

Responding to Generative AI Technologies with Research-through-Design: The Ryelands AI Lab as an Exploratory Study

Jesse Josua Benjamin
Imagination, Lancaster University
Lancaster, United Kingdom
j.j.benjamin@lancaster.ac.uk

Joseph Lindley
Imagination, Lancaster University
Lancaster, United Kingdom
j.lindley@lancaster.ac.uk

Elizabeth Edwards
Lancaster University
Lancaster, United Kingdom
liz.edwards@lancaster.ac.uk

Elisa Rubegni
Lancaster University
Lancaster, United Kingdom
e.rubegni@lancaster.ac.uk

Tim Korjakow
Technische Universität Berlin
Berlin, Germany
tim.korjakow@gmx.de

David Grist
Ryelands Primary and Nursery School
Lancaster, United Kingdom
d.grist@ryelands.lancs.sch.uk

Rhiannon Sharkey
Ryelands Primary and Nursery School
Lancaster, United Kingdom
r.sharkey@ryelands.lancs.sch.uk

ABSTRACT

Generative AI technologies demand new practical and critical competencies, which call on design to respond to and foster these. We present an exploratory study guided by Research-through-Design, in which we partnered with a primary school to develop a constructionist curriculum centered on students interacting with a generative AI technology. We provide a detailed account of the design of and outputs from the curriculum and learning materials, finding centrally that the reflexive and prolonged ‘hands-on’ approach led to a co-development of students’ practical and critical competencies. From the study, we contribute guidance for designing constructionist approaches to generative AI technology education; further arguing to do so with ‘critical responsiveness.’ We then discuss how HCI researchers may leverage constructionist strategies in designing interactions with generative AI technologies; and suggest that Research-through-Design can play an important role as a ‘rapid response methodology’ capable of reacting to fast-evolving, disruptive technologies such as generative AI.

CCS CONCEPTS

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

KEYWORDS

research-through-design, generative AI, HCI education

ACM Reference Format:

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 *Designing Interactive Systems Conference (DIS '24)*, July 1–5, 2024, IT University of Copenhagen, Denmark. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3643834.3660677>

1 INTRODUCTION

Generative AI technologies have rapidly entered public discourse, awareness and usage since late 2022. ChatGPT was the fastest computer system in history to reach 100 million monthly users [40], and the release of Stable Diffusion was the first time such a model was ‘openly’ published [91]. Notable examples of these technologies now residing firmly in the public sphere include widely-shared ‘fake’ images such as the “Balenciaga Pope,”¹ the Writers’ Guild of America strikes [1], and governmental reports such as the UK House of Commons committee report stressing the balance between creators’ rights and innovation [18]. For the HCI community, and design in particular, generative AI technologies have engendered a flood of publications in contexts of use ranging from their integration into design tools or end-user facing components, variously prompting enthusiasm (see e.g. [43]) as well as critique (see e.g. [64]). In both the academic and public sphere, uncertainties around the impact of these technologies abound.

In this paper, we adopt Research-through-Design (RtD) as an overarching methodology as one way to respond to these challenges, testing the approach through an exploratory study in a specific context. In contrast to many canonical examples of RtD that are focused on specific artefacts or devices (e.g., [33]), we embraced RtD’s methodological attributes and qualities including valuing emergence [31] and the aspiration to produce “intermediate-level knowledge” [41] as a primary consideration. In epistemological

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DIS '24, July 1–5, 2024, IT University of Copenhagen, Denmark

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0583-0/24/07.

<https://doi.org/10.1145/3643834.3660677>

¹<https://www.forbes.com/sites/danidiplacido/2023/03/27/why-did-balenciaga-pope-go-viral/>, accessed 10/11/2023.

terms, this necessarily required that we accept our findings would be contingent, and most likely of use for other researchers in a generative rather than prescriptive sense (see [30]). The practical consequence of these epistemological commitments for our contributions is that they should be seen not as proposals for generalizable theory, but rather as one of many examples which will contribute to an emerging “research program” [73] of research relating to the role and use of design in the ongoing adoption of generative AI technologies.

The particular context for testing our approach in an exploratory study is education, as here the adoption of generative AI technologies has caused and is causing significant debates and uncertainties. Text synthesis models have raised fears that students would simply employ such models to ‘cheat’ (e.g., [48]), while technologies intended to detect such usage have been proven to be unable to do so accurately.² Other avenues include the integration of AI technologies for prediction of educational ‘performance,’ which has garnered extensive criticism (e.g., [70]). Further, for many young people growing up today generative AI technologies will be something they interact with throughout their lives. Hence, there is an emerging consensus (see [21, 34]) that considered education strategies are needed to help current and future generations of young people be able to acquire practical and critical competencies for them—in other words, AI literacy (see [60]). In this regard, there is a rich tradition in HCI research in educational contexts that seeks to develop literacies for new technologies; for instance regarding “critical data literacies” [38], and there is nuanced work on AI technologies coming to the fore (e.g., [44, 53, 89]). However, we observe that opportunities for engaging directly with *actual* AI technologies in a prolonged manner have been limited. In turn, we argue that RtD, with its emphasis on novel materials and extended experiences thereof, can provide opportunities for exploration how young people encounter and form competencies with generative AI technologies.

Building methodologically on RtD, we present an exploratory study that designed and delivered an introductory, *constructionist* curriculum to generative AI technology. Constructionist learning posits that models of the world are built through the direct engagement with *materials* (see [45, 69]), and that this is an intrinsic part of how we gain knowledge. In the case of our study this meant that the students were actually *using* the generative AI technology they were learning about, i.e., it was a prolonged, reflective and ‘hands on’ approach to learning. Regarding our overall approach, this means that while we did design several artefacts as part of the project’s delivery (e.g., image generation tool, learning materials and printed outputs), the ‘design’ that is at the center of our use of RtD is the design of the curriculum as a whole. In turn, we argue that this offers intermediate-level knowledge both within the context of the study, i.e. generative AI technology education, as well as methodological impetus on the use of RtD to engage with generative AI technologies reflectively.

The paper is structured as follows. Initially we provide a brief background section discussing related HCI research, focusing on

education and AI technologies to outline our motivation for employing RtD. Next, we discuss the RtD approach in more detail, explaining the rationale for and implications of RtD as the overarching methodological framing for this work. For the exploratory study itself, we subsequently describe the reflexive design, development and delivery of a six week curriculum and the production of several post-hoc artefacts and activities, presenting the work in terms of three key phases. In the study specific findings, we document centrally how practical and critical competencies of students appeared to *co-evolve* and report on the teachers’ views of the curriculum; and gather propositions for HCI research to engage generative AI education in the form of (1) guidance for designing constructionist curricula on generative AI technologies and (2) reflection on the need for ‘critical responsivity’ due to the potential impact of generative AI technologies in educational contexts.³ In our discussion, we then center on the implications stemming from our use of RtD to engage generative AI technologies. Here, we (3) discuss how the design of interactions with generative AI systems may leverage constructionist strategies, and (4) reflect on RtD as a ‘rapid response methodology’ for producing insights pertaining to rapidly evolving disruptive technologies such as generative AI. With these contributions, we aim for HCI researchers to find practical, strategical and methodological support for employing design in the uncertainties characterizing the ongoing adoption of generative AI technologies.

2 RELATED WORK

A full overview of HCI education discourse is beyond the scope of this paper, hence it is important to clarify that we do not refer to education on HCI, but rather HCI *for* education [72]. There are many AI-based resources for education (for example see this collection [66]), however, education-focused HCI research has much to offer in terms of developing accessible tools for teachers and students. While there are many perspectives one could explore this potential from, in this section we initially focus on the constructionist learning theory which is at the heart of this research before summarizing popular approaches to AI literacy in schools.

Among many paradigms, HCI education research has long built on the learning theory of constructionism as proposed by Papert [69]. Closely related to constructivism, which tends to focus on learner experience, constructionism posits that learning can arise from active creation, the development of socially meaningful artefacts, and the combination of self-reflection and interaction with others [19]. The constructionist paradigm explicitly casts students as active learners taking part in interactions with materials which are crucial to learning and knowledge-acquisition (see [4, 46, 56]). The value of constructionism is evidenced in a plethora of toolkits that HCI researchers have developed for researchers, practitioners, students and other stakeholders (see e.g., [14]), and has also served as the guiding philosophy for such hugely successful educational software as *Scratch* [62]. Further, HCI sub-fields such as child-computer interaction have drawn extensively on constructionist principles to combine practical skills and ethics (see [87]),

²<https://openai.com/blog/new-ai-classifier-for-indicating-ai-written-text>, accessed 11/08/2023.

³To further make our work actionable for the field, we provide an OSF repository with all project materials (documentation images, learning materials, design artefacts, anonymized interview transcripts, project outputs) at https://osf.io/9afnr/?view_only=93c8d0e2159847dba2c454519682b85d, accessed 24/04/2024.

leading to teaching approaches that reflect a critical stance on technology (e.g., “critical data literacies” [38]).

There is a high likelihood that generative AI will have a transformative effect on a whole range of social, cultural, and economic issues (see [34]). HCI scholarship can offer a valuable perspective by exploring how to bolster AI literacy with younger stakeholders. The UNESCO K-12 AI education framework defines AI literacy as “some level of competency with regard to AI, including knowledge, understanding, skills, and value orientation” [86]. The consequence of this framing is that AI literacy transcends declarative or explicit technical knowledge alone and must involve a broader awareness of the technology’s societal implications and an ability to critically reflect on these. This aligns with HCI perspectives such as Long and Magerko’s assertion that AI literacy is “a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace” [60].

There is a wide range of HCI research related to education and AI, as shown by Su and Zhong in an extensive review [85]. AI literacy work for middle and high school students includes project-based [44], art-based [89] and ethical impact-focused work [78, 82] as well as more quantitatively oriented work [54]. Druga and colleague’s work employs co-design and prototyping methods to embody an AI agent or prototype potential systems with primary school children [23]. Williams and colleagues note the importance of younger (pre-school) learners gaining AI literacy, citing safety as a motivation as they encounter AI technologies in aspects of everyday life [92], while Han and Cai explore the benefits and challenges of generative AI through interviews with experts (parents, teachers, technologists) discussing an AI-infused storytelling app [37].

There have also already been fruitful examples of HCI design research being used to interrogate AI technologies. An exemplary approach is Dove and Fayard’s work on using a metaphorical, playful comparison of Machine Learning (ML) to monsters such as Frankenstein [22], unfolding concerns, worries, and potential mitigation of adverse effects of ML deployments concerning student mental health. Bilstrup and colleagues used RtD to develop an educational tool for creating ML models [10], and Lindley and colleagues’ employed RtD to interrogate “the emerging reality of living with AI” [58] by designing speculative digital signage. Tamashiro’s use of “design futuring” to explore contemporary AI technologies [3], and Rubegni and colleagues’ investigation into the hopes and fears children have relating to AI using scenarios [75], further demonstrate the value of design-led approaches in this context.

This is clearly a growing and diverse area of research. There are, however, only limited approaches that both involve education practitioners in the process and also center on actual use of the technology as opposed to hypothetical use [94]. Further, exploratory work such as Lee and colleagues’ “Prompt aloud!” [54] has yet to conduct a prolonged engagement of learners on a more qualitative footing. A constructionist approach to learning how to use AI technologies may help address the latter limitation by leveraging a direct and sustained engagement with the technology. It is here that we see a promising opening for testing RtD in the generative AI education space: design-led inquiries of actual technologies have the potential to find radical new applications, concepts and

frameworks, which the disruptive and practically as well as ethically challenging development of generative AI technology calls for. In the next section, we outline our methodological choices before presenting the Ryelands AI Lab project.

3 METHODOLOGY

In this section, we detail the methodological approach behind the Ryelands AI Lab, a constructionist curriculum intervention spanning six weekly lessons with students aged 7 to 9 and a subsequent exhibition of outputs at Ryelands Primary and Nursery School in Northwest England. We also provide contextual detail and provide a consideration of our positionality.

