

Towards A Process Model for Co-Creating AI Experiences

HARIHARAN SUBRAMONYAM, University of Michigan, USA

COLLEEN SEIFERT, University of Michigan, USA

EYTAN ADAR, University of Michigan, USA

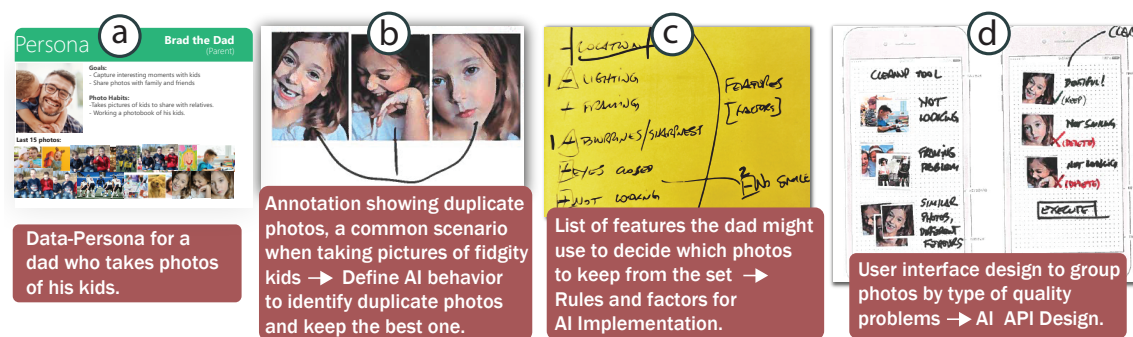


Fig. 1. User Data as a probe to design AI material: (a) The designer uses data-persona to construct (b) a scenario about a parent taking duplicate photos (designs AI behavior), (c) listing the features the parent might use to identify bad photos (defines how AI should implement the behavior), and (d) creating AI-Powered UI for de-cluttering photo albums (defines inputs and outputs to AI).

Thinking of technology as a design *material* is appealing. It encourages designers to explore the material’s properties to understand its capabilities and limitations—a prerequisite to generative design thinking. However, as a material, AI resists this approach because its properties only *emerge* as part of the user experience design. Therefore, designers and AI engineers must collaborate in new ways to create both the material and its application experience. We investigate the *co-creation* process through a design study with 10 pairs of designers and engineers. We find that design ‘probes’ with user data are a useful tool in defining AI materials. Through data probes, designers construct designerly representations of the envisioned AI experience (AIX) to identify desirable AI characteristics. Data probes facilitate divergent design thinking, material testing, and design validation. Based on our findings, we propose a process model for co-creating AIX and offer design considerations for incorporating data probes in AIX design tools.

CCS Concepts: • **Human-centered computing** → **User centered design**; • **Computing methodologies** → **Artificial intelligence**; • **Software and its engineering** → **Collaboration in software development**.

Additional Key Words and Phrases: Design Materials, Human-Centered AI, User Data

ACM Reference Format:

Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. Towards A Process Model for Co-Creating AI Experiences. In *Designing Interactive Systems Conference 2021 (DIS '21), June 28–July 2, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3461778.3462012>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

1 INTRODUCTION

When working with new and unfamiliar technology, designers are encouraged to consider it from a “material” perspective [26, 73, 88]. Just as with wood or fabric, in which the craftsman needs to understand the material to create with it, designers need to know what the technology is capable of, what its limitations are, and what properties are available for design. For instance, when working with Radio Frequency ID (RFID) technology, the designer should first explore its material properties including the signal strength, how much information an RF tag can hold, and how quickly the information can be read [4]. This will allow the designer to *manipulate* its properties in creative ways to generate design solutions [6, 16, 76]. However, unlike other technology materials that are created *before* the user experience (UX) design, artificial intelligence (AI) does not lend itself well to a purely material-driven design approach [49, 94]. Instead, AI’s material properties only emerge through its application experience design [20, 50]. As a simple example, to design an intelligent To-Do List application that automatically creates tasks from emails (e.g., [25]), designers cannot work with AI as a given material that makes predictions from text. To create the AI material, AI engineers need guidance from designers about how end-users think about tasks, who the potential end-users are, what emails mean to users, and so on (i.e., human-centered AI [5]). For both the designer and the AI engineer, this is a challenging *chicken-and-egg* problem [68, 75]. To address this, we investigate what a co-creation process for AI’s *form* and *function* might look like, what AI as a material “under construction” entails, and how the evolving UX design can inform AI development.

A fundamental assumption of the material view in HCI is that materials are a given, and they *possess* specific properties that are amenable to design. To a large extent, this assumption holds. From an engineering standpoint, materials are invented with specific *structure-property* relationships in mind (e.g., [46, 77, 81]). They can be used in any context in which those relationships are desirable [1, 35, 66]. The designer’s job is then to explore the material, understand how end-users might experience it, and thereby acquire knowledge for generative design thinking [26]. For instance, to prototype a To-Do List in mixed-reality (a novel material), the designer can begin by exploring how graphics are rendered spatially, what the visual field of view entails, and which hand gestures are available for interactivity (e.g., [64]). Based on this, designers can prototype alternatives by changing the appearance of graphical elements [61], exploring different gestures and layout options to design the UX. In other words, design with material is accomplished by knowing its *created* properties. Even in extreme cases of customization [54], the design material metaphor holds.

In this vein, AI is also a novel design material [39, 90]. However, there are important distinctions that make it challenging to put this material perspective into practice. First, as a given (or prefabricated) material, AI is *deficient* for design. AI materials are commonly described in abstractions such as techniques (e.g., supervised-learning) and behavior (e.g., prediction), and divorced from contexts in which the AI is applied. Unlike the mixed-reality example, a designer cannot simply explore a “supervised learning” AI to design an intelligent To-Do List. Designing with AI requires *defining* its material properties, including what the AI system should learn using what data, which assumptions and learning rules are appropriate, and how those capabilities should manifest in designed experiences (e.g., data labels). Second, once created by AI engineers, an AI system’s properties cannot be readily manipulated during the application’s design process. In comparison to mixed reality interfaces, in which the designer can change properties (such as color or shape of elements), representational and knowledge barriers prevent designers from directly altering AI to mold it into a ‘designed’ product (e.g., [91]). Third, in many AI systems, its material characteristics can continue to evolve through feedback and learning. The designers must anticipate how the AI will change over time and experience. These capabilities require design across both the application and the AI material.

To consider the intersection of AI creation and human-centered application design, which we call *AI experience design* (AIX), HCI researchers have put forth design guidelines [3, 29], processes [94], and tools [85]. However, these recommendations emphasize the designer’s responsibility to understand the AI design material, but not the role of AI practitioners in AIX. The material design approach assumes the AI, like natural wood, must be taken as given; so, the designer provides the required adaptation. For designers, this challenges their technical expertise and introduces friction to material exploration [93]. For instance, designers cannot prototype the user experience with a “fail fast, fail often [91]” approach. With AI material, *vertical* end-to-end prototyping is required to create, evaluate, and revise design alternatives. Such a process is time-consuming for designers and engineers, it is resource-intensive, and the chicken-and-egg problem remains [50]. This is the primary motivation for our work: *How might designers and engineers co-create AI Experiences through rapid, collaborative design?*

Drawing from prior research in HCI and AI application design, we developed a protocol for co-creating AIX through generative design thinking. Using the protocol, pairs of UX designers and AI engineers worked to design AI material characteristics and the user experience for a given design problem. We observed UX designers’ involvement in AI material creation, including defining AI behavior, specifying the AI architecture, and features related to explainability, failure, and learnability. We found that end-user data played a critical role in shaping AIX. As shown in Figure 1, by using data as a design probe, designers constructed AI-infused scenarios to co-design desired AI behavior. By imagining mental-models for different personas across scenarios and data (i.e., how might the persona perform a task?), they offered inputs to AI architecture design. Through user interface prototyping with data, participants also co-designed the application programming interface (API). Teams created data probes as a scaffold for divergent design thinking, material testing, and design validation. Our key contributions include (1) identifying the role of designerly proxies¹ and data probes in defining AIX material, (2) describing a process model for co-creating AIX, and (3) highlighting a set of design considerations for incorporating data probes in AIX design tools.

2 RELATED WORK

A material framework for design includes (1) *fabrication*—ways to produce materials with specific properties, (2) *application*—ways to transform materials into products, and (3) *appreciation*—reception of material by the end-users [19]. Design requires iteration and feedback across these three aspects. When fabrication and application are cleanly separated—as with natural materials like wood and technological materials like RFID—prior work has considered design material as a *given* [4, 18, 24, 26, 33, 53, 65, 74]. As a consequence, UX design emphasizes understanding material *properties*, developing *processes* to generate material artifacts, and evaluating *expectations* and *values* associated with material encounters. Other work has combined material creation (fabrication) into the design problem and investigated *co-creation* of the material along with its application [7, 48, 55, 71, 76, 84]. We draw from both of these perspectives to develop our understanding of designing the AI material and designing with it. By characterizing ‘design’ as an activity that applies a value system to create objects of reasoning [7], with AI as design material, we identify gaps in guidelines, methods, and representations. Through this discussion, we formulate the research questions for our study.

2.1 Guidelines

Numerous design guidelines for AI applications have emerged from both academic and industry research. The guidelines span across functionality [45], end-user interactions [3, 29, 34], learnability [27], explainability [87], privacy [32, 47],

¹Designerly proxies are the designer’s representations of AI’s technical characteristics.

transparency [21], etc. Several guidelines address the intersections of both fabrication and application design (i.e., AIX). In some cases, the guidelines ask that we consider application context when creating AI capabilities; For instance, PAIR [29] recommends modeling AI after the human expert: “When designing automation, we should consider how a theoretical human ‘expert’ might perform the task today”. Others offer suggestions for repairing AI material flaws through UX enhancements: “Make it easy to edit, refine, or recover when the AI system is wrong [3].” This highlights the inherent dependencies between fabrication and application design for AI [20, 93]. In our work, we look at how designers and AI engineers conceptualize the guidelines from different “points of view” to co-create AIX. To aid our investigation, we consider the language of material engineering which provides a vocabulary for material as *structure*, *surface*, and *properties* [84]. This would allow us to establish connections between material characteristics and AI material experience in its embodiment, encounters, and collaborations [26].

