological characteristic of a Reflective AI design practice is crafting agency by navigating ambiguity and technical uncertainty.

**6.2.3 From Efficiency to Envelopment.** The prevailing “move fast and break things” [20] mentality of mainstream AI tools and technocultures stands in stark contrast to our Reflective AI approach, which offers a concrete alternative through what Hallnäs [18] terms “envelopment”: shifting from fast development for immediate use toward deep understanding. Envelopment treats technology as “a technical and methodological locale we encircle exploring, mapping out an expressional landscape” rather than mere techniques for reaching objectives. This perspective transforms AI interaction from productivity-focused task completion to exploratory practice requiring skillful engagement, where the goal is understanding how technology works, why it works that way, and what consequences follow from its use. This envelopment approach manifests practically through what Scurto and colleagues call “shallow models” [47]. When training a model reflectively, as in our workshop, the goal is not creating perfectly functioning systems based on extensive datasets and high computational power. Instead, these models focus on learnability from smaller datasets, carrying significance for their trainer by virtue of being open to practice-led disentanglement. Through such sustained engagement—circling around and exploring the model rather than optimizing it—practitioners develop intimate knowledge of their models’ behaviors, biases, and boundaries. This knowledge emerges not from performance metrics but from the exploratory process itself. Importantly, this shift from efficiency to envelopment also reveals and challenges the values embedded in conventional AI development. The positivist notion of efficiency privileges particular power relations: studies show that data labeling processes downplay domain experts’ expertise [44] and reflect organizational hierarchies [32]. By contrast, reflective data labeling as a protest and practice has already been proposed to create better AI services for niche communities, where generic labels would fail to cater [2]; and we see a Reflective AI design practice as a connection to this work as well as a potential extension of such an approach from labeling to model training and application.

**6.2.4 Reframing Interpretability and Customization through Reflective AI.** Reflective AI design practice opens up new possibilities for interpretability as user sense-making and customization as practice-specific model adaptation, shifting focus from efficiency and accuracy toward reflection and meaning-making. In the realm of interpretability, this approach offers an alternative for stakeholders without technical expertise, prioritizing values beyond the “accurate” accounts of automated decision-making emphasized in developer-focused techniques. Rather than attempting to “explain away” technical aspects of predictions [3], Reflective AI emphasizes interpretation as a design process. In our workshop, although the AI system remained technically opaque, students’ slow interaction with its components created reflective space for speculation and domain-relevant insights. Their iterative processes of label selection and annotation allowed them to question and redefine meanings within their domains, suggesting that Reflective AI interpretability design could reorient parameters away from efficiency-centered approaches toward supporting annotation practices that scaffold stakeholders’ own “explanation strategies” [5].

Similar opportunities arise in the context of customization. Existing AI customization services (e.g., *exactly.ai*) reproduce the ready-made conundrum by presenting seamless, closed interfaces that streamline workflows but preclude reflection on dataset selection, labeling choices, or annotation strategy development. Reflective AI design instead points toward supporting smaller, domain- and practice-specific model building [26, 53], where “building small” becomes a method for creative exploration. Our approach demonstrates how data collection and annotation can serve as reflective design practices, opening new possibilities for experimentation in generative AI development. Such practices may foster reflexivity and enable designers to regain ownership over creative expression, providing concrete contexts in which Reflective AI can enrich design education and practice.

### 6.3 Limitations and Future Work

Within the context of this workshop study we would like to address some of the challenges and limitations observed. Our students were not expert designers themselves, but rather students. This innate openness due to an as yet unrefined designerly identity was beneficial for our specific setting. To leverage the educational setting further, though, remained an unexplored aspect. The future work could focus on the opportunities of the method for design education in which reflection itself plays an important role. Principles of AI technologies and how to integrate AI in design education are also recent topics worthy of further investigation.