3.1 Research-through-Design as Guiding Methodology

Within our exploratory study, our primary objective was to explore the process of delivering a constructionist approach to teaching practical and critical competencies for generative AI technologies in a primary education context. Such competencies form, according to Long and Magerko, the core of AI literacy: understood not as purely declarative knowledge (i.e., ‘AI is X’) that one receives but rather the kind of knowledge that promotes reflection as well as practical use [60]. To support this kind of AI literacy required the design and development of brand-new learning materials and resources. This is one reason for deciding to employ Research-through-Design. RtD is a methodology that, in general, helps “researchers to investigate the speculative future, probing on what the world could and should be” [84] through the design-led exploration of materials, scenarios and concepts. This is achieved via the process of designing objects, services and other artefacts and leveraging both reflective practice during the process as well as the resultant outcomes to discuss and explore potential impact and implications (see [5, 30, 59, 84, 98, 99]). As highlighted by Williamson, this is an important matter in the generative AI education space insofar as that methods are needed with which to create “alternative social laboratories” [93] that can counter the industry dominance of how these technologies enter society. Coupled with our observation that studies of prolonged engagement of young people with actual generative AI technologies are lacking, this forms the basis for our rationale to engage this context with RtD: The reflective and responsive design *and* simultaneous study of a curriculum intervention as a ‘lab’ for hands-on engagement with adapting materials is made methodologically possible through RtD, offering insights that are not driven by industry headlines or public fears but rather center on how young people engage this disruptive technology directly.

RtD can be a predominantly “artefact-centric” [30] *method*, where it is a “thing-making practice whose objects can offer a critique of the present and reveal alternative futures” [5]. Our work, however, adopted RtD as a guiding *methodology*. This allowed us to lean on one of RtD’s main epistemological strengths; its reflexivity [81], meaning the capacity to engage with a domain or material while continuously “reframing the underlying situation and goal during the design process” [97]. An alternative acknowledgement of the requirement to be flexible about the RtD process is framed as “drifting by intention” [50]. As a result, in RtD knowledge is not expected to

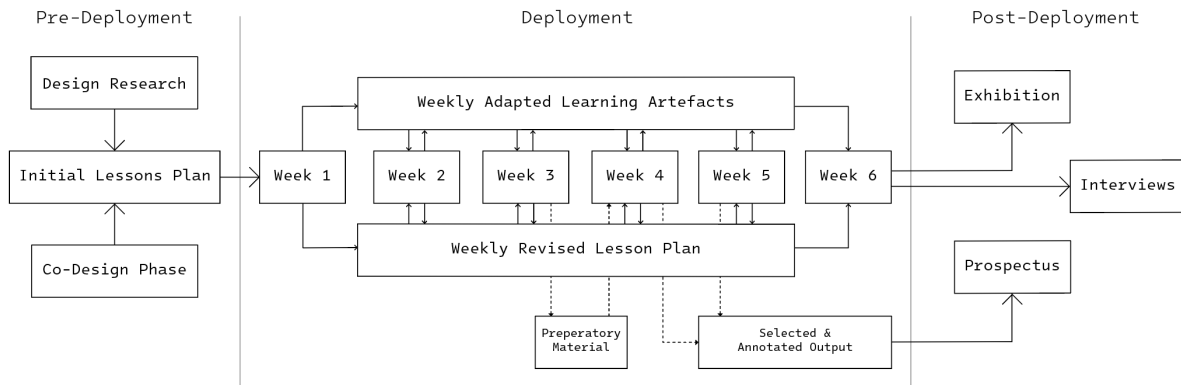


Figure 1: Overview of the exploratory study phases as they manifested *through* the RtD methodology in hindsight.

emerge as a “bounded thing” [84] related to one artefact or hypothesis, but rather as taking shape alongside the various activities of *doing* RtD. In other words, RtD here did not mean that we planned to design one specific artefact and study how its use and contextual embedding would prompt certain propositions or questions. Rather, and similar to epistemological considerations in participatory design (see [36, 79]), we sought to use design to develop a holistic epistemological position on the entirety of a particular context: the design and delivery of a constructionist curriculum on generative AI technologies.

It would have been feasible to imagine adopting a different research methodology that involved creating a curriculum and learning materials, and then studying the efficacy of those externally, for example by establishing a control group and a study group and comparing the impact of the curriculum. Such an approach certainly would have some virtues, however, to facilitate a rapid response to the newly-released technology (in our case we were working with Stable Diffusion which was released only months prior), RtD allowed us to combine the process of designing the curriculum while also assessing its value. An implication of this approach is that the insights we would create would be a type of “intermediate knowledge” [41]. This imports some limitations, for example our findings are based on sample of one (i.e., only one instance of the constructionist curriculum was evaluated). Balancing this limitation is the acceptance that the findings are intended mainly to inform future work, contributing to an ongoing program of research, as opposed to positing that a hypothesis has been evidenced beyond reasonable doubt. Lastly, we were also driven by wanting to better understand how RtD can methodologically respond to a disruptive technology in a specific context—here, generative AI technology in education.

The diagram we present in Figure 1 is a post-hoc representation of how the exploratory study was structured. Elements of this were planned, although some of the structure emerged across the process. The deployment phase depicted was a particularly intense period for the project and involved weekly adaptations to our learning materials and artefacts, ensuring that our learning that developed on a week-by-week basis was incorporated into the subsequent weeks. In the context of this project, this reflexive quality was a clear

benefit of the RtD methodology. Our starting point was undoubtedly provisional, but this allowed us to adapt and shape our findings to provide as holistic a point of view as possible, while accepting the limits of intermediate knowledge (we reflect on further nuances of this process in section 6.2).

3.2 Contextual Overview

Here, we document the specific demands and aspects of engaging in the chosen context of primary school education, and how specific aspects such as recruitment and analysis were shaped by our RtD methodology.

3.2.1 Recruitment and Approach to Context. Prior to the exploratory study, the Head Teacher of the school we were working with suggested working with their year 4 students (aged 7-9). This suggestion was based partly on the considerations that the research intervention would not interfere with any national exams. Additionally, in contrast to younger students this group generally have the skills to use their computers without assistance. After a briefing on the goals and technology involved, we began regular meetings with two year 4 teachers. Each of the teachers has a year 4 class of 25 students, and the teacher deals with all aspects of the curriculum and delivers all lessons for the class they are responsible for.

3.2.2 Demographics. Our participants were preselected as those who were in the two year 4 classes at the school we were working with. The school is situated in an economically and socially disadvantaged area (see section 3.2.4). This is also reflected in the high share of students’ school meals supported by UK government pupil premium grants,⁴ at 63% of the student population.⁵ This makes it roughly representative of the North-West England county it is situated in as a whole, whose GDP sits at 67.7% of the UK national average, and which is also ethnically predominantly White-British.⁶

3.2.3 Research Ethics. In terms of ethical research conduct, we collected informed consent through the year 4 teachers from the

⁴<https://www.gov.uk/government/publications/pupil-premium/pupil-premium>, accessed 08/08/2023.

⁵<https://ryelands.lancs.sch.uk/pupil-premium/>, accessed 24/04/2024.

⁶<https://ryelands.lancs.sch.uk/wp-content/uploads/2022/11/Single-Equalities-Policy-May-2022.pdf>, accessed 24/04/2024.

students' parents or people with parental duties, which we accompanied with an extensive information sheet. We also included various options for withdrawing students from the research aspect of the study without missing the lessons (which formed part of the obligatory curriculum). All participants were free to withdraw from the study at any point on their own accord as well as by decision of their parents, guardians or people with parental duties. We received ethics approval for this study from Lancaster University's ethics committee (reference FASSLUMS-2023-2184-RECR-4). The researchers who delivered the curriculum further obtained a Disclosure and Barring Service (UK) criminal background check.

3.2.4 Teacher Interviews. To ensure we could more conclusively reflect on the pedagogic and didactic elements of our delivery, we also chose to conduct semi-structured interviews with the two year 4 teachers subsequent to the delivery of the Ryelands AI Lab. The prepared questions ranged from general questions on AI and education as well as the experiences during the preparation, delivery and subsequent period of Ryelands AI Lab. The interviews were conducted by the first author and carried out first with each teacher individually for approximately 30 minutes, followed by a 15-minute group discussion so that questions emerging during the individual sessions could be addressed.⁷

3.2.5 Analysis. As mentioned above, we expected the exploratory study to produce “intermediate-level knowledge,” which could in turn be synthesized into “strong concepts” [41] that would inform the study's field; as well as more general implications from using RtD as a research methodology. In this, RtD's reflective capacity extended towards our patterns of analysis as well. As a general and oft-noted strength of design research—going back to Schön's “reflection-in-action” [80]—this meant that the weekly interplay between the initial curriculum, the occurrences of the preceding week, and the development for the upcoming week also shaped on-the-go analyses as well as a schema for findings. This can be further described, therefore, as a “reflection-on-action” [28] beyond the immediacy of designing a particular thing. For instance, while we always assumed that the student's gain of particular practical skills would be an important measure of success for our approach, it only became clear through practice that a crucial qualifier for this measure would be the *intentional* application of skills. We extended this activity-led analysis to the teacher interviews as well, which we considered from an ethnomethodologically informed perspective; meaning that their analysis was primarily directed by the first author's interpretation in relation to other considerations emerging from the delivery and design activities in this project.

3.3 Positionality

In addition to the internal research ethics, we were mindful of further ethical dimensions of generative AI technologies, particularly with regards to the young age of students. While their capacities and the excitement they elicit are tempting, it is now widely understood that AI technologies can discriminate against particular protected categories (e.g., gender, ethnicity, nationality, sexual orientation), while also infringing on intellectual property by scraping



Figure 2: Image from early stage of co-design process, showing second author (middle) and the year 4 teachers.

content (e.g., data, images, texts), creating poorly paid and harmful jobs (e.g., content moderation, labelling), and contributing to the environmental impact of data centers and rare earth mineral extraction. As people without professional primary education expertise, this prompted the first two authors to reflect: is generative AI in schools a good idea in the first place, and are we—two white male Europeans—the people to introduce it to these students? Ultimately, we decided that for this set of students, the practical and critical competencies would aid them dealing with the increasing proliferation of their everyday lives by generative AI technologies. Further, we assumed that the intermediate knowledge we gathered would inform a more substantial critical position that we could add to the generative AI education space (see our proposition on the matter in section 5.2.2) while foregrounding the benefits of constructionist approaches that combine practical and critical competencies.

4 EXPLORATORY STUDY: THE RYELANDS AI LAB

This section details the actual instantiation of the Ryelands AI Lab, heuristically separated into its three significant phases. Again, it is important to note that these phases were only clearly identified in hindsight, and that elements within them overlapped substantially in practice.

4.1 Pre-Deployment: Research and (Co-)Design

First, we conducted *initial research and co-design* activities with the year 4 teachers which produced the first lesson scripts and curriculum shape. This was followed by the *design of learning artefacts* (i.e., image generation tool, slides, worksheets and secondary learning materials) both took shape from and in turn shaped the lesson scripts during actual delivery. Furthermore, all finalized design artefacts can be seen contextualized on a Miro board.⁸

4.1.1 Initial Co-Design. From the start, we were mindful of supplementing the student's ongoing curriculum. To this end, we referred to the UK Department of Education's curriculum guidance; where we found that Keystage 2 Year 4 students' design and technology

⁷Question sheet and anonymized transcripts can be found at https://osf.io/9afnr/?view_only=93c8d0e2159847dba2c454519682b85d, accessed 24/04/2024.