Research Question 1: *How might designers and AI engineers conceptualize shared, and differing, design perspectives arising from human-AI guidelines to co-create AIX?*

2.2 Design Methods

Current AI development workflows consist of critical design-related steps, including identifying model requirements, data labeling, feature engineering, etc. [2]. However, in many cases, UX design and AI development only converge after AI decisions have been made [20]. Consequently, UX designers face challenges in incorporating AI material properties within their design practices [90]. Similarly, engineers find it challenging to obtain ground truth validation for AI-related decisions [38], avoid blind-spots threatening responsible AI needs [41], and incorporate necessary UX inputs for improving model performance [83]. Hence, design methodologies should be symbiotic to produce the best application performance. If AI is meant to replicate human intelligence, UX designers can offer insights to make it practically and emotionally resonant with users [15, 89]. The main challenges to collaboration have been time-related constraints to fabrication and design [92], barriers to immediate feedback [50, 91], and lack of motivation and incentives [17, 50]. Further, AI material challenges conventional prototyping methods because it requires a higher level of commitment and effort to prototype AI applications [91, 93]. We also lack means for UX designers to engage in a “conversation with the materials [90],” and ways for the AI materials to “talk back to the designer [90].”

When considering software code as design material, programming becomes a vital part of the design process and offers necessary “talk-backs” for design [55]. Even in natural materials such as wood, material properties (hardness, grain), and constraints (knots and weak points) offer feedback resulting in design recourse [18]. In co-creating AIX, both designers and engineers require the material and the application experience to respond to each other. Separately, when working with Bluetooth as novel design material, designers generate end-to-end fully working sketches (inspirational bits), allowing them to investigate its properties through form-giving [80]. To co-create AIX, designers and engineers need similar low-cost *vertical-prototyping* strategies to create end-to-end prototypes of the UX and the AI backend [6]. Tools for material prototyping should be accessible (allow developers and designers to think about the material), immediate (support rapid iterative feedback, reflection-in-action, and reflection-on-action), and generative (allow test, probe, and exploration iterations) [33]. Lastly, developers, engineers, and data scientists employ more mechanical, less organic processes focused on application *data* [28]. Consequently, a key task for human-centered designers is promoting user experience data as a bridge between the two fields through a process of translation [28, 37]. We incorporate these perspectives about *data* in developing our study protocol for AIX.

Research Question 2: *How might designers and engineers co-create the design and technical characteristics of AIX?*

2.3 Representations

Design requires creating and comparing alternatives to arrive at a final solution [76]. The challenge with AIX is finding intermediate design representations that can serve as a “lingua franca” easily understood by multi-disciplinary teams of designers and engineers [7]. Representations would ideally allow the design concept to be viewed differently based on functional perspectives. For instance, for engineers, data is represented as a set of variables [11]; but in UX, data is associated with end-users and their situated context [36]. Designing across boundaries requires a show and tell (“I will know it when I see it [8]”) approach. In [50], active discussions between UX designers and ML researchers around mock-ups helped them avoid miscommunication about model capabilities. These rich representations should provide ways to envision viewpoints and resolve differences, and align design needs through negotiation over intermediate representations including words, sketches, physical mock-ups, charts, etc [10, 44]. At the same time, these representations should be easy to create during rapid prototyping. In prototyping web applications, designers use non-functional proxies to negotiate a design that works for both developers and end-users (e.g., Interfake [42], Apiary [67], etc.). Such ‘mocks’ can circumvent the need for more programming effort to creating AI material designs. Our study explores mixed-fidelity prototypes [60] that provide high-fidelity representations in some dimensions and low fidelity in others.

Research Question 3: What types of material and design representations can support co-creating AIX?

3 METHOD

We conducted an in-lab design study in which UX designers paired with AI engineers worked together to co-create AIX (a total of 10 sessions, each with one designer and one engineer). To model the nature of collaboration, we took inspiration from Wizard-of-Oz (WoZ) techniques for AI prototyping [9, 59, 85]. We imagined that designers and engineers would implicitly play the ‘wizard’ (experts) role during co-creation. This allows for rapid feedback about both the AI material being created by engineers and application experience being prototyped by the designer. Therefore, we recruited participants who had prior experience in AI application domains and working in collaborative teams; The designers in our study had an average of 3.4 years of experience ($SD = 2.8$), and AI engineers had 3.9 years of experience on average ($SD = 2.1$). Participants comprised of industry practitioners as well as graduate students with prior work experience (Table 1). Participants were paired based on their availability. Each session lasted 2.5 hrs, and we compensated participants with \$40 for their time. All sessions were video-recorded. We collected all artifacts generated by the participants for our analysis.

3.1 Study Protocol

To develop the protocol, we started with human-AI (HAI) design guidelines, developed by companies, for designers and engineers [3, 29, 45, 69]. Ideally, these guidelines represent ‘best-practice’ advice that is derived from successful practices within the companies. They offer a starting point to explore how UX and AI roles might collaborate in designing the AI experience. We categorized the guidelines into seven steps spanning AI creation, UX design, and AI-UX design processes (see Figure 2). Our steps roughly followed the material design process [49] of proposing material, envisioning material experience, manifesting material experience patterns, and making material product concept designs. In addition, we included material creation in the process. In our instructions to participants, we refrained from using the material metaphor; instead, we worked with terminology specified in the HAI guidelines. Further, we organized the steps into two phases; The first phase aimed at producing initial AI specifications (fabrication) and prototypes of the AI-powered UI (application). This included opportunity spotting, model specification, and UI prototyping. In the second phase,

Session ID	UX Designer	AI Engineer
1	4 months	1 years & 4 months
2	3 months	6 years
3	4 years & 2 months	2 years & 1 month
4	7 years & 5 months	2 years & 9 months
5	4 years & 6 months	1 years & 6 months
6	3 months	7 years
7	4 years & 5 months	2 years
8	6 years & 2 months	4 years
9	5 years & 7 months	5 years & 6 months
10	3 months	6 years & 5 months

Table 1. Participant details indicating years of experience for designers and engineers in our study.

participants iterated over the design to arrive at a ‘pragmatic’ solution by considering errors, explainability, feedback, and expectation-setting for end-users (appreciation). Our first session served as a pilot informing the two phases, and we used the feedback to revise the protocol for the remaining sessions. The second phase would allow teams to consider AI’s uncertainties and offer adaptations to account for AI errors, thereby maximizing AI’s utility for end-users. At each step of the protocol, the study coordinator offered instructions and tools for participants to work on that step. Each step had a time limit, and participants shared and discussed their design with the coordinator throughout the study.

Design Problem Briefing: We selected the problem of *decluttering the photo album* on the phone. We motivated the problem by stating that cameras on smartphones have made it easy to take photos anytime and anywhere. A consequence is that users capture hundreds of photos that may be of little value. Deleting unwanted photos can be tedious and boring. We asked participants to design an AI-powered solution to address this problem. This domain is simple enough for participants to understand, and they can leverage experiential knowledge in brainstorming solutions. In other words, the problem minimizes domain complexities while allowing us to observe collaboration in a single task session. We provided participants with screenshots of the current (non-AI) photo album interface that we created. In our pilot, we observed that the designer spent time creating the same interface by looking at their phone. Providing this design upfront allows us to focus time on more critical steps. We also specified what the back-end delete API looks like (a simple function that takes a list of photo ids to delete and returns a success or failure message).

Like Zhou et al. [94], we offered participants a set of initial persona cards, including a parent, a business traveler, a 3D-artist, and an Instagram influencer. For each persona, we listed their goals and photo-taking habits. Motivated by prior work in data-driven design [37] and parallel work in data visualization design (i.e., “data changed everything [86]”), we included a set of 15 most recent photos taken by each of the personas. Our data-personas align with the “minimum-viable-data [85]” concept that AI engineers typically use in prototyping machine learning models. From a UX standpoint, the data-driven persona is similar to what designers could generate through user research with mixed-method data (e.g., [78]). We carefully curated the photos to include a variety of images that represented a diverse set of photo capturing behavior and photo content. This was done to ensure a diverse set of AIX solutions.

Step 1—Opportunity Spotting (~25 minutes): We asked participants to brainstorm ways in which AI can support the decluttering tasks by aligning AI needs with human needs [29]. We provided them with information about types of AI (predictive, perceptual, generative) [69]. We also gave them guidelines about integrating AI experiences into end-user task workflows (e.g., when the AI should automatically perform a task, and when it should take an assistive role when explicitly invoked by end users) [45]. Participants had access to note pads, sticky notes, and colored markers throughout

the session to brainstorm. They were also free to annotate on any of the printed study materials. At the end of this step, participants converged on the AI capabilities they would design in the next steps.

Step 2—Model Specification (~20 minutes): In this step, participants continued brainstorming ways in which they would implement the AI capabilities that emerged in step 1. We provided them with an ML model design template to brainstorm about training data needs and factors they would consider for implementing the behavior [69]. This forced participants to externalize their thoughts and collaborate. We also provided them with printed spreadsheets with persona images on one column and empty columns to fill out with feature values. This encouraged a WoZ like approach to simulate model predictions. They were free to use it as a worksheet to iterate on the design.