Our chosen technology—object detection using YOLOv5—differs substantially in UX and capability from popular generative systems (e.g., diffusion models). Yet the merits of a Reflective AI design practice are not confined to this model choice. Generative AI systems also operate through complex technical pipelines that are typically hidden from users but could equally be disentangled to cultivate reflection. Extending the principle of disentanglement to the pipelines of generative models presents a promising direction for cultivating reflective practice. Educators might approach Generative AI by designing constraints that disrupt the immediate generation of images. For instance, a “Reflective GenAI” curriculum could require students to curate small datasets for custom model adaptation, shifting the focus from prompt engineering toward data stewardship. Furthermore, educators could introduce material interventions to this pipeline similar to our Objective Portrait process, where students must physicalize their prompts—translating text into tangible materials, like clay or textile—in order to generate an understanding of the logic of AI. These constraints reintroduce the critical friction necessary for learning and prevent the tool from serving solely as an efficiency engine. Finally, while all students followed the same structured process, their Object Portraits varied widely in visual, symbolic, and semantic qualities. This heterogeneity suggests that Reflective AI design practices will generate diverse outcomes, regardless of the specific task (classification, regression, generation). Future work should further investigate how methodological structures can support this diversity across different AI assemblages and design contexts.

## 7 Conclusion

In this paper, we presented an exploration on how the slow technology design approach could inform an alternative design practice to instant, output-oriented AI-powered technologies, that transforms AI tools from productivity tools to reflective media. Building on a three-day workshop in which participants engaged in a reflective, slowed-down process of finetuning and interacting with an object detection model, we derived initial methodological characteristics for Reflective AI education and design practices. Firstly, we argued that the core focus of such practices lies in the material and temporal disentanglement of AI technologies from readymade, content-producing artifacts into technical assemblages of data processing, annotating, training, and prediction components. Secondly, we suggest that specifically with regards to AI technologies, practitioners of Reflective AI education and design practices craft agency over their process by navigating ambiguity (e.g., in labeling) and technical uncertainty (e.g., in prediction probabilities). This approach and the arguments brought forward may help designers to engage with these archetypal “fast” technologies in a more in-depth manner, to integrate them as technical assemblages into their design processes, and to assess this integration and what it leads to with a critical perspective.

### Acknowledgments

We thank all the students who participated in this workshop. We are grateful to studio leaders Rhiarna Dhaliwal, Emmy Bacharach and Ibiye Camp for offering us the space to conduct this research within their Bachelor program, and to Colin Keays for his help on-site. Finally, we thank Design Academy Eindhoven for hosting and facilitating this research.

### References

- [1] Ahmed M. Abuzurairq and Philippe Pasquier. 2024. Towards Personalizing Generative AI with Small Data for Co-Creation in the Visual Arts. In *Joint Proceedings of the ACM IUI 2024 Workshops (CEUR Workshop Proceedings, Vol. 3660)*. CEUR-WS.org, Aachen, Germany, 1–14.
- [2] Anne Arzberger, Maria Luce Lupetti, and Elisa Giaccardi. 2024. Reflexive Data Curation: Opportunities and Challenges for Embracing Uncertainty in Human-AI Collaboration. *ACM Transactions on Computer-Human Interaction* 31, 6, Article 74 (aug 2024), 33 pages. <https://doi.org/10.1145/3689042>
- [3] Jesse Josua Benjamin, Arne Berger, Nick Merrill, and James Pierce. 2021. Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3411764.3445481>
- [4] Jesse Josua Benjamin, Heidi Biggs, Arne Berger, Julija Rukanskaitė, Michael B. Heidt, Nick Merrill, James Pierce, and Joseph Lindley. 2023. The Entoptic Field Camera as Metaphor-Driven Research-through-Design with AI Technologies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3544548.3581175>
- [5] Jesse Josua Benjamin, Christoph Kinkeldey, Claudia Müller-Birn, Tim Korjakow, and Eva-Maria Herbst. 2022. Explanation Strategies as an Empirical-Analytical Lens for Socio-Technical Contextualization of Machine Learning Interpretability. *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (Jan. 2022), 39:1–39:25. <https://doi.org/10.1145/3492858>
- [6] Jesse Josua Benjamin, Joseph Lindley, Elizabeth Edwards, Elisa Rubegni, Tim Korjakow, David Grist, and Rhiannon Sharkey. 2024. Responding to Generative AI Technologies with Research-through-Design: The Ryelands AI Lab as an Exploratory Study. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 1823–1841. <https://doi.org/10.1145/3643834.3660677>
- [7] Karl-Emil Kjør Bilstrup, Magnus Høholt Kaspersen, Niels Olof Bouvin, and Marianne Graves Petersen. 2025. The ML-Machine Toolkit: Empowering Teachers and Education Professionals to Explore Embodied Approaches to Teaching Machine