⁸https://miro.com/app/board/uXjVMw38td8=?share_link_id=503687145119, accessed 24/04/2024.

learning goals especially referred to the study of “past and present design and technology [through which] they develop a critical understanding of its impact on daily life and the wider world” [29]. On this basis, we conducted four in-person and remote co-design sessions with the two year 4 teachers we contacted (see section 3.2.1) over a period of three months. These sessions initially centered on demonstrations of what various generative AI technologies could do, before moving to mapping out important concepts and potential topics that would fit with the above learning goal (see Figure 2). Given our overarching objective that students were to engage with an actual generative AI technology in order to develop practical and critical competencies, we identified multiple possible candidates for the latter—text-to-image (e.g., Dall-E), text-to-text (e.g., GPT3), image-to-text (e.g., CLIP) were all considered, as well as conversational interfaces (e.g., ChatGPT). We eventually settled on a text-to-image generation model as this seemed the most direct and literally generative way to engage students. The specific choice was Stable Diffusion [74], an image diffusion model that was and continues to be highly prominent. Like other text-to-image generation models, Stable Diffusion is capable of using natural text (along with other parameters, see 4.1.2) as an input for a training data guided process of diffusing noise into an image reflecting the text input. Based on the image creation process taking center stage, topics to introduce in the lesson plan were mostly concerned with creativity and ownership, but towards the end also touched upon more complex issues such as generative AI technologies’ transformation of reality (see 4.2.2).

4.1.2 Image Generation Tool. We were aware that in class each student could make use of a Google Chromebook. This meant that we could design custom websites which afforded a less technical interface for accessing a generative AI technology than most current solutions, and allowed for direct interaction with the parameters of the generation process. The students used various versions of the image generation tool (HTML/CSS/JavaScript frontend, Jupyter Notebook backend) for interfacing with our chosen image generation model, Stable Diffusion via the platform replicate.ai.⁹ In response to the lessons, there were four different interface design versions used over the six lessons (see Figure 3 for two examples). In the first week, students used a simple version with textual input only, before moving on to a more advanced version that also included technical settings. In the latter, the seed (i.e., the initial noise image for diffusion), the negative prompt (i.e., textual input that discourages generation of particular things) and the context-free guidance scale (i.e., numerical value that steers the model away from or towards more randomness in relation to the prompt) could also be set. Starting from week 2, this advanced version was used by the students. In week 5, the website confronted students with an ethical question to be interactively explored: whether it was OK to use actual artists’ ‘styles’ for image generation. To this end, we adapted the website to randomly select 5 previous generations and let students choose an artists’ ‘style’ (e.g., van Gogh or Winifred Knights) for a direct comparison between original and ‘adapted.’ To support interaction during delivery, we furthermore implemented a simple querying interface which allowed us to quickly display the image generations featuring a particular term and/or generation

⁹<https://replicate.com>, accessed 11/08/2023.

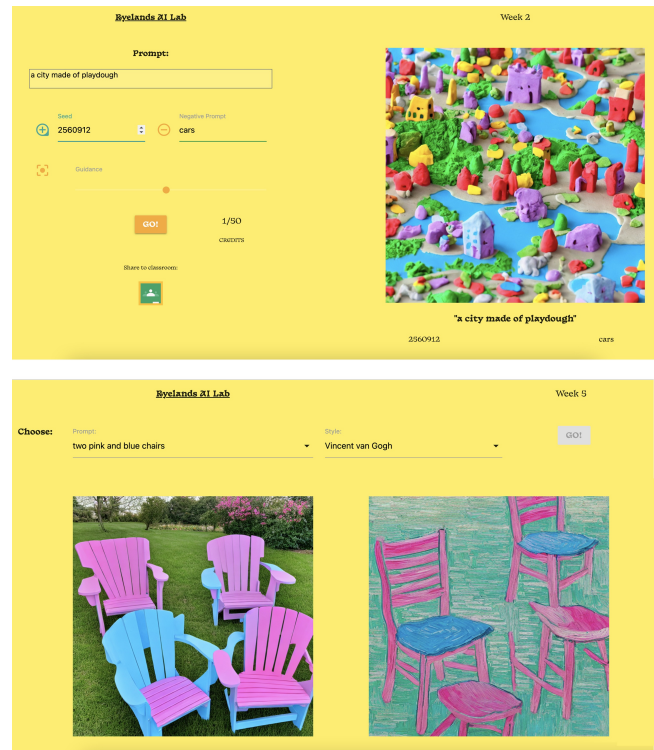


Figure 3: Two versions of the image generation tool used during delivery. Top: In this version used from week 2 onwards, students can input the main text prompt; and optionally specify negative prompt(s), an initial seed, and the context-free guidance scale, and submit to a custom Stable Diffusion API endpoint. **Bottom:** A version created for reflection on issues such as ownership and creativity: students select from 5 randomly sourced preexisting image generations and their prompts and can ‘convert’ these to a particular artist’s style.

date. This way, we could gather students after a period of activity for discussion and reflection. Here, it is also noteworthy that an internal discussion on safeguarding students took place in the development of these web interfaces, led by the concern around explicit content that can be created with image diffusion models. To preempt this, we eventually decided to always frontload a student-initiated API request with a ‘hidden’ negative prompt consisting of ‘gore NSFW creepy adult nudity horror erotic XXX weapon violence gun knife blood’. This was a first practical indication of critical concerns around introducing generative AI technologies to minors, which we reflect on in our discussion (see section 5.2.2).

4.1.3 Further Learning Materials. Other learning materials took shape in weekly iterations in correspondence with the design of lessons and image generation tool. There were slides and, most prominently, folders for which students would receive weekly worksheets (see Figure 4). The worksheets were generally employed to structure lessons by introducing students to basic concepts; for instance first describing an artwork, then using descriptions as a prompt to bring home the ‘back-to-front’ logic of using an image

Table 1: Overview of weekly delivery of lessons, with note on topics as well as materials and process.

Week	Topic	Learning Outcomes and Approach	Notes on Materials and Process
1	Introduction to AI image generation	It is possible to create photorealistic images based on text inputs, the chosen text impacts the resulting image.	Students filled in blanks in preexisting prompts with chosen words on worksheets, then were asked to enter these prompts into the image generation tool.
2	Refined prompt writing	Reinforcing basic prompt writing and learning how to use other parameters (e.g., diffusion seed, negative prompt) to help control image generation.	Students were asked to write their own prompts but were provided with some inspiration. They were also asked to experiment with advanced parameters and record their choices on worksheets.
3	Focusing on intention in prompt writing	Demonstrating and practising the use of descriptive language to translate specific intentions into prompts (e.g., a ‘blue door’ vs ‘a door’).	Introduced the notion of a group project (‘reimagining our school’) as a means to focus prompt generation around a specific topic.
4	Advanced vocabulary	Practising how to use specialised vocabulary in complex prompts to achieve deliberate outcomes relating to a specific theme (‘reimagining our school’).	The teachers did some ‘pre-teaching’ prior to the classes where students were introduced to advanced vocabulary.
5	Ownership and styles	Showing particular styles of image can be obtained from AI image generators and that using artist names can accomplish this. Using this demonstration to explain how training data impacts upon generated images and opening a discussion on training data ethics and ownership of outcomes.	In this week the image generation tool was used to re-generate images from previous weeks using the same prompts, seed and other parameters but with the addition of keywords to achieve specific styles. Worksheets were used to capture students’ thoughts on whether stylised images were preferable and on the ethical implications of creating images based on artist styles.
6	Retrospective, discussion, exploration of other types of AI	The final week was an opportunity to recap and assess previous weeks’ learning outcomes. We also demonstrated other types of generative AI to assess whether the students could transfer their knowledge to other domains.	The final lesson mainly consisted of demonstrations at the front of the class. A worksheet was used to allow students to reflect on one of their chosen images, describing that image’s attributes.

description to create a previously nonexistent image. Additionally, the worksheets were designed to let students keep track of their generations and decisions by noting prompts, parameter settings, etc. Aesthetically, the image generation tool and learning materials developed together through the initial co-design phase, with choices of font and colour for instance constrained by browser affordances.

4.2 Deployment: In-Person Delivery and Reflexive Adaptations


Here, we detail the actual *deployment* of our learning materials and scripts in the six lessons.

4.2.1 Lesson Overview. While the overall thematic and conceptual pedagogical approach remained largely unchanged, actual lesson scripts, slides, worksheets and iterations of our image generation tool were designed, developed or adapted on a weekly turnaround basis. While this led to substantial time demands, it also allowed us to more directly adapt to actual needs and challenges which could not have exhaustively or even accurately been formulated

prior to deployment. The final lesson plan is summarized in Table 1. Additionally, it is noteworthy that the lesson foci shifted significantly following the introduction of the lab project ‘Reimagining Ryelands’ in week 3, which we decided to pursue in response to the high quality of student engagement and to challenge students with a concrete goal relevant to their everyday lives. Subsequently, using the technical skills gained previously, the week 3-6 lessons gained a more constant and sustained focus.

4.2.2 In-Class Delivery. The two first authors delivered the six lessons along a timed but responsive lesson script according to each week, without specifying who spoke on what topic (see Figure 6). The modes of interaction generally proceeded from an interactive hands-up or call-out session, such as a game using the “Which Face is Real?” website,¹⁰ to an instructional presentation period followed by a related worksheet. This typically more didactically focused half of the lesson would then be followed by more or less guided individual image generation tool usage on students’ Chromebooks. The lessons were mostly concluded by reviewing recent image

¹⁰<https://www.whichfaceisreal.com/>, accessed 23/07/2023.



Week 1

Week 2

Week 3

Week 4

Week 5

Week 6

Figure 4: Exemplary scans of the worksheets for each week which students would keep in folders. Weeks 1 and 2 feature an introductory worksheet first (e.g., “Describe the artworks below” in week 1). Weeks 3 and 4, with the introduction of the “Reimagining Ryelands” lab project, first inspires students (e.g., by prompting them to list all things they can think of relating to their school) and then asks students to log their reimagination experiments. Week 5 first asks students to log their experiments with artist’ styles, and then provides a questionnaire for reflection. Week 6 features individual worksheets which show their chosen prospectus contribution and asks them to write a persuasive accompanying text.

generations, which brought about the development of the ‘listing’ interface for the tool (see section 4.1.2).

The year 4 teachers and teaching assistants remained present, and would support the delivery by, for instance, reformulating concepts in terms of previously encountered curriculum matter, or by focusing on a particular student according to emergent or ongoing needs. Particularly noteworthy is their activity outside

of the Ryelands AI Lab lessons between week 3 and 4. Here, a custom vocabulary list with architectural and aesthetic terms was pre-taught to the students to support more specific and intentional generations. This reflected the student’s architectural and aesthetic vocabulary deficiencies due to the Covid-19 pandemic, as well as their own desire for more ‘resources’ for image generation. Further, in week 6 we delved deeper into examples that dealt with the effects

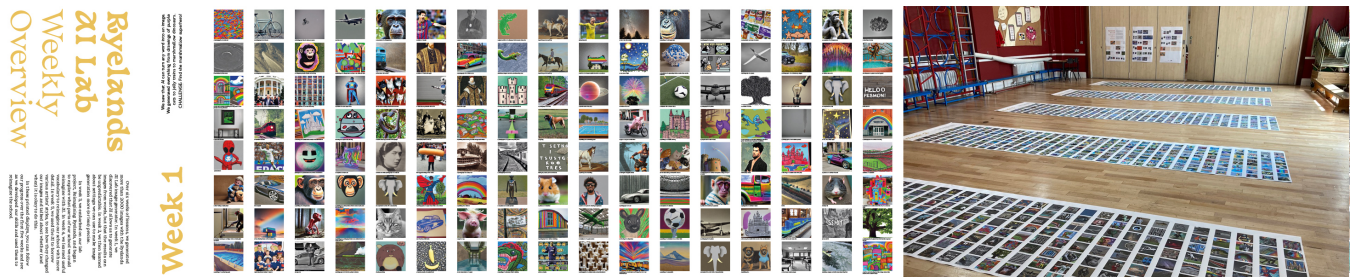


Figure 5: Left: A cropped print file of the 5x1 meter strip showing an introduction to the week’s delivery contents and a random selection of 250 images with their respective prompts. Right: Overview of the week strips laid out in the school gymnasium for the exhibition. In the background, three posters can be seen that show the spreads of the Ryelands AI Lab prospectus.