Step 3—Vertical Prototyping (~30 minutes): At this point, we instructed the UX designers to prototype the user interface design using output from steps 1 and 2. We provided them with printed templates for wireframing mobile interfaces. We asked the designers not to use placeholders for text or images in their prototype. We provided them with printed images for each persona (both thumbnail size and screen-size images). Participants could cut and glue the images onto their prototypes (i.e., make medium-fidelity prototypes [22]). We intended to see how participants worked with factual data when designing the interface. For text and labels, when they were unclear what the content needs to be, they were asked to annotate with a question mark for later discussion with the engineer. In parallel, we asked the AI practitioner to fill out model API cards (one for each AI capability), providing details about API name, model inputs, model outputs, behavior description, and details about the training data (adapted from [62]). This comprised a low-cost realization of the AI “material” for feedback and iteration. At the end of this step, they each explained the model API design and the UI design to each other and the study coordinator. We provided participants with translucent sheets (vellum paper) to place on top of the prototypes to annotate and discuss. The goal was to map user interactions to API calls and align model inputs and outputs to the prototype. During this stage, engineers revised the model details through negotiation, and designers updated the interface when required.

Step 4—Identify AI Errors (~15 minutes): We created design cards explaining different types of AI errors and potential sources of errors [3, 29]. Using this information and the prototypes (UI and Model cards), we asked participants to brainstorm AI and UI specific errors for their design. We provided them with a template to document errors along with different categories (system limitation, context error, background error [29]), but they were free to use the notepad. For this step, participants had to generate a set of potential errors for their AIX design.

Step 5—Design for Explainability (~20 minutes): In this step, we instructed participants to consider explainability as a solution to the errors generated in the previous step. According to guidelines, context errors are a type of AIX errors in which the system is working as intended, but the user might perceive an error due to lack of understanding, or mismatch with their own mental model [29]. We asked participants to incorporate explainability into their design (both interface prototype and model API) to resolve AI context errors. We provided them with six design cards listing techniques and examples for designing explainable interfaces [29]. We also provided participants with vellum sheets which they could use to annotate over the prototype to design explainable solutions collaboratively.

Step 6—Design for Learnability and Feedback (~15 minutes): We asked participants to consider learnability and end-user feedback to improve the model performance. Participants had to design ways to elicit feedback from the users. The key here was to design feedback in a way that can be used for model improvement. We provided participants with information about the types of feedback and guidelines for designing explicit feedback [29].

Step 7—Setting Expectations for End-Users (~10 minutes): In this final step, participants had to design ways to communicate AI capabilities to end-users. We asked them to consider how they might design for end-user trust and

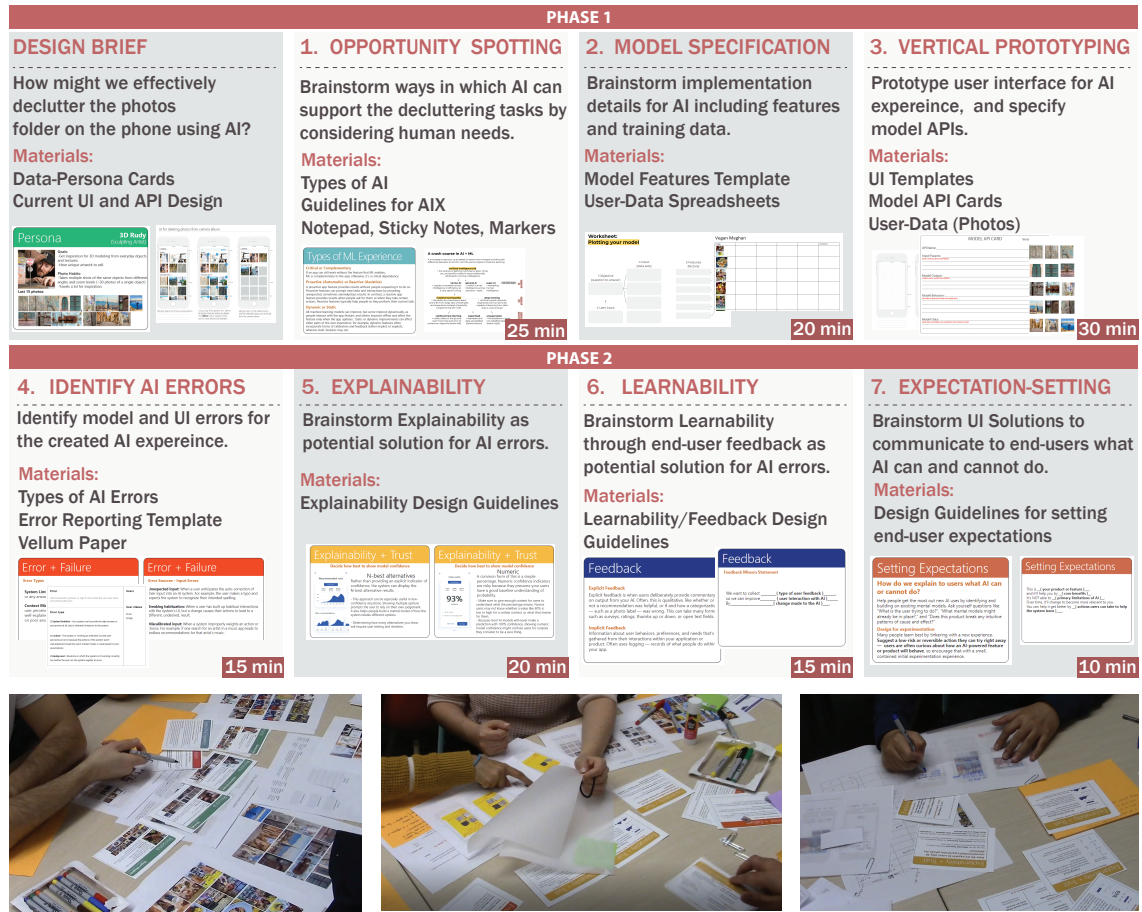


Fig. 2. Overview of our study protocol for co-creating AIX. Top: Design brief and high-level objective for each of the seven steps. Bottom: Visuals from In-Lab sessions.

how they might design to support end-user control over the data. We provided participants with guidelines about trust and expectation setting [3, 29]. Participants could create new wireframes or annotate over existing ones.

Collectively these steps follow the material design process by considering fabrication, application, and appreciation and would allow teams to offer adaptations for AI's uncertainties during the co-creation process. At the end of the study, we debriefed participants about our motivation to investigate AI co-creation process based on HAI guidelines. Participants had the opportunity to ask us questions and provide feedback on the study protocol.

3.2 Data Analysis

To prepare the data for analysis, the first author manually transcribed all video recordings. This allowed them to annotate and capture necessary metadata, such as who created the artifacts and how designers and engineers engaged with the study materials. During transcription, they included screen captures of the video to indicate pointing and show-and-tell actions. They also added scanned copies of corresponding artifacts at appropriate points in the transcript.

This was done in-line in a word document (one for each session). We then conducted qualitative coding utilizing a combination of deductive and inductive codes [23]. From literature and our protocol steps, we generated an initial set of codes (e.g., AI fabrication, application design, structure, properties, surface, etc.). For instance, we coded AI implementation details as material structure, and discussions about aligning the UI prototypes and model cards as material surface, i.e., the model API. After coding the transcripts using these codes, we carried out the second round of inductive *in-vivo* coding to analyze the data within each category. The generated codes included references to data and types of communication between designers and engineers (knowledge sharing, negotiation, artifact purpose, validation, guidelines, mentions of UI and AI design elements, etc.). After coding, the authors collectively reviewed and discussed the coded transcripts to identify higher-level themes to answer our research questions on co-creating AIX.

4 FINDINGS

We asked participants to co-create an experience for decluttering photo albums using AI (in the abstract). By considering human-centered needs, design guidelines, and user-data context, designers approached the AI material in terms of its *experiential* traits. Engineers, who are technically trained, approached defining AI material in terms of *structural* traits, such as learning algorithms, model features, and model architecture. To bridge these differing viewpoints, teams engaged in rich discussions to ascribe material characteristics to the AI, co-create the application experience and evaluate its fit for end-users (i.e., the user experience). In this process, designers concretized their expectations for the AI material through ‘designerly’ representations, such as scenarios, mental-models, and wireframes. These *shareable* instantiations served as the designers’ *proxies* for their desired AI material characteristics. For engineers, these proxies offered human-centered requirements that allowed them to derive the AI material’s technical characteristics. We summarize our study findings in terms of (1) *designerly proxies* for articulating AI material needs based on human needs, (2) *data probes* to shape AI material design, and (3) role of *representational artifacts* as realizations of AIX.

4.1 Designerly Proxies for Articulating AI Material Needs

4.1.1 Material Properties: User Scenarios as a Proxy for Designing AI Behavior. Based on the protocol, teams started the design activity by exploring the intersections of user needs and AI strengths. By looking at the personas and their data (photos), the designers constructed different *scenarios* (user vignettes) to identify reasons behind photo clutter. These scenarios captured varied perspectives, including photo-taking (creation context), photo usage, and photos as memory artifacts (archive). For instance, in session 5, the designer (D5) constructed a scenario where a ‘Dad’ persona takes a burst of photos to capture his fidgety kids, intending to keep only the best one. Using such scenarios as anchors, the teams then explored how the AI might support decluttering. The designer then asked the engineer (E5) whether the AI could detect similar photos and identify the best one to keep. This was a *conversational* process consisting of “thinking out loud” about different scenarios, supplemented with annotations over the photos (Figure 1 b). The designer questioned the engineer about AI capabilities:

D5: ‘*Look at what this Dad is doing, he takes lots of photos of his kids and forgets to delete it, so this is one of the main challenges. Can AI identify duplicate photos and find the best one to keep?*’

E5: “*Yes, we can cluster the images based on similarity...*”

By constructing such user scenarios (i.e., material application), designers could ascribe potential capabilities for AI in the abstract and co-create the AI’s desirable behavioral characteristics (fabrication). Across all sessions, the scenarios led to instantiations of AI with different behaviors, including parsing text information, image quality assessment, and object