- Learning. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 535–551. <https://doi.org/10.1145/3715336.3735737>
- [8] Virginia Braun, Victoria Clarke, Nikki Hayfield, Louise Davey, and Elizabeth Jenkinson. 2023. Doing Reflexive Thematic Analysis. In *Supporting Research in Counselling and Psychotherapy*. Springer International Publishing, Cham, 19–38. [https://doi.org/10.1007/978-3-031-13942-0\\_2](https://doi.org/10.1007/978-3-031-13942-0_2)
- [9] Jenna Burrell. 2016. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society* 3, 1 (Jan. 2016), 205395171562251. <https://doi.org/10.1177/2053951715622512>
- [10] 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, CSCW2 (2022), 477:1–477:24. <https://doi.org/10.1145/3555578>
- [11] Michele Cremaschi, Max Dorfmann, and Antonella De Angeli. 2024. A Steampunk Critique of Machine Learning Acceleration. In *Proceedings of the ACM Designing Interactive Systems Conference (DIS '24)* (Copenhagen, Denmark). Association for Computing Machinery, New York, NY, USA, 246–257. <https://doi.org/10.1145/3643834.3660688>
- [12] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* 10, 28 (2024), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- [13] Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2024. The AI ghostwriter effect: When users do not perceive ownership of AI-generated text but self-declare as authors. *ACM Transactions on Computer-Human Interaction* 31, 2 (2024), 1–40. <https://doi.org/10.1007/s13347-022-00543-1>
- [14] Christopher Frauenberger. 2020. Entanglement HCI: The Next Wave? *ACM Transactions on Computer-Human Interaction* 27, 1, Article 2 (feb 2020), 27 pages. <https://doi.org/10.1145/3364998>
- [15] William Gaver. 2012. What Should We Expect from Research Through Design?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 937–946. <https://doi.org/10.1145/2207676.2208538> event-place: Austin, Texas, USA.
- [16] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference (DIS '22)*. Association for Computing Machinery, New York, NY, USA, 517–526. <https://doi.org/10.1145/3532106.3533524>
- [17] William W. Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a Resource for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)* (Ft. Lauderdale, Florida, USA). Association for Computing Machinery, New York, NY, USA, 233–240. <https://doi.org/10.1145/642611.642653>
- [18] Lars Hallnäs. 2015. On the Philosophy of Slow Technology. *Acta Universitatis Sapientiae, Social Analysis* 5, 1 (2015), 29–39. <https://acta.sapientia.ro/en/series/social-analysis/publications-acta-social-analysis-contents-of-volume-5-no-1-2015/on-the-philosophy-of-slow-technology>
- [19] Lars Hallnäs and Johan Redström. 2001. Slow technology—designing for reflection. *Personal and ubiquitous computing* 5 (2001), 201–212. <https://doi.org/10.1007/PL00000019>
- [20] Brett A. Halperin and Erin McElroy. 2023. Temporal Tensions in Digital Story Mapping for Housing Justice: Rethinking Time and Technology in Community-Based Design. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)* (Pittsburgh, PA, USA). Association for Computing Machinery, New York, NY, USA, 2469–2488. <https://doi.org/10.1145/3563657.3596088>
- [21] 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. <https://doi.org/10.1145/3715336.3735843>
- [22] Angel Hsing-Chi Hwang. 2022. Too Late to be Creative? AI-Empowered Tools in Creative Processes. In *CHI '22 Extended Abstracts: CHI Conference on Human Factors in Computing Systems Extended Abstracts* (New Orleans, LA, USA). Association for Computing Machinery, New York, NY, USA, Article 201, 9 pages. <https://doi.org/10.1145/3491101.3503549>
- [23] Kristina Höök and Jonas Löwgren. 2012. Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction* 19, 3 (Oct. 2012), 23:1–23:18. <https://doi.org/10.1145/2362364.2362371>
- [24] Don Ihde and Lambros Malafouris. 2019. Homo faber Revisited: Postphenomenology and Material Engagement Theory. *Philosophy & Technology* 32, 2 (2019), 195–214. <https://doi.org/10.1007/s13347-018-0321-7>
- [25] David G Jansson and Steven M Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11. [https://doi.org/10.1016/0142-694X\(91\)90003-F](https://doi.org/10.1016/0142-694X(91)90003-F)
- [26] Issie Lapowsky. 2023. Why Timnit Gebru wants AI giants to think small. FastCompany. <https://www.fastcompany.com/90952740/why-timnit-gebru-wants-ai-giants-to-think-small>. Accessed Feb 2024.
- [27] Alex Leitch and Celia Chen. 2025. Unlimited Editions: Documenting Human Style in AI Art Generation. 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, 1–9. <https://doi.org/10.1145/3706599.3716214>