Figure 6: Image from in-class delivery during discussion of students’ creations in week 5, showing the ‘listing’ interface on the whiteboard and the week’s image generation tool (see Figure 3) on the students’ Chromebooks.



Figure 8: The output of the “Reimagining Ryelands” lab project were transformed into an Ryelands AI Lab prospectus and printed copies were handed out at the exhibition.



Figure 7: As part of the Ryelands AI Lab, students were given a lab project of “Reimagining Ryelands”. The above shows three such reimaginings concerning their school hall, cafeteria food, and library using a generative AI text-to-image model, with respective prompts below.

of generative AI technologies, such as Bogost’s article on wildfire skies turning gray in smartphone photographs [15]), and issues of representation (e.g., advanced ‘beauty’ filters on social media [13]).

4.3 Post-Deployment: Processing and Sharing Outputs

In this section, we describe the activities undertaken after the delivery of the six-lesson-plan which consisted of processed outputs that we used for celebrating and sharing the students’ achievements. It should be noted that the aim to hold an exhibition and to design an ‘imaginary’ school prospectus only took shape during delivery in response to the outputs and ongoing questions.

4.3.1 *Exhibition.* The exhibition was designed to give a near-complete overview of the work students had put in. For each week, we printed 5x1 meter strips showing 256 randomly chosen images and their prompts for each of the five weeks in which image generation was center stage (see Figure 5). We created a version of the image generation tool which incorporated all interactive elements (i.e., prompts, negative prompts, seeds, guidance-scale, artist styles) for the exhibition and set up interactive stations where students could show people with parental duties, siblings or other students their newly gained skills.

4.3.2 Prospectus. The design of the prospectus (see Figure 8) was inspired by existing materials found at other schools, such as info-brochures for parents/guardians as to what to expect from the school. In week 5, we asked students to pick one image to be included in this prospectus from among those created for the lab project ‘Reimagining Ryelands’ from week 4 (see Figure 7). In the following, final week, we further asked students to provide brief, persuasive descriptions of these images for the purpose of the prospectus. This material was then copy-edited into the Ryelands AI Lab prospectus, including an introduction by the two lead authors giving an overview of the project. The prospectus further included school group advertisements that integrated some of the common, unusual and amusing themes found in students’ images (e.g., things made from marshmallows).



Figure 9: Example of an image brought up for discussion by a student in the week 5 lesson. Note the extra hand circled on the far right, and the empty sleeve circled to the left of the extra hand.

5 STUDY FINDINGS AND GENERATIVE AI EDUCATION PROPOSITIONS

In this section, we gather the findings generated within and through the exploratory Ryelands AI Lab study and close by synthesizing them in the form of “strong concepts” [41] (see section 3.2.5) that HCI can draw from as propositions related to the generative AI education context. Note that all findings are scaffolded by reflections of the teachers (T1, T2) gathered from the post-deployment semi-structured interviews, and that we have included descriptive elements to accurately represent the reflective interplay between a priori methodological decision-making and actual deployment of this project.

5.1 Reviewing Constructionist Attributes of the Ryelands AI Lab

Here, we consider a central constructionist aspect of our exploratory curriculum, and propose guidance on how HCI research in the educational context of generative AI technologies may build on our exploratory study.

5.1.1 Findings on the Co-Development of Practical and Critical Competencies. The major constructionist attribute of the curriculum is that we observed practical and critical competencies developing not distinctly, but rather exhibiting an interrelationship. Essentially, as practical know-how grew through continuous and diverse interactive sessions, students became more capable of reflecting on and even independently identifying critical aspects of their activity and the involved image diffusion technology. Over the first five weeks in the school the students created more than 2,000 images using corresponding prompts and parameters (week 6 was about reflection and discussion, so involved no image generation). To illustrate the diversity of the images the students created, Figure 11 shows randomly selected images from weeks 2 to 4. Students’ choices of prompts in the first two weeks seemed either random or unlikely things from everyday life or pop culture (e.g., Harry Potter, TikTok memes, footballers); providing a baseline for our observation of competency development.

Following the introduction of the ‘Reimagining Ryelands’ lab project in the third week, the image generations tended towards more deliberate choices. At the same time, this is not a clear-cut change as is shown by the three lower images in the week 3 column (see again Figure 11) which do not directly seem to relate to the goal of re-imagining the primary school. However, it can also not be ruled out that these were part of students’ figuring out how to use the parameters (i.e., negative prompts, guidance scale, seeds) introduced in week 2, and thereby reflect the ‘model-probing’ nature of constructionism—as students explored how far and where they can push a particular technology, they gain competency with the latter which in turn opens up more opportunities for intentional and reflective engagement. At this stage of the curriculum, the development of practical competencies (in terms of achieving desired, task-oriented results) indicated latent critical competencies in the probing interactions the former also supported.

In week 5, students used that week’s image generation tool version (see 3, bottom) to re-generate previously generated images in particular artists’ styles. This led to some students recognizing their previous prompts and noting the differences and distinct appearance of the artists’ styles, which led to conversation among authors and students why that may be—particularly in contrast to some of their previous attempts to generate things familiar to them (e.g., their teacher but flying a plane). This allowed for a consideration of the celebrity of artists leading to many pixels of their work being on the internet and that these, unlike clearly identifiable portraits of *their* teacher, were likely in the training data for the image diffusion model we were using. Most indicative is an example where a student excitedly asked to show their latest generation from this interaction on the class screen. The student showed the image in Figure 9, where a superfluous hand appeared after the re-render in the style of Van Gogh. This led to a conversation in class on why the AI technology would show this hand, which, in combination with




Figure 10: Share distribution of student answers for the Week 5 questionnaire.

the previous conversation on artists' styles, led to discussing that AI technologies aggregate pixels rather than actually know what they generate. We argue that this is important to note: the student's growing practical competency with a particular image diffusion model made it possible to see the creation of images as more than 'mindblowing,' but rather to scrutinize and question what appeared in front of them, which in turn opened opportunities for shared critical reflection on generative AI technologies. In other words, the familiarity and ease with which students then engaged the image generation tool made space for the strengthening of critical competencies such as observation and reflection.

After this engagement, we shared the story of artists' disgruntlement and anger with generative AI technologies.¹¹ Then, we asked students (n=48 on the day) a series of questions on a worksheet, with two relating to ownership ("Is it OK to generate images based on artists' styles?" and "Which of the following [Prompt | Settings | Generated Image | Training Data] do you feel are 'yours'?") and one relating to creativity ("Is using an AI image generator creative in the same way as doing a painting?"). The results of these questions are visualized in Figure 10.

The responses for all three questions were ambiguous, but show consideration and reflection in this ambiguity. As is shown in the figure, a combined majority of students (n=33) stated that it was okay to use artists' styles, however of these a slight majority (n=17) further qualified this statement by saying it was only okay given a particular condition, such as the artist giving their consent or if they were dead. More idiosyncratic answers also argued that artists might like someone making pictures in their style. However, the students answering that it wasn't OK to use artists' styles were much less ambiguous when expanding on this selection, arguing that it would be a threat to their livelihood and theft. The results for the question whether generating images with an AI technology was as creative as doing a painting was equally ambiguous, with a slight majority (n=26) saying it wasn't. However, here again some answers were unexpected and not clear-cut, for instance some students said it was not as creative because they themselves had used the image generator, implying that they did not consider themselves creative.

¹¹Specifically, the case discussed here: <https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/>, accessed 25/07/2023.

Concerning ownership, a slightly larger share of students (n=18) expressed feelings of ownership regarding the prompt, which "came from my head." There was some overlap with feelings of ownership of the generated image (n=16), with some students (n=7) selecting both. The ambiguity (and indeed idiosyncrasies) in these responses indicate that students' competencies were not only repeated following instruction, but developed through the actual engagement with the technology at hand.

In the final week 6, we shared further controversial aspects of generative AI technologies as mentioned above (see section 4.2.2), but furthermore also acknowledged that we built in negative prompts to the image generation tool to keep the students, in our minds, safe from explicit imagery. This led to a different kind of conversation than we anticipated, with students and teachers reflecting on the similarities to search engines and the fact that, verbatim, if you search for upsetting things you will get upset; but also allowed us to point out the differences in that the image diffusion model could have generated upsetting things 'unprompted.'

T2 describes the observed development towards a more reflective stance as a shift from "the primary view [which is] basic systems stuff like 'how do I make you change color and how do I make it bigger?' [towards getting] into the nuts and bolts of kind of like ownership and art and creativity" (T2). They attribute this to the prolonged, hands-on approach of the project, stating that the students "were able to find different layers to it with every single session, with every single session, their level of understanding of it grew" (T2).

In sum, these findings indicate (1) that practical and critical competencies of generative AI technologies cannot be clearly disentangled but, as per constructionism, co-develop; and (2) that the critical competency component of AI literacy, which may seem as the most challenging to convey to young people, can be significantly scaffolded by practical know-how. Building on this insight, we lay out our first proposition for HCI research to pursue constructionist generative AI education.

5.1.2 Proposition I: Constructionist Curricula for Co-Developing Practical and Critical Engagement with Generative AI Technologies. We propose that the observed co-development of practical and








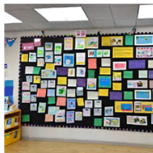

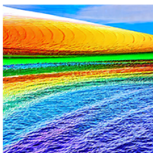
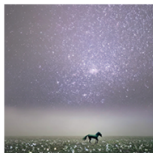




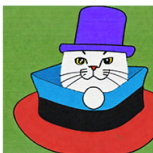


Week 2 (parameters introduced)	Week 3 (lab project introduced)	Week 4 (new vocabulary introduced)
 <p>a rainbow flying in space</p>	 <p>a maths lesson in a red classroom</p>	 <p>a wall with arches and with faces which is unlimited</p>
 <p>a harry potter and gin from hogwarts geating married</p>	 <p>a chair that can massage a neon green hand while flying</p>	 <p>blue pizza food</p>
 <p>a pinapple chacinga train upsidedown on a train track</p>	 <p>a really big display for classroom</p>	 <p>a futuristic school uniform with a red lions face</p>
 <p>a rainbow water</p>	 <p>A photo of a gloomy glittery sky with a uni-corn flying in it</p>	 <p>modern bright coloured drawer. Wide-angle shot.</p>
 <p>puple blue rainbow</p>	 <p>windows on a car non see through</p>	 <p>modern organic school playground made of birch wood</p>
 <p>cat in hat</p>	 <p>the blue sea galaxy</p>	 <p>a modern library that is huge</p>

Figure 11: Overview of exemplary, randomly selected images and their respective prompts from weeks 2 to 4. Note that even though the lab project was introduced in week 3, the lower three images show unrelated, more exploratory prompts.

critical competencies may encourage HCI researchers to utilise constructionist approaches to education in relation to generative AI technologies. Specifically, the instantiation of our exploratory study phases (see Figure 1) show that, as highlighted in other work, in-depth and prolonged hands-on engagement with generative AI technologies is needed to really pursue the development of competencies. For instance, the referenced work by Williams and colleagues notes that time and exploration leads to ‘better’ performance [92], which are factors yet to be explored in the emerging approaches to generative AI education (e.g., the “Prompt aloud!” project by Lee and colleagues [54] lasting for a single 90 minute lesson). Accordingly, we propose that the responsive study structure and our

provided reasoning can assist HCI researchers and practitioners in various educational contexts (e.g., students in primary, secondary or tertiary education) that may on build constructionist-style education approaches such as “experiential,” “active” or “inquiry-based learning” (see [49, 52, 95], respectively).