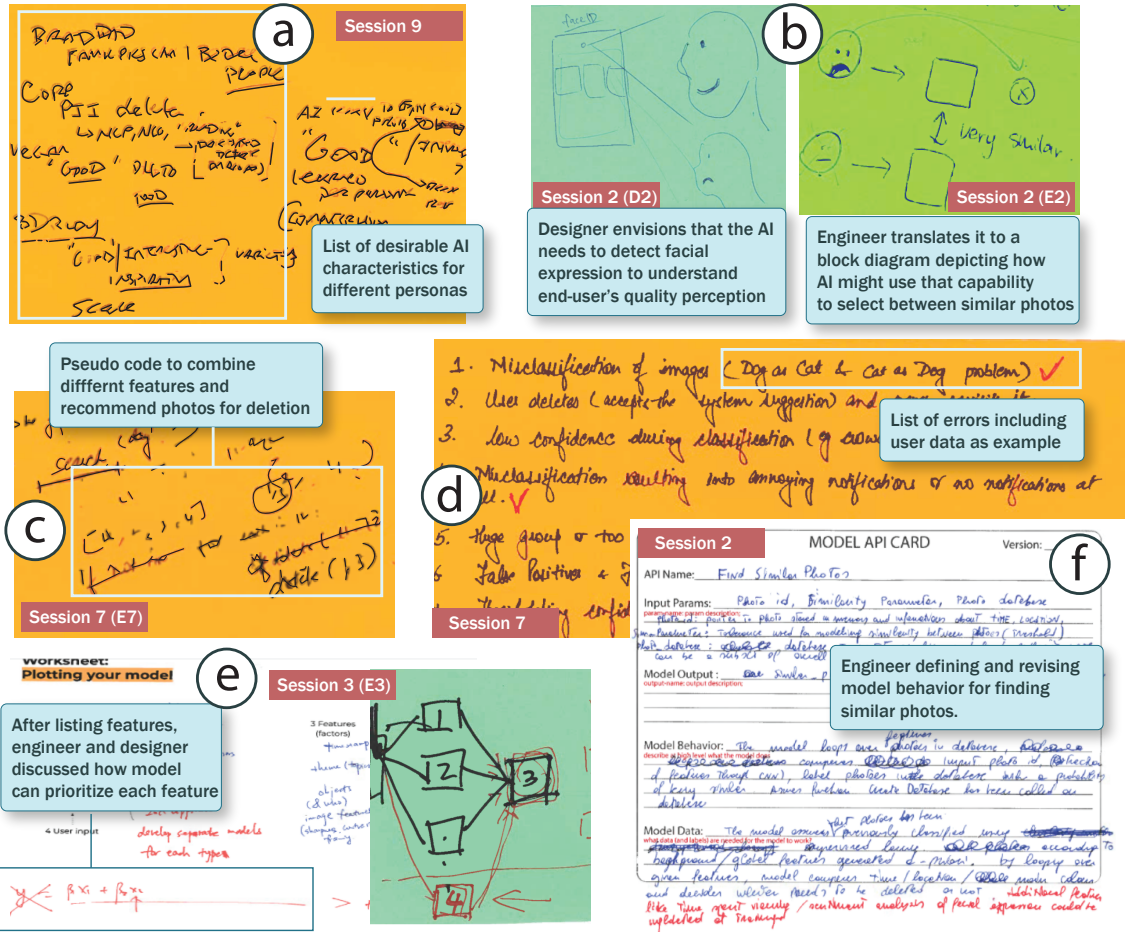


Fig. 3. Designerly proxies and AI material prototypes created by study participants.

recognition (Figure 3a). Further, in the course of defining AI behavior, designers would revise their initial scenarios to incorporate AI capabilities from engineers, thus creating “AI-infused” scenarios. As an example, D4 added to their vignette that the AI could intervene immediately after the person takes the photo: “After they take photos you wait until they turn off the phone, and then you have a dialog with the user: ‘Hey these photos, the eyes are shut.’...” (D4).

4.1.2 *Material Structure: End-User Mental Models as a proxy for Designing AI Implementation.* Design would be incomplete without human-centered considerations about the structure of AI material; that is, *How should the AI do what it is supposed to do?* In step 2 of the protocol, we observed that designers approached this issue by simulating in-depth data walkthroughs with previously defined scenarios [70]. During these walkthroughs, the team attempted to construct novel ‘mental models’ about how end-users might make judgments about decluttering their photos (i.e., which to keep and which ones to delete?), and what the AI can be expected to do for them. We call this an *expectation model* of the end-user. For instance, in session 2, the designer considered the moment immediately after taking a photo

and the end user's thought process for deciding whether to take a second (Figure 3b). This led to defining *logic* for whether a photo should be deleted: "... the Dad takes a photo of their kids and then views it for three seconds, which means it might be a good photo or a really, really bad photo and they want to improve it... there is some intention behind it... Can you understand from their facial expression [while looking at the photo] whether this is a good photo?" (D2).

During these walkthroughs, engineers listened and translated user expectations and decision factors into features, rules, and even pseudo-code for training the model (by writing it down in the model design template - Figure 3f). They would question designers about the importance of each feature to end-users and how the model can assign different weights to different features. For example, in session 3, the designer first talked about different attributes of a 'bad photo.' The engineer then visualized a linear model (See Figure 3e) to discuss ways to combine those different attributes:

D3: "Imagine these four similar pictures, but in this one, I cut his face just a little bit. Can AI identify that? Or if the eyes are closed, and then these are not useful..., and what about lighting or blurriness?"

E3: "From machine learning perspective $y = f(x)$... and each x can be a feature you put into a small model, and we can aggregate the outputs of each small model into a bigger model... you can run the same picture through each model and weigh the decision from each model whether to delete or not..."

From these examples, it is apparent that considering both expectation models and associated technical details helped participants co-create the model's structure.

4.1.3 Material Surface: User Interface as a Proxy for Designing AI's API. The third aspect of the AI material involves how people interact with it through its surface (i.e., the API). For AI, the API drives the user interface for end-user to engage with AI behavior. In our sessions, designers used the interface prototypes as a proxy when co-creating the model APIs. As with scenarios and walkthroughs, prototyping with concrete data points (as opposed to abstract placeholders such as 'lorem ipsum²') allowed designers to articulate specific API-level needs. For example, in session 8, the designer created the interface for viewing a set of recommended photos to delete. By examining this experience, they requested that the API also show key features about *why* the photo was recommended for deletion: "Let us think about the workflow, it is time to delete, how do you think they are going to process the photos to delete? Are they going to skim it by looking at the thumbnail, or enlarge it to focus on details? Can you provide a smart thumbnail with the features identified by the AI?" (D8). In their final design, they proposed 'smart-thumbnails' as the AI output. Here, the UI prototypes provided talk-backs for iterating on the material APIs. In most instances, engineers responded by revising the model API card or creating a new API version.

4.2 Data Probes to Shape AI Material Design

While designerly proxies offer a medium for articulating AI needs, design also requires generative thinking about alternative solutions. We observed that both the user-data associated with provided personas and participants' 'imagined' data (based on their prior knowledge) facilitated generative design processes. Specifically, participants used individual user-data points as *probes* to construct varied designerly proxies, explore the capabilities and limitations of AI, and evaluate created AI material against different HAI guidelines from the protocol. In human-centered design, design probes promote generative thinking and allow designers to explore the design space [58]. We found that end-user data (e.g., the photos associated with data personas) played the role of design probes in crafting the AI experience.

²https://en.wikipedia.org/wiki/Lorem_ipsum

4.2.1 Data probes for Divergent Thinking. In constructing user scenarios, participants constantly referred to personas and their photos to brainstorm different AI behavior types. In session 9, referencing the business traveler persona and their photo receipts, the team imagined a natural language understanding behavior to declutter old receipts. As the designer described, *“For this person who is trying to reimburse something, can we delete automatically by reading the text?”* (D9) E9 responds, *“Yes, we can identify text and what is in it, we can use natural language understanding...”* Besides, the data probes allowed participants to think about different ways to implement AI behavior through *mental models and implementation rules*:

D7: *“Something that came to me when I was looking at the personas was the emotional connections to these pictures like these pictures have values (pointing at pictures of kids). The application has to acknowledge the value for the users and save them instead. How can you classify pictures that have short term values and those that have long term value?”*

E7: *“On the back-end, what I hear is that there are different clusters of pictures, and I understand that different pictures have different value...but there is always a possibility of misclassification”*

D7: *“Could we have overarching rules, like faces might fall into personal attachment bin?”*

Data probes also allowed participants to explore various surface-level features like explainability and end-user feedback to the model. For example, in session 6, the designer did not want to display confidence scores to end-users. To this, the engineer used example data points to illustrate why a lack of explanation may result in distrust for end-users: *“Would the user be displeased if they took four photos and all four had bad lighting, but the system showed the photo with the highest score as this is the best photo because of lighting, but in actuality, they are all bad... Would that cause a loss in trust?”* (E6) Then the designer agreed *“That is a good point, maybe we have 4-5 categories [features], and you show a score underneath for each ..., so you know that all four photos are bad.”* (D6). Similarly, in all sessions, participants used data probes to determine the type of feedback that might be useful for model improvements. In session 3, the engineer explained that binary feedback of whether the recommendation was good or bad might not be sufficient for the model to improve. According to the engineer, *“How to use feedback is a difficult task for ML...the thing is users deciding to keep an image may not just be because the prediction is wrong, but because they just like it...for example in this picture the eyes are closed, but she looks cute... One way to design the interface is to have more than two buttons; in addition to keep or delete, if there could be a third button, like yes, eyes closed, but I still want it...”* (E3).

4.2.2 Data Probes for Exploring Material Boundaries and Limitations. A key aspect of designing AI is understanding the “edge cases,” where the AI might fail, or the designed behavior may not work. This understanding is essential for selecting the best alternative at design time (i.e., maximize potential for material appreciation) and designing for uncertainty during use, including failure, error handling, explainability, learnability, and setting expectations. In many instances, designers used data probes to decide whether to incorporate specific AI behaviors into the design. For instance, in session 4, the engineer suggested using smile detection to determine ‘goodness’ of the photo. To this, the designer pushed back by commenting: *“...if they had a stroke and are not smiling [in the picture], deleting that would be bad...”* (D4). In this case, they incorporated their understanding of the broader domain to think about examples beyond our data personas.