- [28] Peter R. Lewis and Ștefan Sarkadi. 2024. Reflective Artificial Intelligence. *Minds and Machines* 34, 2 (May 2024), 14. <https://doi.org/10.1007/s11023-024-09664-2>
- [29] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User experience design professionals' perceptions of generative artificial intelligence. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 1–18. <https://doi.org/10.1145/3613904.3642114>
- [30] Martin Lindrup, Rune Moberg Jacobsen, Joel Wester, Niels van Berkel, Dimitrios Raptis, and Peter Axel Nielsen. 2025. Prompt Machine: A Tangible Generative AI Tool for Supporting Children's Learning and Literacy. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 489–505. <https://doi.org/10.1145/3715336.3735673>
- [31] Fang Liu, Junyan Lv, Shenglan Cui, Zhilong Luan, Kui Wu, and Tongqing Zhou. 2024. Smart "Error"! Exploring Imperfect AI to Support Creative Ideation. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–28. <https://doi.org/10.1145/3532106.3533524>
- [32] Milagros Miceli, Martin Schuessler, and Tianling Yang. 2020. Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–25. <https://doi.org/10.1145/3532106.3533524>
- [33] Midjourney Inc. 2023. *Midjourney*. Midjourney Inc., San Francisco, CA, USA. <https://www.midjourney.com>
- [34] Dave Murray-Rust, Maria Luce Lupetti, Johanna Nicenboim, and Wouter Van Der Hoog. 2024. Grasping AI: Experiential Exercises for Designers. *AI & Society* 39, 6 (2024), 2891–2911. <https://doi.org/10.1007/s00146-023-01794-y>
- [35] William Odom. 2015. Understanding long-term interactions with a slow technology: An investigation of experiences with FutureMe. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 575–584. <https://doi.org/10.1145/2702123.2702221>
- [36] William Odom, Mark Selby, Abigail Sellen, David Kirk, Richard Banks, and Tim Regan. 2012. Photobox: on the design of a slow technology. In *Proceedings of the ACM Designing Interactive Systems conference*. Association for Computing Machinery, New York, NY, USA, 665–668. <https://doi.org/10.1145/2317956.2318055>
- [37] William Odom, Ron Wakkary, Jeroen Hol, Bram Naus, Pepijn Verburg, Tal Amram, and Amy Yo Sue Chen. 2019. Investigating slowness as a frame to design longer-term experiences with personal data: A field study of olly. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3290605.3300264>
- [38] OpenAI. 2023. *ChatGPT*. OpenAI, San Francisco, CA, USA. <https://chat.openai.com>
- [39] Seymour Papert and Idit Harel. 1991. Situating Constructionism. In *Constructionism*, Idit Harel and Seymour Papert (Eds.). Ablex Publishing, Norwood, NJ, 1–11. [https://web.media.mit.edu/~calla/web\\_comunidad/Reading-en/situating\\_constructionism.pdf](https://web.media.mit.edu/~calla/web_comunidad/Reading-en/situating_constructionism.pdf)
- [40] Joon Sung Park, Rick Barber, Alex Kirlik, and Karrie Karahalios. 2019. A slow algorithm improves users' assessments of the algorithm's accuracy. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–15. <https://doi.org/10.1145/3359204>
- [41] Brittany L. Peterson. 2017. Thematic Analysis/Interpretive Thematic Analysis. In *The International Encyclopedia of Communication Research Methods*, Jörg Matthes, Christine S. Davis, and Robert F. Potter (Eds.). John Wiley & Sons, Hoboken, NJ, USA, 1–9. <https://doi.org/10.1002/9781118901731.iecrm0249>
- [42] Julien Porquet, Sitong Wang, and Lydia B Chilton. 2025. Copying style, Extracting value: Illustrators' Perception of AI Style Transfer and its Impact on Creative Labor. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3706598.3713854>
- [43] Johan Redström and Heather Wiltse. 2018. *Changing Things: The Future of Objects in a Digital World*. Bloomsbury Academic, London, UK. <https://doi.org/10.5040/9781350004382>
- [44] Nithya Sambasivan and Rajesh Veeraraghavan. 2022. The Deskillling of Domain Expertise in AI Development. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)* (New Orleans, LA, USA). Association for Computing Machinery, New York, NY, USA, Article 587, 14 pages. <https://doi.org/10.1145/3491102.3517578>
- [45] Hauke Sandhaus, Qiuan Gu, Maria Teresa Parreira, and Wendy Ju. 2025. Co-Designing with Transformers: Unpacking the Complex Role of GenAI in Interactive System Design Education. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1228–1243. <https://doi.org/10.1145/3715336.3735805>
- [46] Donald A Schon. 1983. *The reflective practitioner: How professionals think in action*. Basic Books, New York.
- [47] Hugo Scurto, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Prototyping machine learning through diffractive art practice. In *Proceedings of the ACM Designing Interactive Systems Conference (Virtual Event)*. ACM, New York, NY,