First, during an *initial research and (co-)design phase* prior to deployment, a shared baseline understanding with teachers can be established. The goal of this should not be to establish a concrete specification for what would happen next (e.g., fixed lesson plans or designs for the image generation tool) but rather to understand more deeply the underlying challenges relating to the unique elements of teaching generative AI competencies and how

those challenges relate to the wider curriculum, pedagogic strategies used, and experiences with technology that students make in their everyday lives. It is noteworthy that while the teachers and ourselves were surprised by the quality and high degree of engagement of students during the second phase, this can also be considered more proactively to feed into the initial phases of future work. On reflection, we posit that despite their age, the students have grown up in a world dominated by the ‘reality-on-demand’ logic of social media and streaming services. That a technology can serve up particular and strange images at which is arguably a norm of this generation inasmuch as search engines—yet the particular differences to the latter should equally be stressed. It is therefore important for HCI research to more fully and precisely consider the competencies students can *already* build on. Accordingly, future work may choose to intensify the participatory involvement of teachers in this phase (as Williamson argues, see [93]) or by also involving students or people with parental duties to discern how ‘baseline’ competencies relate to generative AI technologies.

Second, the actual *in-person delivery and reflexive adaptations* of our exploratory study can be drawn from. As illustrated in Figure 1, the weekly lessons essentially formed responsive, interrelated components that fed into each other as the delivery progressed. We argue that the co-development of competencies was especially fostered in this regard. Facilitated by regular interactions directly with the image generation tool, the students seemed to have ‘small revelations’ (e.g., if something appears in the image they don’t like, negative prompts may remove it) which they would then experiment with further to find the limits of this new perspective (e.g., if the image generated from the seed just isn’t a good one, then no amount of negative prompting will help). The example of the superfluous hand spotted and made into a topic of class discussion by a student above is a case in point: by that point in the curriculum, basic practical competencies were so far established that the creation of an image in the style of Van Gogh (in this case) was not the all-consuming focus, and therefore space for a further ‘small revelation’ was made. In response, the authors could then adapt the content and form of the learning materials of the subsequent lesson accordingly. It is in this interplay that HCI researchers and practitioners can also identify where precisely to include more generalizable principles such as the relationship between training data, input prompts, and outputs that hold for generative AI technologies. If structured correctly, this has potential for providing the ‘sticky learning’ that the teachers in our study referred to (see section 5.2.1).

We therefore suggest that responsivity and longitude are key attributes of a constructionist curriculum on generative AI technology, and that this may in turn actionably contribute to more expansive understandings of AI literacy—not as a single discrete set of knowledge to be transferred (see [60]), but rather made of competencies to constantly develop and hone through prolonged, playful and ‘revelatory’ engagement.

Third, the final phase of *processing and sharing outputs* of the exploratory study was key to both explicate as well as celebrate the ‘*flow of the work*’ with our participants and related stakeholders; here students, siblings, and people with parental duties. To us, the significance of this phase lies in the framing of the students’ achievements as core competencies *for life*, rather than a one-off

school art project with individual achievements. By giving students’ work a concrete form (i.e., the imaginary prospectus built on the ‘Reimagining Ryelands’ lab project) and a venue (i.e., the handing over of a prospectus copy for each student in an exhibition setting), another setting was provided in which students could reflect on and share the competencies they have built up; particularly with those in their lives who may take a longer view on how the developed competencies could help them navigate everyday life. This brings us to a final argument for the significance of HCI pursuing constructionist curricula on generative AI technologies. Seeing the latter as a contemporary “foundational technology” [17] that has already permeated much of the everyday activities of children (see [71]), practical and critical competencies will be required of current and future generations as a prerequisite to twenty-first century citizenship. While we do not claim that this exploratory study has actually achieved this, we hope to have contributed to the emerging field that seeks to develop pedagogical approaches which consider the challenge of generative AI education in depth.

5.2 Findings on Responsivity in the Curriculum Design Process

Here, we stress the importance of responsivity in the design process of a constructionist curriculum on generative AI technologies to the particular educational context. This is shown by providing space to consider the teachers’ perspectives on the Ryelands AI Lab, and their reactions to their own and their students’ increasing competencies with generative AI technologies.

5.2.1 Perspective of Teachers. The teachers reported that the Ryelands AI Lab project had a significant impact on them and their perspectives on teaching about technology on three levels. First, concerning their own knowledge of and attitudes towards AI technologies, they previously regarded AI in terms of how “people are worried about it” (T1) and as “this kind of big, scary monster like skirting around the periphery of everything” (T2). Early on, the teachers still felt that “my mind was blown just by the very concept of what [generative AI technologies can do]” (T2), whereas by the end T1 stated that “I feel more confident to talk about [AI technologies], I can lead those conversations now” (T1). Teachers delivering this type of curriculum intervention themselves is not as unlikely as may be thought. The realities of being a primary school teacher were evocatively described by T2: “I am a 40-year-old man from Wigan. Is my French accent amazing? No. Will I teach French? Yes.” (T2). This refers to the fact that especially primary school teachers need to gain confidence and expertise in subject matters which they do not have extensive knowledge of themselves. Both teachers referred to a whole industry that has sprung up to address this circumstance, mentioning “existing models from, like, music companies [...] that literally give a whole package for teachers that are: ‘This is the prior knowledge you need to have’” (T2). The Ryelands AI Lab was seen as having potential to be adapted in a similar way, because “the skeleton was able to be adapted so much and also the sharing of ideas and things enabled it to be really kind of a fluid approach and, and a really creative approach” (T1).

Second, from a pedagogical perspective, T2 highlighted the constructionist element of the curriculum, noting that it was “really rare that we get to be in a situation where we know just as little as

the kids do when we start off” (T2). This led to them being “able to [...] really lean into it with the kids, and their questions were our questions” (T2). T2 further reflected that “[this was] probably my first real experience sort of [this kind of mutual, constructionist learning] (T2). Concerning what it was that made this curriculum intervention successful, the teachers couldn’t draw out one thing: “It’s like the demystification of the process, it’s about the fact that it was a shared experience, it is about the fact that [...] you brought the kids in to be experts right from the beginning” (T2). Akin to the description of the students’ progress, they saw the intervention resulting in “[having] the [actual technical] process of it demystified” (T2); now thinking how “AI is a bit of like an amalgamation of stuff that’s already out there, which I didn’t know before” (T1).

Third, the teachers began re-evaluating how AI and other technologies can be taught and to what benefit for their students. Concerning a more general media literacy encouraged by the curriculum, T1 stated that “we’ve had a lot more in-depth conversations about art” (T1) following the project. To them, the importance cannot be overstated, as “we want them to look at things around them and not just take them at face value to actually question them because there is such a mental health issue with young people” (T1), referring to social media and image filters specifically. This indicates that *through* literacy gains surrounding AI technologies, other aspects of media (in this case, images) can become targets for more reflective engagement as well.

In sum, the impact on the teachers reported here is indicative that a constructionist, hands-on engagement with generative AI technologies is not only valuable within its own lifetime and solely for students, but also seemed valuable to the teachers, helping them to reflect on their teaching strategies. While the above quotes are essentially exclusively positive with regards to the developed curriculum, we use them as an impulse for our proposition that . Further below, parts of this finding also inform limitations we see for future work (see section 6.3).

5.2.2 Proposition II: Approaching Generative AI Technology Education with Critical Responsivity. Above, we documented how the teachers framed the responsivity of the curriculum ‘skeleton’ and the gradually more complex hands-on engagement as an overwhelmingly positive experience. However, from a more distanced perspective and regarding our experience as researchers, we here reflect on the opportunity for more critical engagement in design that we and other HCI researchers may need to pursue. As mentioned above, generative AI technologies in educational contexts are frequently seen through either a punitive (e.g., preventing students from cheating) or utilitarian (e.g., improving teacher efficiency) lens. Our practical experience with the Ryelands AI Lab has brought further concerns to the fore, which connect the generative AI education space to more critical scholarship such as calls for a “decolonial approach to AI in higher education teaching and learning” [96]. In short, aside from the above argument for intertwining the development of practical and critical competencies *within* a constructionist curriculum, we propose that the design of the latter calls for ‘critical responsivity.’ By this we mean that researchers engage in an ongoing noticing of the emerging concerns of stakeholders which crucially could not have preceded the actual interaction with a generative AI technology.

The difficulty in critically negotiating the capacities of generative AI technologies in an educational context came to the fore for us swiftly in the form of our ‘frontloading’ of negative prompts (*‘gore NSFW creepy adult nudity horror erotic XXX weapon violence gun knife blood’*). While this ‘need’ was shaped by our particular participants, it is noteworthy that it was *our* need—whereas the teachers were at ease with considering it similar to the use of search engines by students. At the same time, this could have provided us with an opportunity to stress in which ways this ease may need to be tempered, given the difference between targeted search and the barely noticeable “pattern leakage” [7] that may crop up in AI technologies due to their probabilistic abstracting from training data. That we did not act on this opportunity is, in hindsight, an indicator where ‘critical responsivity’ could have informed our design of subsequent lesson materials to act on this faultline. Similarly, a limitation of the current set of learning materials (particularly, the more instructive slides) is that more ‘sharp end’ examples are needed to represent and educate on harms brought about by AI technologies not dissimilar from the kind of technology the students were using. While there is potential for more in-depth discussions to arise in a possible multi-year programme (e.g., a year 4 module later supplemented by an advanced year 6 module) as advocated by the teachers,¹² there were multiple opportunities for a more attuned ad hoc response. For instance, this could have been to stress that we were using an ‘off the shelf’ image diffusion model (i.e., non-finetuned Stable Diffusion); and discuss with students in which contexts such generic solutions would not work—then extending it considering harms of generic models in arguably higher stakes contexts such as facial recognition, criminal recidivism prediction or credit scoring. This echoes concerns from critical voices such as Nemorin and colleagues who have noted that often little attention is paid to how AI technologies in educational contexts may reify knowledge along extractivist lines [65].

In sum, the ambition for “lifelong learning” [27] regarding AI technologies needs to bring values beyond the purely economic to the forefront. We argue that in the long term, constructionist HCI research efforts may help in addressing this need, but the intricate nature of critical concerns and rapid changes also necessitate action in current design efforts. In practice, and reflecting on our disclosure of the frontloaded negative prompt and the subsequent exchange with students, we suggest that HCI research could be proactive and responsive here by adopting a practice of *disclosure*, rather than optimization. A parallel can be drawn here to the field of “explainable AI” (XAI), particularly regarding how Benjamin and colleagues’ have stressed that while machine learning uncertainty is usually explained or engineered away, it can serve as a design material in its own right when being actively engaged [7]. In other words: in addition to putting ‘safety’ of stakeholders first (e.g., by adding locked negative prompts to an image diffusion model), the disclosure of having done so (i.e., of having *further* biased a biased model) can function as a way to critically ground reflection

¹²This would also allow for an introduction of more intricate and challenging aspects of AI technologies and society, such as shown in the works of, for example, Bender and colleagues [6] or Birhane and colleagues [12]. Such work would help to further underline the economic motivations behind much of AI technological development (e.g., the human element of training data production, extraction, curation and/or moderation), and also distinguish educational from utilitarian efforts aimed at ‘efficiency’ such as learning how to write prompts in order to capitalize on generated images.

on the capacities and dangers of generative AI technologies. The challenge and potential remedy, we propose, in pursuing this lies in researchers—or educators—designing generative AI technology education provisions with critical responsivity in mind.

6 DISCUSSION

In this section, we present a discussion relating to broader implications for HCI researchers regarding firstly the design of generative AI systems, and secondly the methodological strength of RtD to respond swiftly to disruptive technologies.