We also observed several instances in which engineers used example data points to highlight the limitations of AI. For instance, E4 commented, *“The model might predict that she is not smiling because of a missing tooth, or braces...”* E5 cautioned about AI limitations on classifying images of kids: *“For adults, facial recognition works well, but I am not*

sure whether it will work for kids. These all may be photos of the same kids, but I mean kids grow... the clustering may not work.” These examples helped participants determine whether the user experience with AI should be automated or assistive: “For duplicate photos, I think it should be assistive and complementary... if the user deletes one photo, we can say here are other photos like this, would you like to delete those as well?” (E10).

4.2.3 Data Probes for Design Evaluation. In addition to exploring material qualities, data probes also supported design convergence and confirmatory evaluation. Similar to identifying edge cases, participants made use of data probes to check whether a behavior ‘scaled’ across different personas. For example, in session 10, participants identified the desired AI behavior of detecting duplicates and selecting the best photo to keep. The designer then considered an Instagrammer persona for testing whether the same behavior might be useful to them. According to D10: “Same goes here... after they upload to Instagram they probably do not need it. Especially for Instagram people, they edit their photos a lot, and each time they edit, it will create new copies if it...”. Similarly, after making decisions about *not* including certain behaviors in their design, teams discussed why that decision was right by considering other personas. For instance, in session 9, they rejected their initial idea about predicting images’ temporal utility. In acknowledging their decision, the engineer E9 commented: “I am somewhat nervous these [frequency] metrics... for instance, a big failure in which case is deleting a baby’s birth picture, that would be bad...”

4.3 Representational Artifacts as Realizations of the AI Material

Prototyping with user-data allowed designers to experience the AIX ‘first-hand’ as they were creating it—material appreciation. It also allowed them to concretely communicate their use of the created AI material back to the engineers by “representing it to make the solution transparent [76]”. Given the simplicity of the design brief, we were surprised to see the wealth of design variations generated across the sessions. As shown in Figure 4, the entry points for the declutter UX included explicitly invoking the AI via a button click, use of an AI-triggered notification, end-user (delete) actions, providing seed images as a starting point, and conversational dialog. Participants also designed various AI-powered presentation views, including intelligent clustered lists, comparison views highlighting feature differences, interactive thumbnail views, etc. Across all prototypes, the selection of images was driven by imagining AI-infused scenarios and using it to ‘play-out’ the experience as it was created. Here, we discuss three key benefits to AI material prototyping, including: (1) aligning AI and UI through translation and feedback, (2) addressing misconceptions and gaps in understandings about AI, and (3) helping designers perceive the complex nature of the AI material.

4.3.1 Aligning AI and UI through Translation and Feedback: In constructing AIX prototypes, designers had to translate their understanding of the AI model’s structure and APIs into their own knowledge of user interface design. Drawing from their design expertise, they began by representing both the AI behavior and output through familiar design patterns. In many sessions, this provided a starting point for envisioning AIX. In session 2, to recommend to users which photos to delete, the designer started with a familiar photo album interface: “How to separate good one from all the junk ones? In Apple, they show you all the photos, and you select which ones you want to keep or not...” (D2). Similarly, in session 4, the designer started with their familiarity about ‘snack-bar’ UI (i.e., a notification with quick action buttons) to prototype an ‘assistive’ AI experience: “... they have what is called a snack bar... where it is a confirmation message, but within the confirmation message you can have a like an undo button, that is really common now...” (D4).

During the prototyping process, designers concretized their understanding of the model by considering model input and feedback controls within the context of user-data. Their choice of text (such as labels and dialog) corresponded with the AI model implementation rules and features; e.g., in session 4 prototype: “it looks like it’s a little blurry, do you

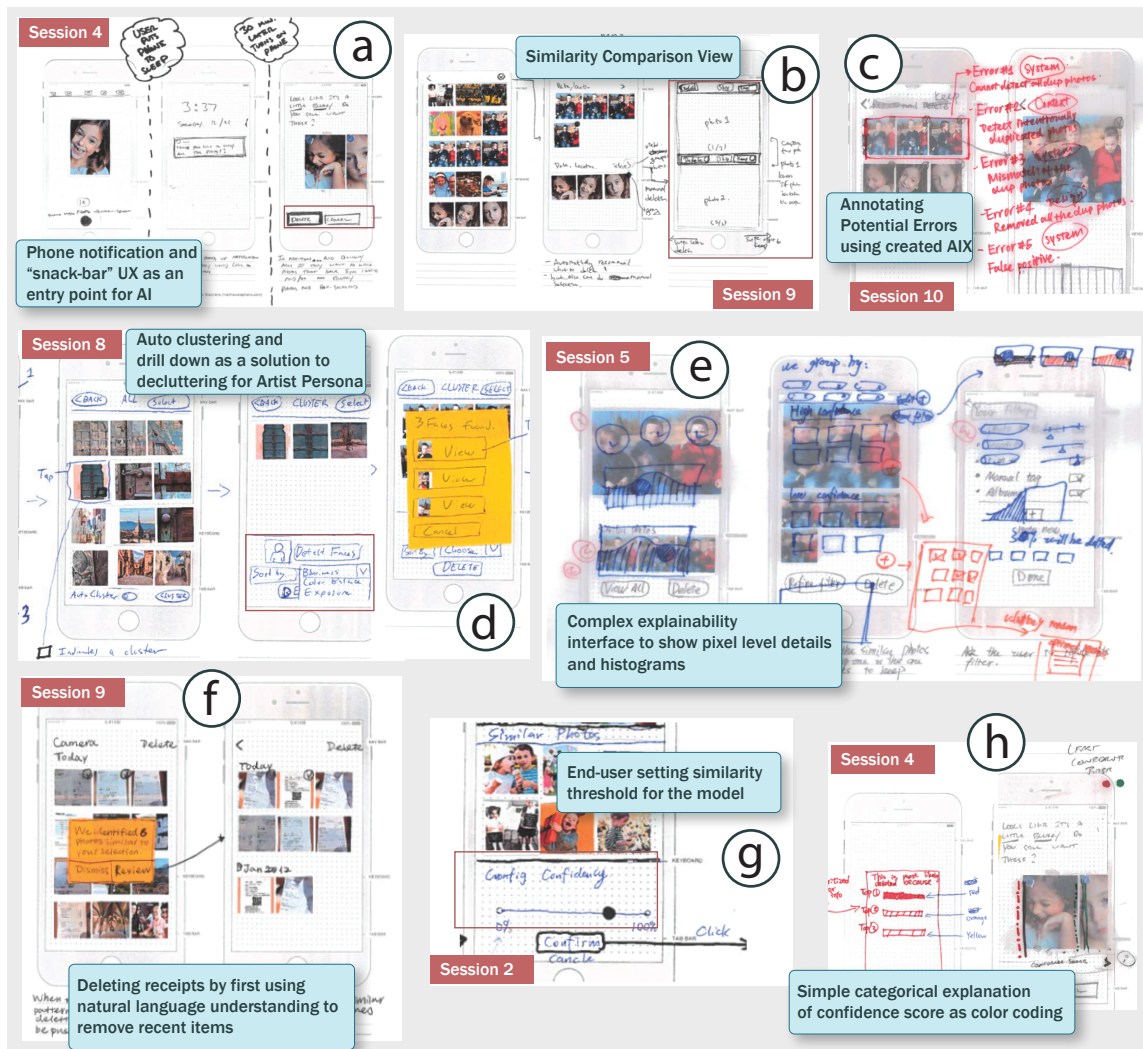


Fig. 4. AIX Representational Artifacts showing interface design solutions including explainability, and model feedback.

want to keep this?”. We also observed that designers intuitively designed certain presentation features that led to team discussions and subsequent changes to the model API. In session 4, the designer D4 decided to bin confidence scores into higher-level categories and color code them in photos (Figure 4h). Their rationale was to make it more accessible to end-users who may not understand the meaning of differences in confidence scores. This design decision prompted a discussion with the engineer about how the model might categorize confidence scores for a more intuitive presentation:

D4: “you can give them like some sort of color highlight, so you can give them a threshold...”

E4: “That’s a good point, you can order them by confidence, and you can give them a threshold... we can test it out by having a training dataset and test dataset, and we can say from our average, a typical use case the

most evenly divided percentages are for example anything below 50 is low, 50 to 75% is medium, anything above is high confidence.”

The material prototype allowed the team to anticipate user needs and redesign the AI material to address them.

4.3.2 Addressing Misconceptions and Gaps in Understanding. Representational prototypes allowed both designers and engineers to identify gaps in each others’ understanding. Through discussions and annotation overlays (using vellum paper), they negotiated differences in their understanding of the AI material. For instance, in session 8, the designer had prototyped detailed text explanations about why a set of images were clustered together. Looking at the prototype, the engineer commented: *“I am not sure technically machines are capable of that... people are good at generating semantic explanations, for example people can tell others in natural language why these photos are similar or why they are clustered together... But like I am not sure even state of the art ML models are capable of that”* (E8). By contrast, in session 9, the AI engineer suggested using personalized sorting algorithms to present the photos to be deleted. The designer then annotated onto the prototype to clarify that sorting should be objective, and subjective sorting would likely make the user not trust the recommendations.