- 2013–2025. <https://doi.org/10.1145/3461778.3462163>
- [48] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph 'Jofish' Kaye. 2005. Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility (CC '05)* (Aarhus, Denmark). Association for Computing Machinery, New York, NY, USA, 49–58. <https://doi.org/10.1145/1094562.1094569>
- [49] Phoebe Sengers and Bill Gaver. 2006. Staying Open to Interpretation: Engaging Multiple Meanings in Design and Evaluation. In *Proceedings of the 6th Conference on Designing Interactive Systems (DIS '06)*. ACM, New York, NY, USA, 99–108. <https://doi.org/10.1145/1142405.1142422>
- [50] Hanieh Shakeri, Ye Yuan, Bennett Axtell, Denise Y. Geiskovitch, and Carman Neustaedter. 2024. Designing Smart Home Technology For Passive Co-Presence Over Distance. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (Copenhagen, Denmark) (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 3389–3406. <https://doi.org/10.1145/3643834.3661508>
- [51] Christian Sivertsen and Anders Sundnes Løvlie. 2024. Exploring Aesthetic Qualities of Deep Generative Models through Technological (Art) Mediation. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 2738–2752. <https://doi.org/10.1145/3643834.3661498>
- [52] Alexa Steinbrück and Aeneas Stankowski. 2023. Creative ownership and control for generative AI in art and design. In *Generative AI in HCI Workshop, CHI 2023*: <https://github.com/GenerativeAIandHCI/GenerativeAIandHCI.github.io/blob/main/papers/2023/genaichi-creativeownership.pdf>.
- [53] Eliza Strickland. 2022. Timnit Gebru Is Building a Slow AI Movement: Her new organization, DAIR, aims to show a more thoughtful mode of AI research. <https://spectrum.ieee.org/timnit-gebru-dair-ai-ethics>. Accessed Feb 2024.
- [54] Jakob Tholander and Martin Jonsson. 2023. Design ideation with ai-sketching, thinking and talking with Generative Machine Learning Models. In *Proceedings of the ACM Designing Interactive Systems Conference (Pittsburgh, PA, USA)*. ACM, New York, NY, 1930–1940. <https://doi.org/10.1145/3563657.3596014>
- [55] Ultralytics. 2023. Comprehensive Guide to Ultralytics YOLOv5. <https://docs.ultralytics.com/yolov5/>. Accessed: 2024-02-01.
- [56] Anna Vallgåda and Johan Redström. 2007. Computational Composites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 513–522. <https://doi.org/10.1145/1240624.1240706> event-place: San Jose, California, USA.
- [57] Vera Van der Burg, Almila Akdag Salah, Senthil Chandrasegaran, and Peter Lloyd. 2022. Ceci n'est pas une chaise: Emerging practices in designer-AI collaboration. *Design Studies* 78 (2022), 101073. <https://doi.org/10.1016/j.destud.2021.101073>
- [58] Vera van der Burg, Gijs de Boer, Alkim Almila Akdag Salah, Senthil Chandrasegaran, and Peter Lloyd. 2023. Objective Portrait: A practice-based inquiry to explore AI as a reflective design partner. In *Proceedings of the ACM Designing Interactive Systems Conference (Pittsburgh, PA, USA)*. ACM, New York, NY, 387–400. <https://doi.org/10.1145/3563657.3595974>
- [59] Samangi Wadinambiarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 1–18. <https://doi.org/10.1145/3613904.3642919>
- [60] Justin D. Weisz, Michael Muller, Jessica He, and Stephanie Houde. 2023. Toward General Design Principles for Generative AI Applications. In *Joint Proceedings of the ACM IUI Workshops (Sydney, Australia)*. ACM, New York, NY, 13 pages.
- [61] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–13. <https://doi.org/10.1145/3313831.3376301>