6.1 Constructionist Strategies in the Design of Generative AI Systems

While the Ryelands AI Lab carries direct implications for constructionist AI education, there is also intermediate knowledge we can build upon in this regard for more general design concerns. Here, we consider two primary aspects of designing generative AI systems where constructionist strategies may support HCI: first, during development as a way to build understanding and requirements of a system; and second, as guidance offered in the design of artefacts (e.g., onboarding user interfaces). Both aspects connect to the larger discourse on XAI or ‘interpretable machine learning (ML),’ but importantly weigh the importance of explanation and interpretation differently. Similarly to the difference of constructionist to ‘instructionist’ (e.g., frontal lecturing, passive students) learning methods, the goal of designing systems using constructionist strategies does not lie in transferring one definitive explanation or interpretation but rather in providing opportunities for engagement; whether this be playful or serious. To clarify the below, we consider constructionist strategies as ways to foreground particular *affordances* of a technological artefact; i.e. the relations through which the world and the ways in which to perceive and act in it takes shape [42, 67].

First, constructionist strategies may benefit both designers and stakeholders in the early phases of system development, such as in participatory co-design methods or place-based inquiries. Already in and of themselves, these types of methods are not generally present in XAI research, which predominantly engages with expert audiences, specifically people with formal ML education (e.g., [57, 61]). Following proposals for the field to shift focus (cf. [24, 90]), work in XAI has begun to diversify its methods and audiences. For instance, Benjamin and colleagues have conducted co-design workshops with employees at a research institution where the goal was not to convince stakeholders of the suitability of a proposed ML-driven visualization system, but rather to understand how precisely explanations for that system’s outputs shape the highly specific contextual understanding of the stakeholders, with participants literally constructing representations of their context from materials such as playdough, which in turn led their analysis for the design of their system (see [9]). Such constructionist strategies therefore further stand in contrast to the generally ‘instructionist’ stance of XAI even when non-experts are considered, where stakeholders are often left unclear about how generic explanation metrics such as fairness map to their specific context (see [76]). In this vein, we argue that constructionist strategies which closely couple opportunities for practical and critical engagement

may further assist HCI in developing and pursuing such context-aware methods.

Second, constructionist strategies can also serve in the design of system *artefacts* such as onboarding interfaces that can help stakeholders comprehend the particular generative AI technology’s capacities rather than disguising or obfuscating them. Such HCI work is becoming more pressing, for instance to contextualize and weigh the outputs of text synthesis models such as ChatGPT that are particularly prone to “fabrications and falsifications” [25]. However, it is unlikely that there will ever be a perfectly trustworthy form of presenting current generative AI technology outputs—when developers themselves can frequently not fully account for their systems and/or routinely misjudge and over-accept explanations for AI technologies [47, 83]. Thus, we suggest that constructionist strategies for generative AI systems offers a plausible route forward: rather relying on explanatory elements which puts people interacting with a generative AI system into a passive role, being given the means (e.g., through interactive examples) to playfully work through generative AI technologies’ capacities within a given system may empower stakeholders. For instance, in contrast to the highly specialized explanatory interfaces dominant in the field (cf. [39]), a constructionist design strategy may pursue incremental increases in the complexity of an interface (e.g., by introducing parameters such as seeds or guidance along specified routes). This could firstly support the practical understanding of a generative AI technologies’ capacities, and secondly scaffold the multitude of “explanation styles” [11] that people bring with them by allowing various thresholds for reflection. In this light, we argue that design can learn from educational approaches more generally to design for playful, ‘tinkering’ engagements, and note that importantly such ‘ludic’ design of generative AI systems can build on a long research tradition tracing back to Gaver and colleagues’ *Drift Table* [33] and contexts such as playful encounters with robots [55].

6.2 Research-through-Design as a “Rapid Response Methodology”

Our main methodological interest was the effects of choosing RtD as our guiding project methodology for engaging generative AI technologies in this context. The relative openness of the topic of generative AI technologies in terms of its novelty and connections to all kinds of subject areas in the UK curriculum (e.g., art, citizenship, design, engineering) led to a burgeoning of design possibilities. During the initial collaborative process, this led us to envision a whole range of generative AI technologies (particularly text models such as ChatGPT) to be included in the delivery. Up until the first delivery, this was seen as feasible and the first lesson plans as well as scripts included details on how the various technologies would tie into each other. For instance, in one version the students were to create an avatar of some sort (e.g., a superhero) using a text-to-image model, then write scenes for the character using a text prediction model, which then would be used for further image generation. Additionally, we also considered making a hardware component; specifically a micro:bit¹³-based camera which would leverage Stable Diffusion’s image-to-image mode for an even more

¹³<https://microbit.org/>, accessed 01/08/2023.

direct (and less desk-based) hands-on engagement with the AI technology based on Benjamin and colleagues' *Entoptic Field Camera* [8]—with us thinking that these *artefacts* would be the focus of the RtD methodology. As stated repeatedly above, these initial plans did not 'survive first contact' with the realities of introducing fundamental concepts alongside the 'mind-blowing' effect of the sheer possibility to create never-seen-before images from words. While we, therefore, did not achieve the same scope of topics and technologies we had initially sought out, we could rapidly adapt to the flow of the interventions rather than having to stick to a pre-given and potentially less impactful plan. Further, the RtD methodology we followed allowed for a previously unplanned introduction of a lab project ('Reimagining Ryelands') and the subsequent creation of the prospectus and exhibition.

This experience now leads us to consider the choice of RtD as our project methodology. The 'inward' responsivity of RtD projects and related design research methodologies is well known, and has been extensively detailed above in terms of weekly adaptations of learning materials, reshaped project goals, etc. Here, we reflect briefly on the *outward* responsivity of RtD. To put this into perspective: the text-to-image diffusion model we used was released only months prior (August 2022¹⁴) to our first meeting with teachers (November 2022). We find this noteworthy due to the vast space of concerns that was opened up in short succession to the release of it and similar models (e.g., ChatGPT, Midjourney, etc.)—as highlighted in the Writers' Guild of America strikes [1], or a UK House of Commons committee report [18]. We argue that it was exactly this rapidly unfolding and ongoing event which made it feasible for us to choose RtD as a methodology. That is not to say that the widespread integration and proliferation of generative AI technologies won't require large scale studies—however, given the limitless contexts in which generative AI technologies as *software* artefacts can be applied, HCI also needs to be able to quickly set up, deliver, and understand small-scale and single context studies.

We argue that RtD is eminently suited to study the impacts of disruptive technologies while their socio-cultural-technical ramifications are still being negotiated. This is likely particularly true with AI technologies, where for instance the gap between user-facing apparent 'intelligence' and vast socio-technical realities (e.g., content moderation, bias, data sourcing) can be particularly pronounced. At the same time, any technological innovation or adaption brings about uncertainty regarding its socio-cultural ramifications. This fits into a strong tradition within RtD and HCI design research where ambiguity [32] or multiple meanings [81] are seen as resources rather than obstacles. Yet, we also think that the *outward* responsivity of RtD is perhaps obvious, but underspecified. In turn, we propose to term this attribute of RtD as its suitability for being a *rapid response methodology*, a term we borrow from museum studies. "Rapid response collecting" was initiated in 2013 as a curatorial strategy by The Victoria and Albert Museum (see [26, 88]) to deal with significant quickly unfolding events in an increasingly connected world saturated by scaling information technologies—for instance, 3D-printed guns.¹⁵ In contrast to rapid response collecting, where

museological ordering and contextualizing takes precedence, it is the strength of RtD to synthesize intermediate knowledge from the materials it gathers that is not only reflective but *generative* towards new theories, products, design strategies, or subsequent research. But in contrast to other types of design processes (e.g., the double-diamond framework), RtD here can also be seen as responding to the questions thrown up by *historical* phenomena rather than the demands of a product for particular ends (e.g., efficiency, UX, profit). This becomes all the more pressing given the general difficulty in preventing negative outcomes of technology once these have been identified (often termed the 'Collingridge dilemma,' see [51]), and specifically the already significant outpacing of regulatory efforts by releases of AI technologies (see [77]). Accordingly, we argue that seeing RtD as a rapid response methodology can lend further legitimacy to the studies and researchers subscribing to it while still maintaining its specificity as a research methodology first and foremost.

6.3 Limitations and Future Work

One of the main limitations of this project is its relatively small sample—one exploratory study in one school with one specific set of teachers and students. As such, it cannot be ruled out that parts of our observations came about due to the "novelty effect" (see [2, 63]) of generative AI technologies for students and teachers. While we think that it is exactly this novelty which needs to be addressed, and that RtD is in a position to do so (see section 6.2 above), it remains a concern within the context of generative AI education approaches. This effect may be alleviated with further studies, for instance conducting a more complex intervention with the same year group of students at a later date to assess the sustainability of our approach; or with studies that are explicitly designed for triangulation to test the approach for scalability. Regarding the latter specifically, and with encouragement from the teachers, we are planning to extend the approach to other schools in the area and other age groups (e.g., secondary and high schools). Given that one of the main learning artefacts, the image generation tool, can be accessed very easily by any device capable of web browsing, such efforts could focus on *scale* primarily—there are as many as 16,783 primary schools in England alone.¹⁶ However, there are multiple design challenges such future work needs to navigate, especially given that it would mean pivoting from our responsive approach towards building the kind of package which are provided by platforms to rapidly upskill teachers on particular subjects, such as Twinkl.¹⁷ To frame this challenge productively, we are considering to frame it through the established HCI lens of "research products" [68] developed by Odom and colleagues and recently expanded precisely in the dimension of "scale" [16] by Boucher, as this clarifies the particular demands of scaling the Ryelands AI Lab by articulating a design space with constraints. This would also allow for further exploration of a specific RtD concept as a way to rapidly respond to emergent *design* challenges of disruptive technologies. However, a packaged module could also easily become just another generic resource commodity pushed on overworked and underpaid educators without enough

¹⁴<https://stability.ai/news/stable-diffusion-announcement>, accessed 24/11/2023.

¹⁵More recently, rapid response collecting resurfaced during Covid-19 pandemic focusing on the related paraphernalia and side-effect of widespread shifts to online socio-cultural activities [20].

¹⁶<https://explore-education-statistics.service.gov.uk/data-tables/permalink/74bb0ee9-712c-4820-ad0e-08dbb395de42>, accessed 14/09/2023.

¹⁷<https://www.twinkl.co.uk/>, accessed 03/08/2023.

critical perspectives or representative examples; which would have to be carefully considered and critically reflected upon (see also [35]).

7 CONCLUSION

In this paper we detailed how we used a Research-through-Design project methodology to engage the uncertainties surrounding generative AI technologies with an exploratory study in the educational context; which took the shape of the Ryelands AI Lab as a constructionist pilot curriculum for generative AI primary education. Adopting RtD's capacity for inquisitive and reflexive development, we designed, produced, and delivered six lessons on generative AI technologies which appears to demonstrate improvements in practical and critical competencies for students; and drew intermediate knowledge from our observations for future work in HCI. Concerning the project's context directly, we (1) offer guidance for HCI research on developing constructionist generative AI curricula, and (2) critically reflect on the role of AI technologies in education based on our experience. Reflecting on wider implications for the field, we further (3) consider the value of constructionist strategies in designing generative AI systems and (4) look at RtD as a 'rapid response methodology' that is particularly suited to unfolding and unsettled socio-technical developments surrounding disruptive technologies. Especially this latter aspect shows, we argue, why design can attend to emerging uncertainties in unique and valuable ways.

ACKNOWLEDGMENTS

We cannot thank the amazing children of classes 4SB and 4G enough for their enthusiasm, curiosity and wild imaginations. We also want to thank headteacher Mrs Linda Pye for her support and encouragement in making this project happen, as well as Willow Mitchell and Joe Bourne for their invaluable input. This work is supported by UK Research and Innovation (grant MR/T019220/1, "Design Research Works").