E9: *“We can start from a fixed equation and adjust parameters based on user interactions..., for example, this image has more exposure than others, then we can personalize the sorting algorithm by changing parameters.”*

D9: *“This is supposed to be an objective way of sorting the value, so if I knew that this sorts blurriness [order photos by blur level] to my choice, I would not trust it anymore.”*

4.3.3 Perceiving AI Complexities. The act of constructing an AI material prototype made designers more aware of the complexities of AI-powered interfaces. Across many sessions, by looking at material flaws (e.g., material uncertainties that do not communicate rationale to end-users), designers discovered additional needs, such as setting default model parameters, feedback features, explainability, setting expectations, etc. Working through several iterations on their design, the designer in session 6 commented: *“Now we are ‘frankensteining’ this sucker, but you can have a settings page that pops open over here and allows them to say, ‘keep top N photos’...”*

While the co-creation process clarified the underlying material structure for designers, some found it hard to separate from their increased knowledge about the AI in order to design the UX, or to communicate their AI understanding to end-users through their design. We observed instances in which the user-to-AI feedback mechanism became very complicated, reflecting the designer’s understanding of the AI model but failing to use that understanding in designing the UX. For example, in session 5, the designer and engineer engaged in rich discussions about categorization failures based on pixel-level features and personalization. In applying this new understanding, the designer created a very complicated user interface prototype without considering trade-offs between improving AI model accuracy and the user effort required for feedback (see Figure 4e). These examples demonstrate changes in conceptualizations of AI material arising through the co-design of AIX.

5 DISCUSSION

Design materials are central to design processes. In conventional UX design, the graphical user interface (GUI) is the prevalent “design material” that every UI designer understands [14, 24]. If AI were like any other material defined by nature or convention, designers would learn how to work with its given properties to generate design solutions for human users. However, AI resists this approach. Had we provided participants a ‘created’ AI material (e.g., a closed-box ML model that assigned a quality score to each image [82]), they likely would not have produced the range of expressive,

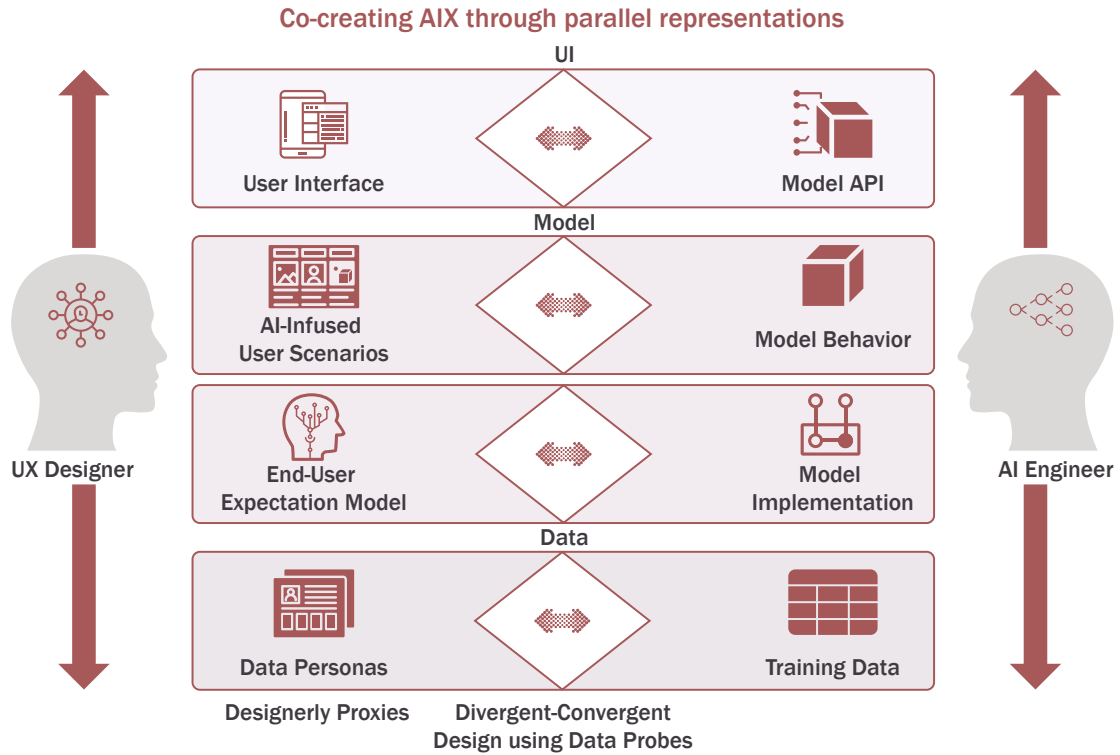


Fig. 5. A process model for co-creating AIX.

human-centered designs that we observed. Instead, as our findings show, AI material must be defined by investigating the human user's envisioned experiences. In our study, we aimed to answer: (1) How do designers and engineers conceptualize design guidelines from AI and UX perspectives? (2) How do they co-create the design and technical characteristics of AI materials; and (3) What representations are invented during this process? Based on our findings, we respond to these questions by proposing a process model for co-creating AIX and by reflecting on the role of end-user data as a design probe for generative design thinking. Through this discussion, we offer design considerations for data probes within AIX design tools.

5.1 Towards a Process Model for Co-Creating AIX

In the study, our teams displayed a design progression occurring across a *spectrum* of materiality. Teams started the design activity by exploring the intersections of user needs and AI strengths. First, with wholly *imaginary and abstract* material, the team worked together to envision the AI. The designers negotiated this initial form-giving by constructing user scenarios, allowing them to approach AI through its potential capabilities without detailed engineering knowledge. This is similar to the role of visualization as proxies in designing 'immaterial' materials [4]. Designers used personas, data points, vignettes, and user scenarios—their designerly proxies—to create initial instantiations of the desired AI behavior. These proxies offered design representations as abstractions of *planned behaviors* that allowed engineers to

define the technical characteristics of the AI material. In the course of co-designing AI behavior, designers also revised their initial scenarios to incorporate new AI capabilities, thus creating “AI-infused” scenarios.

Equipped with an invented form, teams moved forward to *enact and specify* the AI material. Following human-centered walkthroughs of scenarios, the designers constructed novel *expectation models* capturing how the AI’s end-users might make judgments. The engineers translated these identified user expectations and decision factors into features and rules for training the AI. This co-creation process led to discussions about the attributes, priorities, and values important to users and the technical capabilities the AI needed to support them. While the AI material has taken on behavior and structural characteristics, its envisioned design was only fully *realized* through the team’s use of material prototypes. Designers used interface prototypes as a proxy when co-designing the model’s inputs and outputs (i.e., the material’s surface). These designerly proxies allowed the team to align the AI and UI through translation and feedback. Identifying specific AI and human behaviors allowed the evaluation of material flaws, misconceptions, and scalability issues. This is similar to *Replay Enactments* [40] that use authentic data to make complex system behavior tangible to designers. Only at this late stage could the full scope of the AI material and UX design be made visible in its interactive complexity.

Clearly, the *dynamic* nature of AI material is unlike other design materials; consequently, the design process for AIX differs from standard design approaches. In conventional human-centered design (HCD), like the double diamond framework [16], the design process is linear or top-down. Designers mainly work at the user interface layer to specify the end-user experience. They then hand-off the created specifications to engineers to build [75]. However, when designing AI experiences, design extends *beyond* the interface and into the design of AI components, including the model’s behavior, learning characteristics, assumptions, and nature of training data. UX professionals lack the means to engage in designing these AI components. Instead, current AI development workflows take an “AI-first” approach in which the AI material is created before envisioning its use. Such an approach is problematic because any changes to align the AI’s properties to human needs will require costly rework (e.g., addressing disparities in gender classification [12]). Further, there are instances in which the AI behavior itself does not align with human needs, values, and concerns (e.g., using facial recognition to expose political orientation [52]).

In order to address these issues, we need a process in which AI material creation and its application experience design can happen in parallel through iteration and feedback. As shown in Figure 5, we propose a process model that combines top-down (UX-first) and bottom-up (AI-first) workflows to distribute agency between designers and engineers. As represented by the *bidirectional* arrows in our model, the AI and UX components are designed in parallel, a critical insight from our study. Our approach shifts engineers’ mindsets towards more proactive engagement through accessible user-data proxies and data probes during the co-creation process. Designers engage in co-creating AI behavior without technical roadblocks, operationalize HAI guidelines, and reduce time to feedback (a concern with AI design [91]). Our model’s parallel process affords immediate feedback for both material creation and design, obviating the significant rework costs (from collecting and training with new training dataset, retraining the models, etc.) when even small changes arise later. However, we also emphasize the “towards” in the paper’s title; Our study lays groundwork for a collaborative process to align AI’s form and its function in the early stages of design. Future research should build on this parallel process model to investigate specific data and representation needs across different application domains and AI capabilities.

5.2 Role of Data Probes in AIX Design

Data probes served as a “content common ground” for designers and engineers to collaborate across the different stages in the process model. A characterization of this collaboration is that designers are immediate consumers of AI materials. Their objective is to ensure that the material specification meets their end-users’ UX needs. Using data probes, designers advocated for end-users during the material creation process and simultaneously tested the AI material under construction. Similar to [72], end-user data offered necessary grounding for designers to advocate for centering people in the design of AI, including its behavior, implementation, and APIs.

As shown in Figure 5, each of the parallel stages in our model involved both divergent and convergent processes. Designers and engineers ideated on UX needs and AI capabilities together, and they mutually constrained convergence towards a design solution. User data as probes played a critical role in this divergent-convergent process of creating the AI experience. We can extend the material metaphor and borrow from the language of physical material design to characterize the role of data probes: (1) data *molds*, (2) data *vulcanizers*, and (3) data *coupons*. In the early stage of the study, data probes functioned as “molds” for AI’s initial form-giving. By constructing AI-infused scenarios with data probes, designers and engineers explored different forms the AI could potentially take in supporting the declutter experiences. After identifying the initial form, designers used data probes to define the AI material’s internal properties. By constructing expectation models with data probes, designers and engineers “solidified” the AI’s implementation requirements. This step in AI material creation is analogous to ‘vulcanization’ chemicals in the rubber manufacturing process to solidify its internal structure. Finally, by constructing AIX interface representations with real data, designers produced coupons (test samples) of the material to assess the AI experience. In traditional material design processes, ‘coupons’ are samples of the material used to test its properties at a small scale (e.g., [43]). The designers’ mixed-fidelity prototypes served as coupons to test the AI material and address gaps in the desired AI experience.