## A Appendix

### A.1 Workshop Task Descriptions

For transparency regarding the design of the workshop, we provide the specific questions and assignments given to students, organized by the workshop timeline.

#### Day 1: Dataset Preparation and Annotation.

*Dataset Collection Criteria.* Prior to the workshop, we instructed students to collect a digital dataset based on the following criteria to ensure feasibility for the object detection model:

- **Quantity:** At least 100 images (collecting more was encouraged).

- **Subject:** Images must depict one specific topic, phenomenon, or concept (e.g., locks, tidiness, surrealism, seating in public space) closely related to the student's research interest.
- **Consistency in Style:** All images must share a visual medium (e.g., all studio photos, or all screenshots).
- **Consistency in Scale:** The images should function as a 'scene' containing elements for the algorithm to detect, rather than isolated close-ups.

*Initial Group Sharing.* We facilitated a collective review where each student presented their chosen dataset and articulated the rationale behind their selection to the group.

*Label Selection Task.* We guided the quadrant-based label selection using four steps:

- (1) Categorize 10 images from your printed set into the quadrants based on their general feel.
- (2) Find details in the images that evoke that feeling most.
- (3) For each quadrant, write a list of words that describes this feeling best.
- (4) Select the four words that capture the underlying judgment most accurately.

*Post-Labeling Reflection and Debrief.* After the labeling task, students responded to the following written prompt: "For each quadrant, write some thoughts on how your view of these terms changed during the process." Following this individual reflection, we held a collective moment focusing on how the students experienced the labeling task and how it influenced their perspective on the data.

#### Day 2: Object Portrait Creation.

*Design Brief.* Students were tasked with creating an artifact to be read by their model. The design constraints required the Object Portrait to be:

- An A2 photo print or a 10-second Full HD video.
- Composed of objects related to their specific work.
- Styled aesthetically to match the training dataset.
- Feasible to produce within a single afternoon session.

*Collective Reflection on Making.* We concluded the day with a brief collective moment where students discussed their experience of translating their dataset aesthetics into a physical Object Portrait.

#### Day 3: AI Analysis and Reflection.

*Model Interpretation.* Upon running their trained models on their Object Portraits, students were asked to share in the group: "Look at the AI analysis. What do you see? What do you think the interpretations mean?"

*Final Individual Reflection.* We concluded the workshop with a structured reflection card containing the following questions:

- What did you learn about your topic and yourself?
- What are moments where you felt your position (views of your topic) changed? (Consider labeling, portrait-making, and interpreting).
- How do you relate to your algorithm? Did that change how you relate to AI in general?
- Highlight one thing you want to share about your process.

## A.2 Final Thematic Structure

Table 3 outlines the final thematic structure that emerged from our reflexive analysis. These themes correspond to the findings detailed in Section 5 .

**Table 3: Final thematic structure derived from the analysis of student reflections.**

Identified Theme		Description of Pattern in Data
<i>Category 1: Learning Through AI (Reflective Medium)</i>		
Articulating Judgments	Subjective	Students using the label selection process to distill complex, fluid emotions into specific vocabulary (e.g., distinguishing 'enticing' from 'cloying').
Anticipating AI Analysis		Moments where students altered their creative work (the Object Portrait) because they internalized how the model might 'see' it.
Symbolic Discovery		Instances where students interpreted model errors or unexpected detections as metaphorical or poetic meaning rather than technical failure.
<i>Category 2: Learning About AI (Technical Awareness)</i>		
Discovering ML Principles		Students deducing technical concepts (e.g., pattern recognition, generalization) through observing model performance (e.g., "it reads pixels, not text").
Negotiating Consistency	Annotation	The struggle between maintaining a consistent dataset for the machine and the student's evolving personal judgment during the process.