REFERENCES

- [1] 2023. WGA Negotiations Status as of May 1, 2023. <https://www.wgacontract2023.org/the-campaign/wga-negotiations-status-as-of-5-1-2023>
- [2] Ritu Agarwal and Elena Karahanna. 2000. Time Flies When You're Having Fun: Cognitive Absorption and Beliefs about Information Technology Usage. *MIS Quarterly* 24, 4 (2000), 665–694. <https://doi.org/10.2307/3250951> Publisher: Management Information Systems Research Center, University of Minnesota.
- [3] Mariana Aki Tamashiro. 2021. How do we teach Emerging Technologies in K-9 Education? Using design fiction and constructionist approaches to support the understanding of emerging technologies' societal implications in formal K-9 education. In *Proceedings of the 20th Annual ACM Interaction Design and Children Conference (IDC '21)*. Association for Computing Machinery, New York, NY, USA, 637–640. <https://doi.org/10.1145/3459990.3463402>
- [4] Heather Banchi and Randy Bell. 2008. The many levels of inquiry. *Science and children* 46, 2 (2008), 26. Publisher: National Science Teachers Association.
- [5] Jeffrey Bardzell, Shaowen Bardzell, and Lone Koefoed Hansen. 2015. Immodest Proposals: Research Through Design and Knowledge. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2093–2102. <https://doi.org/10.1145/2702123.2702400>
- [6] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [7] 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>
- [8] 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>
- [9] 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>
- [10] Karl-Emil Kjør Bilstrup, Magnus Hoholt Kaspersen, Marie-Louise Stisen Kjerstein Sørensen, and Marianne Graves Petersen. 2022. Opportunities and Challenges of Teaching Machine Learning as a Design Material with the micro:bit. In *Adjunct Proceedings of the 2022 Nordic Human-Computer Interaction Conference (NordCHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3547522.3547689>
- [11] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. 'It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3173951>
- [12] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. <https://doi.org/10.48550/arXiv.2110.01963> arXiv:2110.01963 [cs].
- [13] Charlotte Bitmead. 2023. The problem with TikTok's 'Bold Glamour' filter. <https://www.cosmopolitan.com/uk/beauty-hair/beauty-trends/a43177683/tiktok-bold-glamour-filter/> Section: Beauty Trends.
- [14] Paulo Blikstein. 2013. Gears of our childhood: constructionist toolkits, robotics, and physical computing, past and future. In *Proceedings of the 12th International Conference on Interaction Design and Children (IDC '13)*. Association for Computing Machinery, New York, NY, USA, 173–182. <https://doi.org/10.1145/2485760.2485786>
- [15] Ian Bogost. 2020. Your Phone Wasn't Built for the Apocalypse. <https://www.theatlantic.com/technology/archive/2020/09/camera-phone-wildfire-sky/616279/> Section: Technology.
- [16] Andy Boucher. 2023. Research Products at Scale: Learnings from Designing Devices in Multiples of Ones, Tens, Hundreds and Thousands. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3544548.3581540>
- [17] Jean-Lou Chameau, William F. Ballhaus, and Herbert S. Lin. 2014. Foundational Technologies. In *Emerging and Readily Available Technologies and National Security: A Framework for Addressing Ethical, Legal, and Societal Issues*. National Academies Press (US). <https://www.ncbi.nlm.nih.gov/books/NBK216326/>
- [18] Culture, Media and Sport Committee. 2023. *Connected tech: AI and creative technology*. Technical Report HC 1643. House of Commons, London, UK. <https://publications.parliament.uk/pa/cm5803/cmselect/cmcmds/1643/report.html>
- [19] Aayushi Dangol and Sayamindu Dasgupta. 2023. Constructionist approaches to critical data literacy: A review. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference (IDC '23)*. Association for Computing Machinery, New York, NY, USA, 112–123. <https://doi.org/10.1145/3585088.3589367>
- [20] Sandro Debono. 2021. Collecting Pandemic Phenomena: Reflections on Rapid Response Collecting and the Art Museum. *Collections* 17, 2 (June 2021), 179–185. <https://doi.org/10.1177/1550190620980844> Publisher: SAGE Publications Inc.
- [21] Department for Education. 2023. *Generative artificial intelligence in education*. Technical Report. UK Government. <https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education>
- [22] Graham Dove and Anne-Laure Fayard. 2020. Monsters, Metaphors, and Machine Learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–17. <https://doi.org/10.1145/3313831.3376275>
- [23] Stefania Druga, Sarah T. Vu, Eesh Likhith, and Tammy Qiu. 2019. Inclusive AI literacy for kids around the world. In *Proceedings of FabLearn 2019 (FL2019)*. Association for Computing Machinery, New York, NY, USA, 104–111. <https://doi.org/10.1145/3311890.3311904>
- [24] Upol Ehsan and Mark O. Riedl. 2020. Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach. In *HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence (Lecture Notes in Computer Science)*, Constantine Stephanidis, Masaaki Kuroso, Helmut Degen, and Lauren Reinerman-Jones (Eds.). Springer International Publishing, Cham, 449–466. https://doi.org/10.1007/978-3-030-60117-1_33
- [25] Robin Emsley. 2023. ChatGPT: these are not hallucinations – they're fabrications and falsifications. *Schizophrenia* 9, 1 (Aug. 2023), 1–2. <https://doi.org/10.1038/s41537-023-00379-4> Number: 1 Publisher: Nature Publishing Group.
- [26] Rose Etherington. 2013. Interview with Kieran Long on rapid response collecting at the V&A. <https://www.dezeen.com/2013/12/18/rapid-response-collecting-victoria-and-albert-museum-kieran-long/> Section: all.

- [27] Rebecca Eynon and Erin Young. 2021. Methodology, Legend, and Rhetoric: The Constructions of AI by Academia, Industry, and Policy Groups for Lifelong Learning. *Science, Technology, & Human Values* 46, 1 (Jan. 2021), 166–191. <https://doi.org/10.1177/0162243920906475> Publisher: SAGE Publications Inc.
- [28] M. A. Fitzgerald. 1994. Theories of Reflection for Learning. In *Reflective Practice in Nursing: The Growth of the Professional Practitioner*. Blackwell Scientific Publications, 63–84. https://figshare.utas.edu.au/articles/chapter/Theories_of_Reflection_for_Learning/23054621/1
- [29] Department for Education. 2013. *The national curriculum in England: Key stages 1 and 2 framework document*. Technical Report DFE-00178-2013. Department for Education.
- [30] 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.
- [31] 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>
- [32] William W. Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a Resource for Design. In *Proceedings of CHI 2003*. ACM Press, 269–270.
- [33] William W. Gaver, John Bowers, Andrew Boucher, Hans Gellerson, Sarah Pennington, Albrecht Schmidt, Anthony Steed, Nicholas Villars, and Brendan Walker. 2004. The Drift Table: Designing for Ludic Engagement. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. ACM, New York, NY, USA, 885–900. <https://doi.org/10.1145/985921.985947> event-place: Vienna, Austria.
- [34] Stefania Giannini. 2023. *Generative AI and the future of education*. Technical Report ED/ADG/2023/02. UNESCO, 8 pages.
- [35] Kalervo N. Gulson and Kevin Witzemberger. 2022. Repackaging authority: artificial intelligence, automated governance and education trade shows. *Journal of Education Policy* 37, 1 (Jan. 2022), 145–160. <https://doi.org/10.1080/02680939.2020.1785552> Publisher: Routledge _eprint: <https://doi.org/10.1080/02680939.2020.1785552>
- [36] Kim Halskov and Nicolai Brodersen Hansen. 2015. The diversity of participatory design research practice at PDC 2002–2012. *International Journal of Human-Computer Studies* 74 (Feb. 2015), 81–92. <https://doi.org/10.1016/j.ijhcs.2014.09.003>
- [37] Ariel Han and Zhenyao Cai. 2023. Design implications of generative AI systems for visual storytelling for young learners. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference (IDC '23)*. Association for Computing Machinery, New York, NY, USA, 470–474. <https://doi.org/10.1145/3585088.3593867>
- [38] Samantha Hautea, Sayamindu Dasgupta, and Benjamin Mako Hill. 2017. Youth Perspectives on Critical Data Literacies. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 919–930. <https://doi.org/10.1145/3025453.3025823>
- [39] Fred Hohman, Andrew Head, Rich Caruana, Robert DeLine, and Steven M. Drucker. 2019. Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 579:1–579:13. <https://doi.org/10.1145/3290605.3300809> event-place: Glasgow, Scotland Uk.
- [40] Krystal Hu. 2023. ChatGPT sets record for fastest-growing user base - analyst note. *Reuters* (Feb. 2023). <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- [41] 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>
- [42] Tim Ingold. 2018. Back to the future with the theory of affordances. *HAAU: Journal of Ethnographic Theory* 8, 1-2 (March 2018), 39–44. <https://doi.org/10.1086/698358> Publisher: The University of Chicago Press.
- [43] Ellen Jiang, Kristen Olson, Edwin Toh, Alejandra Molina, Aaron Donsbach, Michael Terry, and Carrie J Cai. 2022. PromptMaker: Prompt-based Prototyping with Large Language Models. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3491101.3503564>
- [44] Sarah Judd. 2020. All Means All: Bringing Project-based, Approachable AI Curriculum to More High School Students through AI4ALL Open Learning. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*. Association for Computing Machinery, New York, NY, USA, 1409. <https://doi.org/10.1145/3328778.3372554>
- [45] Yasmin B. Kafai. 2005. Constructionism. In *The Cambridge Handbook of the Learning Sciences*, R. Keith Sawyer (Ed.). Cambridge University Press, Cambridge, UNITED KINGDOM, 35–46. <http://ebookcentral.proquest.com/lib/lancaster/detail.action?docID=261112>
- [46] Yasmin B. Kafai and Quinn Burke. 2015. Constructionist Gaming: Understanding the Benefits of Making Games for Learning. *Educational Psychologist* 50, 4 (Oct. 2015), 313–334. <https://doi.org/10.1080/00461520.2015.1124022>
- [47] Harmanpreet Kaur, Harsha Nori, Samuel Jenkins, Rich Caruana, Hanna Wallach, and Jennifer Wortman Vaughan. 2020. Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–14. <https://doi.org/10.1145/3313831.3376219>
- [48] Alyson Klein. 2023. ChatGPT Cheating: What to Do When It Happens. *Education Week* (Feb. 2023). <https://www.edweek.org/technology/chatgpt-cheating-what-to-do-when-it-happens/2023/02>
- [49] David A. Kolb, Richard E. Boyatzis, and Charalampos Mainemelis. 2001. Experiential learning theory: Previous research and new directions. In *Perspectives on thinking, learning, and cognitive styles*. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, US, 227–247.
- [50] Peter Gall Krogh and Ilpo Koskinen. 2020. *Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research*. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-030-37896-7>
- [51] Olya Kudina and Peter-Paul Verbeek. 2019. Ethics from Within: Google Glass, the Collingridge Dilemma, and the Mediated Value of Privacy. *Science, Technology, & Human Values* 44, 2 (March 2019), 291–314. <https://doi.org/10.1177/0162243918793711>
- [52] F. Lau. 1997. A Review on the Use of Action Research in Information Systems Studies. In *Information Systems and Qualitative Research: Proceedings of the IFIP TC8 WG 8.2 International Conference on Information Systems and Qualitative Research, 31st May–3rd June 1997, Philadelphia, Pennsylvania, USA*, Allen S. Lee, Jonathan Liebenau, and Janice I. DeGross (Eds.). Springer US, Boston, MA, 31–68. https://doi.org/10.1007/978-0-387-35309-8_4
- [53] Irene Lee, Safinah Ali, Helen Zhang, Daniella DiPaola, and Cynthia Breazael. 2021. Developing Middle School Students' AI Literacy. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, New York, NY, USA, 191–197. <https://doi.org/10.1145/3408877.3432513>
- [54] Unggi Lee, Ariel Han, Jeongjin Lee, Eunseo Lee, Jiwon Kim, Hyeoncheol Kim, and Cheolil Lim. 2023. Prompt Aloud!: Incorporating image-generative AI into STEAM class with learning analytics using prompt data. *Education and Information Technologies* (Sept. 2023). <https://doi.org/10.1007/s10639-023-12150-4>
- [55] Wen-Ying Lee and Malte Jung. 2020. Ludic-HRI: Designing Playful Experiences with Robots. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. Association for Computing Machinery, New York, NY, USA, 582–584. <https://doi.org/10.1145/3371382.3377429>
- [56] Ilya Levin and Dina Tsybulsky. 2017. The Constructionist Learning Approach in the Digital Age. *Creative Education* 8, 15 (Dec. 2017), 2463–2475. <https://doi.org/10.4236/ce.2017.815169> Number: 15 Publisher: Scientific Research Publishing.
- [57] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376590>
- [58] Joseph Lindley, Haider Ali Akmal, Franziska Pilling, and Paul Coulton. 2020. Researching AI Legibility through 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/10.1145/3313831.3376792>
- [59] Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017. Implications for Adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 265–277. <https://doi.org/10.1145/3025453.3025742>
- [60] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3313831.3376727>
- [61] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, Honolulu, HI, USA, 1–14. <https://doi.org/10.1145/3313831.3376445>
- [62] John Maloney, Mitchel Resnick, Natalie Rusk, Brian Silverman, and Evelyn Eastmond. 2010. The Scratch Programming Language and Environment. *ACM Transactions on Computing Education* 10, 4 (Nov. 2010), 16:1–16:15. <https://doi.org/10.1145/1868358.1868363>
- [63] Alexander G. Mirnig, Magdalena Gärtner, Alexander Meschtscherjakov, and Manfred Tscheligi. 2020. Blinded by novelty: a reflection on participant curiosity and novelty in automated vehicle studies based on experiences from the field. In *Proceedings of Mensch und Computer 2020 (MuC '20)*. Association for Computing Machinery, New York, NY, USA, 373–381. <https://doi.org/10.1145/3404983.3405593>
- [64] Yisroel Mirsky and Wenke Lee. 2021. The Creation and Detection of Deepfakes: A Survey. *Comput. Surveys* 54, 1 (Jan. 2021), 7:1–7:41. <https://doi.org/10.1145/3425780>
- [65] Selena Nemorin, Andreas Vlachidis, Hayford M. Ayerakwa, and Panagiotis Andriotti. 2023. AI hyped? A horizon scan of discourse on artificial intelligence