5.3 Design Considerations for AIX Design Tools

Prototyping is an essential step in software development [75]. Through iterative prototyping, teams incorporate increasing details to define different software aspects [6]. The mixed-fidelity approach in our study is an initial step to iterative prototyping. As details increase, teams need to increase the fidelity of their prototypes as well. In this regard, UX design tools should escape the “closed-box” view to make AI more accessible and transparent to designers [92]. Beyond helping designers understand AI (i.e., educational goals), designers should be able to work with AI material during AIX design. The insights from our study suggest that data probes offer useful design considerations for this goal. With this in mind, we offer a set of design considerations for incorporating data probes into design tools.

Support for creating data probes: In the current study, we constructed the data for each persona to include a variety of solution alternatives. Participants also imagined their own additional data points during the design process. Design tools should allow designers to incorporate data from user research into AIX design processes directly. This could include data collected from participants (similar to Wizard-of-Oz prototyping [9]), through mixed methods persona creation (e.g., Data-Assisted Affinity Diagramming [78]), or from dedicated data collection and annotation pipelines [28]. In addition, tools should support accessible ways to generate user data with desired properties. In the data visualization community, tools exist to create datasets with desirable statistical properties (e.g., [31, 57]), allowing designers to select charts to fit their data needs. Our study teams imagined varied data—blurry images, variations in size, and time-progression photos—using their understanding of the task set within use contexts. Design tools should support such expressive ‘queries’ to find or generate *just-in-time* data probes for designers.

Support for interactive AI & UX design workflows: To work with the AI material under construction, designers can use data probes to receive “talk-back” from the AI design workflow. Currently, end-user machine learning tools such as RunwayML [63] allow novices to interact with machine learning models and for visual exploration of machine learning behavior (e.g., the What-If Tool [30]). However, these tools do not support the use of data probes for divergent thinking about AI behavior. To be effective, probes should be integrated into generative prototyping workflows to provide input and feedback to designers. Moreover, interactions using data probes allow designers to propose desired outputs based on human needs. This can be in the form of ground truth data, annotated labels, output format, etc. For instance, designers can curate a set of diverse data points and ground truth outputs and compare the working model’s output against ground-truth values.

Support for constructing designerly proxies: Currently, when prototyping UIs, designers typically work with static placeholder content. As with data visualization, in AI, “data changes everything [86].” Prototyping tools should allow designers to construct AIX design candidate representations by incorporating data probes and AI material talk-backs (e.g., [51]). This allows consideration of alternative choices [56], along with design for AI uncertainty through explainability, learnability, and edge-case analysis[3].

Support for communication during co-creation: A final consideration for co-creating AIX is to share intermediate proxies of AIX, including scenarios, mental models, and interface prototypes with AI engineers. For engineers, they need to offer descriptions of AI properties, assumptions, learning rules, and API details back to designers. This is fundamentally different from standard workflows in which designers primarily share final design specifications with engineers. Both UX design tools and AI creation tools should incorporate features to import, translate, and share intermediate design proxies. Similar to the transparent vellum paper during our study, digital tools should support annotation overlays, generate new examples with different data probes, and communicate failures and constraints (i.e., explainability for designers).

5.4 Limitations and Future Work

Our study’s design problem was to define an AI-powered experience for decluttering photo albums. While this simple problem allowed us to observe co-creation processes in an accessible domain, other data types may be more challenging for co-creation. More complex design problems may be difficult to represent with pen-and-paper approaches. We plan to conduct co-design sessions with other data types and problem domains to iterate on the process model. We aim to assess (and address) the fit and shortcomings of our process model for different AIX problems through these sessions. Second, we provided participants the data probes for use in our study. We plan to investigate inclusive and participatory approaches to creating diverse data probes for AIX design. Third, our protocol demonstrates a low-cost, rapid prototyping approach to co-creating AIX. As described, this is the first step to iterative prototyping with increasing levels of fidelity. We are currently exploring high-fidelity prototyping tools for AIX to understand how teams might continue to evolve their designs. As an example, we developed ProtoAI [79] to allow designers to invoke models and services with concrete data during prototyping. Fourth, we recruited participants with prior experience in designing AI applications. In many industries, both designers and engineers are new to AI [13]. We are now using the study protocol in classrooms to teach design and engineering students about AIX design. Through these efforts, we will investigate the types of training and scaffolding designers need to effectively participate in AIX’s rapid prototyping. Finally, future work should investigate how other stakeholders, including domain experts, representative end-users, and data analysts, might participate in the AIX co-creation process.

6 CONCLUSION

Treating technology as a design material encourages designers to explore its properties for UX design. However, when working with AI as design material, neither a form-follows-function nor a function-follows-form approach is practical. Instead, the AI material and its application UX need to be co-created through collaboration between designers and AI engineers. In this work, we investigated such an approach by conducting an in-lab design study with ten pairs of designers and engineers. Our protocol combines a vertical prototyping approach with talk-backs from AI and UX to facilitate co-creation. We identified the crucial role of end-user data as a tool for co-creating AI design material. By using data probes, designers were able to construct designerly proxies and specify material needs for AI. Data probes facilitated divergent thinking, material testing, and design validation. Based on these findings, we propose a process model for collaborative AIX design and offer considerations for incorporating data probes in AIX design tools.

ACKNOWLEDGMENTS

We thank our reviewers and study participants for their time and feedback. We also thank Ken Holstein, Steven Drucker, and Steve Oney for their inputs on the paper.

REFERENCES

- [1] Christopher Alexander. 1964. *Notes on the Synthesis of Form*. Vol. 5. Harvard University Press.
- [2] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. 2019. Software engineering for machine learning: A case study. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 291–300.
- [3] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 3.
- [4] Timo Arnall. 2014. Exploring ‘Immaterials’: Mediating Design’s Invisible Materials. *International Journal of Design* 8, 2 (2014).
- [5] Jan Auernhammer. 2020. Human-centered AI: The role of Human-centered Design Research in the development of AI. (2020).
- [6] Michel Beaudouin-Lafon and Wendy E Mackay. 2009. Prototyping tools and techniques. In *Human-Computer Interaction*. CRC Press, 137–160.
- [7] Eli Blevis, Youn-kyung Lim, Erik Stolterman, et al. 2006. Regarding software as a material of design. In *Proceedings of Design Research Society International Conference*. 1–18.
- [8] Barry Boehm. 2000. Requirements that handle IKIWISI, COTS, and rapid change. *Computer* 33, 7 (2000), 99–102.
- [9] Jacob T Browne. 2019. Wizard of Oz Prototyping for Machine Learning Experiences. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [10] Louis L Bucciarelli. 2002. Between thought and object in engineering design. *Design studies* 23, 3 (2002), 219–231.
- [11] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, et al. 2013. API design for machine learning software: experiences from the scikit-learn project. *arXiv preprint arXiv:1309.0238* (2013).
- [12] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [13] Carrie J Cai and Philip J Guo. 2019. Software Developers Learning Machine Learning: Motivations, Hurdles, and Desires. In *2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 25–34.
- [14] John M Carroll, Robert L Mack, and Wendy A Kellogg. 1988. Interface metaphors and user interface design. In *Handbook of human-computer interaction*. Elsevier, 67–85.
- [15] M Ceconello, D Spallazzo, and M Sciannamè. 2019. Design and AI: prospects for dialogue. (2019).
- [16] Design Council. 2005. The ‘double diamond’ design process model. *Design Council* (2005).
- [17] Henriette Cramer and Juho Kim. 2019. Confronting the tensions where UX meets AI. *interactions* 26, 6 (2019), 69–71.
- [18] Kristin N Dew and Daniela K Rosner. 2018. Lessons from the woodshop: Cultivating design with living materials. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [19] Dennis P Doordan. 2003. On materials. *Design Issues* 19, 4 (2003), 3–8.
- [20] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 278–288.
- [21] Malin Eiband, Hanna Schneider, Mark Bilandzic, Julian Fazekas-Con, Mareike Haug, and Heinrich Hussmann. 2018. Bringing transparency design into practice. In *23rd international conference on intelligent user interfaces*. 211–223.