- in education (AIED) and development. *Learning, Media and Technology* 48, 1 (Jan. 2023), 38–51. <https://doi.org/10.1080/17439884.2022.2095568> Publisher: Routledge _eprint: <https://doi.org/10.1080/17439884.2022.2095568>.
- [66] Chrissi Nerantzi, Sandra Abegglen, Marianna Karatsiori, and Antonio Martínez-Arboleda (Eds.). 2023. *101 creative ideas to use AI in education, A crowdsourced collection*. Zenodo. <https://doi.org/10.5281/zenodo.8072950> Version Number: 2023 1.0.
- [67] Donald A. Norman. 1999. Affordance, conventions, and design. *interactions* 6, 3 (1999), 38–43. Publisher: ACM New York, NY, USA.
- [68] William Odum, Ron Wakkary, Youn-kyung Lim, Audrey Desjardins, Bart Hengeveld, and Richard Banks. 2016. From Research Prototype to Research Product. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2549–2561. <https://doi.org/10.1145/2858036.2858447>
- [69] Seymour Papert. 1991. Situating constructionism. In *Constructionism*, Seymour Papert and Idit Harel (Eds.). Vol. 36. Ablex Publishing Corporation, 1–11. <https://pirun.ku.ac.th/~btun/papert/sitcons.pdf>
- [70] Carlo Perrotta and Neil Selwyn. 2020. Deep learning goes to school: toward a relational understanding of AI in education. *Learning, Media and Technology* 45, 3 (July 2020), 251–269. <https://doi.org/10.1080/17439884.2020.1686017> Publisher: Routledge _eprint: <https://doi.org/10.1080/17439884.2020.1686017>.
- [71] Pew Research Center. 2020. Children’s engagement with digital devices, screen time. <https://www.pewresearch.org/internet/2020/07/28/childrens-engagement-with-digital-devices-screen-time/>
- [72] Fabio Pittarello, Gualtiero Volpe, and Massimo Zancanaro. 2017. HCI and education in a changing world: from school to public engagement. In *Proceedings of the 12th Biannual Conference on Italian SIGCHI Chapter (CHIItaly '17)*. Association for Computing Machinery, New York, NY, USA, 1–2. <https://doi.org/10.1145/3125571.3125576>
- [73] Johan Redström. 2017. *Making Design Theory*. The MIT Press. <https://doi.org/10.7551/mitpress/11160.001.0001>
- [74] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. <https://doi.org/10.48550/arXiv.2112.10752> arXiv:2112.10752 [cs].
- [75] Elisa Rubegni, Laura Malinverni, and Jason Yip. 2022. “Don’t let the robots walk our dogs, but it’s ok for them to do our homework”: children’s perceptions, fears, and hopes in social robots.. In *Proceedings of the 21st Annual ACM Interaction Design and Children Conference (IDC '22)*. Association for Computing Machinery, New York, NY, USA, 352–361. <https://doi.org/10.1145/3501712.3529726>
- [76] Debjani Saha, Candice Schumann, Duncan C. McElfresh, John P. Dickerson, Michelle L. Mazurek, and Michael Carl Tschantz. 2020. Measuring Non-Expert Comprehension of Machine Learning Fairness Metrics. In *PMLR*, Vol. 119. <http://arxiv.org/abs/2001.00089> arXiv: 2001.00089.
- [77] Adam Satariano and Cecilia Kang. 2023. How Nations Are Losing a Global Race to Tackle A.I.’s Harms. *The New York Times* (Dec. 2023). <https://www.nytimes.com/2023/12/06/technology/ai-regulation-policies.html>
- [78] Marie-Monique Schaper, Rachel Charlotte Smith, Mariana Aki Tamashiro, Maarten Van Mechelen, Mille Skovhus Lunding, Karl-Emil Kjær Bilstруп, Magnus Høholt Kaspersen, Kasper Løvborg Jensen, Marianne Graves Petersen, and Ole Sejer Iversen. 2022. Computational empowerment in practice: Scaffolding teenagers’ learning about emerging technologies and their ethical and societal impact. *International Journal of Child-Computer Interaction* 34, C (Dec. 2022). <https://doi.org/10.1016/j.ijcci.2022.100537>
- [79] Douglas Schuler and Aki Namioka. 1993. *Participatory Design: Principles and Practices*. CRC Press. Google-Books-ID: pW0Ek6Sk4YkC.
- [80] Donald A. Schön. 1983. *The reflective practitioner: how professionals think in action*. Basic Books, New York.
- [81] 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>
- [82] Rachel Charlotte Smith, Marie-Monique Schaper, Mariana Aki Tamashiro, Maarten Van Mechelen, Marianne Graves Petersen, and Ole Sejer Iversen. 2023. A research agenda for computational empowerment for emerging technology education. *International Journal of Child-Computer Interaction* 38 (Dec. 2023), 100616. <https://doi.org/10.1016/j.ijcci.2023.100616>
- [83] Aaron Springer, Victoria Hollis, and Steve Whittaker. 2017. Dice in the black box: User experiences with an inscrutable algorithm. In *The 2017 AAAI Spring Symposium Series*.
- [84] Pieter Stappers and Elisa Giaccardi. 2019. Research through Design. In *The Encyclopedia of Interaction Design* (2 ed.). Interaction Design Foundation. <https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/research-through-design>
- [85] Jiahong Su and Yuchun Zhong. 2022. Artificial Intelligence (AI) in early childhood education: Curriculum design and future directions. *Computers and Education: Artificial Intelligence* 3 (Jan. 2022), 100072. <https://doi.org/10.1016/j.caeai.2022.100072>
- [86] UNESCO. 2022. *K-12 AI curricula: a mapping of government-endorsed AI curricula*. Technical Report 0000380602. UNESCO, Paris. <https://unesdoc.unesco.org/ark:/48223/pf0000380602>
- [87] Maarten Van Mechelen, Gökçe Elif Baykal, Christian Dindler, Eva Eriksson, and Ole Sejer Iversen. 2020. 18 Years of ethics in child-computer interaction research: a systematic literature review. In *Proceedings of the Interaction Design and Children Conference (IDC '20)*. Association for Computing Machinery, New York, NY, USA, 161–183. <https://doi.org/10.1145/3392063.3394407>
- [88] Victoria and Albert Museum. 2017. Rapid Response Collecting | V&A. https://www.youtube.com/watch?v=_9AQXrKlKVY
- [89] Benjamin Walsh, Safinah Ali, Francisco Castro, Kayla Desportes, Daniella Di-Paola, Irene Lee, William Payne, Scott Sieke, and Helen Zhang. 2022. Making Art with and about Artificial Intelligence: Three Approaches to Teaching AI and AI Ethics to Middle and High School Students. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V, 2 (SIGCSE 2022)*. Association for Computing Machinery, New York, NY, USA, 1203. <https://doi.org/10.1145/3478432.3499157>
- [90] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y. Lim. 2019. Designing Theory-Driven User-Centric Explainable AI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Glasgow, Scotland Uk, 1–15. <https://doi.org/10.1145/3290605.3300831>
- [91] Kyle Wiggers. 2023. Stability AI releases its latest image-generating model, Stable Diffusion XL 1.0. <https://techcrunch.com/2023/07/26/stability-ai-releases-its-latest-image-generating-model-stable-diffusion-xl-1-0/>
- [92] Randi Williams, Hae Won Park, Lauren Oh, and Cynthia Breazeal. 2019. PopBots: Designing an Artificial Intelligence Curriculum for Early Childhood Education. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (July 2019), 9729–9736. <https://doi.org/10.1609/aaai.v33i01.33019729> Number: 01.
- [93] Ben Williamson. 2023. Degenerative AI in education. <https://codeactsineducation.wordpress.com/2023/06/30/degenerative-ai-in-education/>
- [94] Ben Williamson and Rebecca Eynon. 2020. Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology* 45, 3 (July 2020), 223–235. <https://doi.org/10.1080/17439884.2020.1798995> Publisher: Routledge _eprint: <https://doi.org/10.1080/17439884.2020.1798995>.
- [95] Nesra Yannier, Scott E. Hudson, Kenneth R. Koedinger, Kathy Hirsh-Pasek, Roberta Michnick Golinkoff, Yuko Munakata, Sabine Doebel, Daniel L. Schwartz, Louis Deslauriers, Logan McCarty, Kristina Callaghan, Elli J. Theobald, Scott Freeman, Katelyn M. Cooper, and Sara E. Brownell. 2021. Active learning: “Hands-on” meets “minds-on”. *Science* 374, 6563 (Oct. 2021), 26–30. <https://doi.org/10.1126/science.abj9957> Publisher: American Association for the Advancement of Science.
- [96] Michalinos Zembylas. 2023. A decolonial approach to AI in higher education teaching and learning: strategies for undoing the ethics of digital neocolonialism. *Learning, Media and Technology* 48, 1 (Jan. 2023), 25–37. <https://doi.org/10.1080/17439884.2021.2010094> Publisher: Routledge _eprint: <https://doi.org/10.1080/17439884.2021.2010094>.
- [97] John Zimmerman and Jodi Forlizzi. 2014. Research Through Design in HCI. In *Ways of Knowing in HCI*, Judith S. Olson and Wendy A. Kellogg (Eds.). Springer New York, New York, NY, 167–189. https://doi.org/10.1007/978-1-4939-0378-8_8
- [98] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research Through Design As a Method for Interaction Design Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 493–502. <https://doi.org/10.1145/1240624.1240704> event-place: San Jose, California, USA.
- [99] John Zimmerman, Erik Stolterman, and Jodi Forlizzi. 2010. An Analysis and Critique of Research Through Design: Towards a Formalization of a Research Approach. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems (DIS '10)*. ACM, New York, NY, USA, 310–319. <https://doi.org/10.1145/1858171.1858228> event-place: Aarhus, Denmark.

Received 8 February 2024; revised 3 May 2024