- [22] Daniel Engelberg and Ahmed Seffah. 2002. A framework for rapid mid-fidelity prototyping of web sites. In *IFIP World Computer Congress, TC 13*. Springer, 203–215.
- [23] Jennifer Fereday and Eimear Muir-Cochrane. 2006. Demonstrating rigor using thematic analysis: A hybrid approach of inductive and deductive coding and theme development. *International journal of qualitative methods* 5, 1 (2006), 80–92.
- [24] Ylva Fernaeus and Petra Sundström. 2012. The material move how materials matter in interaction design research. In *proceedings of the designing interactive systems conference*. 486–495.
- [25] David Flink. 2020. The Wire: Your AI-Powered ‘To Do’ list. <https://people.ai/blog/the-wire-your-ai-powered-to-do-list/>
- [26] Elisa Giaccardi and Elvin Karana. 2015. Foundations of materials experience: An approach for HCI. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2447–2456.
- [27] Yolanda Gil, James Honaker, Shikhar Gupta, Yibo Ma, Vito D’Orazio, Daniel Garijo, Shruti Gadewar, Qifan Yang, and Neda Jahanshad. 2019. Towards human-guided machine learning. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 614–624.
- [28] Fabien Girardin and Neal Lathia. 2017. When User Experience Designers Partner with Data Scientists. In *2017 AAAI Spring Symposium Series*.
- [29] Google. 2019. People + AI Guidebook. <https://pair.withgoogle.com/>
- [30] Google. 2020. Visually probe the behavior of trained machine learning models, with minimal coding. <https://pair-code.github.io/what-if-tool/>
- [31] Robert Grant. 2020. DrawMyData a tool for teaching stats and data science. <http://robertgrantstats.co.uk/drawmydata.html>
- [32] Thilo Hagendorff. 2019. The ethics of AI ethics—an evaluation of guidelines. *arXiv preprint arXiv:1903.03425* (2019).
- [33] Lise Amy Hansen. 2011. Full-body movement as material for interaction design. *Digital Creativity* 22, 4 (2011), 247–262.
- [34] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850.
- [35] Michael Heidt, Andreas Bischof, and Paul Rosenthal. 2014. Deconstructivist Design within HCI. In *International Conference of Design, User Experience, and Usability*. Springer, 115–122.
- [36] Joseph M Hellerstein, Vikram Sreekanti, Joseph E Gonzalez, James Dalton, Akon Dey, Sreyashi Nag, Krishna Ramachandran, Sudhanshu Arora, Arka Bhattacharyya, Shirshanka Das, et al. 2017. Ground: A Data Context Service.. In *CIDR*.
- [37] Karey Helms, Barry Brown, Magnus Sahlgren, and Airi Lampinen. 2018. Design methods to investigate user experiences of artificial intelligence. In *AAAI Spring Symposium Series, March 26–28, 2018*. Association for the Advancement of Artificial Intelligence.
- [38] Charles Hill, Rachel Bellamy, Thomas Erickson, and Margaret Burnett. 2016. Trials and tribulations of developers of intelligent systems: A field study. In *2016 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 162–170.
- [39] Lars Erik Holmquist. 2017. Intelligence on tap: artificial intelligence as a new design material. *interactions* 24, 4 (2017), 28–33.
- [40] Kenneth Holstein, Erik Harpstead, Rebecca Gulotta, and Jodi Forlizzi. 2020. Replay Enactments: Exploring Possible Futures through Historical Data. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 1607–1618.
- [41] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [42] Daneil Hough. 2020. Interfake: Quick JSON APIs. <https://github.com/basiclydan/interfake>
- [43] Yuner Huang and Ben Young. 2014. The art of coupon tests. *Journal of Constructional Steel Research* 96 (2014), 159–175.
- [44] Thomas L Huber, Maïke AE Winkler, Jens Dibbern, and Carol V Brown. 2019. The use of prototypes to bridge knowledge boundaries in agile software development. *Information systems journal* (2019).
- [45] Apple Inc. 2019. Designing the UI and User Experience of a Machine Learning App. <https://developer.apple.com/design/human-interface-guidelines/machine-learning/overview/introduction/>
- [46] Yuhua Jin, Isabel Qamar, Michael Wessely, and Stefanie Mueller. 2020. Photo-Chromeleon: Re-Programmable Multi-Color Textures Using Photochromic Dyes. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Emerging Technologies*. 1–2.
- [47] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [48] Heekyoung Jung and Erik Stolterman. 2012. Digital form and materiality: propositions for a new approach to interaction design research. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*. 645–654.
- [49] Elvin Karana, Bahareh Barati, Valentina Rognoli, Anouk Zeeuw Van Der Laan, et al. 2015. Material driven design (MDD): A method to design for material experiences. (2015).
- [50] Claire Kayacik, Sherol Chen, Signe Noerly, Jess Holbrook, Adam Roberts, and Douglas Eck. 2019. Identifying the intersections: User experience+ research scientist collaboration in a generative machine learning interface. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [51] Scott R Klemmer, Anoop K Sinha, Jack Chen, James A Landay, Nadeem Aboobaker, and Annie Wang. 2000. Suede: a Wizard of Oz prototyping tool for speech user interfaces. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*. 1–10.
- [52] Michal Kosinski. 2021. Facial recognition technology can expose political orientation from naturalistic facial images. *Scientific reports* 11, 1 (2021), 1–7.
- [53] Hanna Landin. 2005. Fragile and magical: Materiality of computational technology as design material. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*. 117–120.
- [54] Germán Leiva, Nolwenn Maudet, Wendy Mackay, and Michel Beaudouin-Lafon. 2019. Enact: Reducing designer–developer breakdowns when prototyping custom interactions. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 3 (2019), 1–48.

- [55] Rikard Lindell. 2012. Code as Design Material. In *Participatory Materialities Workshop and Symposium at Aarhus University*.
- [56] Allan MacLean, Richard M Young, Victoria ME Bellotti, and Thomas P Moran. 1991. Questions, options, and criteria: Elements of design space analysis. *Human-computer interaction* 6, 3-4 (1991), 201–250.
- [57] Miro Mannino and Azza Abouzied. 2019. Is this Real? Generating Synthetic Data that Looks Real. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 549–561.
- [58] Tuuli Mattelmäki et al. 2006. *Design probes*. Aalto University.
- [59] David Maulsby, Saul Greenberg, and Richard Mander. 1993. Prototyping an intelligent agent through Wizard of Oz. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. 277–284.
- [60] Michael McCurdy, Christopher Connors, Guy Pyrzak, Bob Kanefsky, and Alonso Vera. 2006. Breaking the fidelity barrier: an examination of our current characterization of prototypes and an example of a mixed-fidelity success. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 1233–1242.
- [61] Microsoft. 2020. Start Designing and Prototyping. <https://docs.microsoft.com/en-us/windows/mixed-reality/design>
- [62] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [63] Runway ML. 2018. Runway ML: Machine Learning for Creators. <https://runwayml.com/>
- [64] Michael Nebeling, Janet Nebeling, Ao Yu, and Rob Rumble. 2018. Protoar: Rapid physical-digital prototyping of mobile augmented reality applications. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [65] William Odom and Tijds Duel. 2018. On the design of OLO Radio: Investigating metadata as a design material. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–9.
- [66] Gregory B Olson. 2000. Designing a new material world. *Science* 288, 5468 (2000), 993–998.
- [67] Oracle. 2020. Powerful API Design Stack. <https://apiary.io/>
- [68] Kayur Patel. 2010. Lowering the barrier to applying machine learning. In *Adjunct proceedings of the 23rd annual ACM symposium on User interface software and technology*. 355–358.
- [69] Nadia Piet. 2019. AI Meets Design. <http://aimeets.design/>
- [70] Peter G Polson, Clayton Lewis, John Rieman, and Cathleen Wharton. 1992. Cognitive walkthroughs: a method for theory-based evaluation of user interfaces. *International Journal of man-machine studies* 36, 5 (1992), 741–773.
- [71] Johan Redström. 2005. On technology as material in design. *Design Philosophy Papers* 3, 2 (2005), 39–54.
- [72] Yim Register and Amy J Ko. 2020. Learning Machine Learning with Personal Data Helps Stakeholders Ground Advocacy Arguments in Model Mechanics. In *Proceedings of the 2020 ACM Conference on International Computing Education Research*. 67–78.
- [73] Erica Robles and Mikael Wiberg. 2010. Texturing the "material turn" in interaction design. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*. 137–144.
- [74] Donald Schon and John Bennett. 1996. Reflective Conversation with Materials in Bringing Design to Software, Winograd T.
- [75] Ahmed Seffah, Jan Gulliksen, and Michel C Desmarais. 2005. *Human-centered software engineering-integrating usability in the software development lifecycle*. Vol. 8. Springer Science & Business Media.
- [76] HA Simon. 1969. *The sciences of the artificial* (1966; Orig. ed. 1969; 2nd, rev).
- [77] David Canfield Smith. 1975. *Pygmalion: a creative programming environment*. Technical Report. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE.
- [78] Hariharan Subramonyam, Steven M Drucker, and Eytan Adar. 2019. Affinity Lens: Data-Assisted Affinity Diagramming with Augmented Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [79] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. ProtoAI: Model-Informed Prototyping for AI-Powered Interfaces. In *26th International Conference on Intelligent User Interfaces*. 48–58.
- [80] Petra Sundström, Alex S Taylor, and Kenton O'Hara. 2011. Sketching in software and hardware Bluetooth as a design material. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. 405–414.
- [81] Ivan E Sutherland. 1964. Sketchpad a man-machine graphical communication system. *Simulation* 2, 5 (1964), R–3.
- [82] Hossein Talebi and Peyman Milanfar. 2018. NIMA: Neural image assessment. *IEEE Transactions on Image Processing* 27, 8 (2018), 3998–4011.
- [83] Sandeep Tata, Alexandrin Popescu, Marc Najork, Mike Colagrosso, Julian Gibbons, Alan Green, Alexandre Mah, Michael Smith, Divanshu Garg, Cayden Meyer, et al. 2017. Quick access: building a smart experience for Google drive. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1643–1651.
- [84] Anna Vallgård and Johan Redström. 2007. Computational composites. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 513–522.
- [85] Philip van Allen. 2018. Prototyping ways of prototyping AI. *interactions* 25, 6 (2018), 46–51.
- [86] Jagoda Walny, Christian Frisson, Mieka West, Doris Kosminsky, Søren Knudsen, Sheelagh Carpendale, and Wesley Willett. 2019. Data changes everything: Challenges and opportunities in data visualization design handoff. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 12–22.
- [87] 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*. 1–15.

- [88] Mikael Wiberg. 2016. Interaction, new materials & computing—Beyond the disappearing computer, towards material interactions. *Materials & design* 90 (2016), 1200–1206.
- [89] Qian Yang. 2017. The role of design in creating machine-learning-enhanced user experience. In *2017 AAAI Spring Symposium Series*.
- [90] Qian Yang. 2018. Machine learning as a UX design material: How can we imagine beyond automation, recommenders, and reminders?. In *AAAI Spring Symposia*.
- [91] Qian Yang, Justin Cranshaw, Saleema Amershi, Shamsi T Iqbal, and Jaime Teevan. 2019. Sketching NLP: A Case Study of Exploring the Right Things To Design with Language Intelligence. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [92] Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating how experienced UX designers effectively work with machine learning. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 585–596.
- [93] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
- [94] Zhibin Zhou, Lingyun Sun, Yuyang Zhang, Xuanhui Liu, and Qing Gong. 2020. ML Lifecycle Canvas: Designing Machine Learning-Empowered UX with Material Lifecycle Thinking. *Human-Computer Interaction* (2020), 1–25.