"It Is Not Always Discovery Time" : Four Pragmatic Approaches in Designing AI Systems

'It Is Not Always Discovery Time': Four Pragmatic Approaches in Designing AI Systems

Maximiliane Windl
LMU Munich
Munich, Germany
maximiliane.windl@ifi.lmu.de

Sebastian S. Feger
LMU Munich
Munich, Germany
sebastian.feger@ifi.lmu.de

Lara Zijlstra
Utrecht University
Utrecht, Netherlands
l.e.zijlstra@uu.nl

Albrecht Schmidt
LMU Munich
Munich, Germany
albrecht.schmidt@ifi.lmu.de

Paweł W. Woźniak
Chalmers University of Technology
Gothenburg, Sweden
pawel.wozniak@chalmers.se

ABSTRACT

While systems that use Artificial Intelligence (AI) are increasingly becoming part of everyday technology use, we do not fully understand how AI changes design processes. A structured understanding of how designers work with AI is needed to improve the design process and educate future designers. To that end, we conducted interviews with designers who participated in projects which used AI. While past work focused on AI systems created by experienced designers, we focus on the perspectives of a diverse sample of interaction designers. Our results show that the design process of an interactive system is affected when AI is integrated and that design teams adapt their processes to accommodate AI. Based on our data, we contribute four approaches adopted by interaction designers working with AI: a priori, post-hoc, model-centric, and competence-centric. Our work contributes a pragmatic account of how design processes for AI systems are enacted.

CCS CONCEPTS

• **Human-centered computing → Human computer interaction (HCI)**.

KEYWORDS

artificial intelligence, design, process, data work

ACM Reference Format:
Maximiliane Windl, Sebastian S. Feger, Lara Zijlstra, Albrecht Schmidt, and Paweł W. Woźniak. 2022. 'It Is Not Always Discovery Time': Four Pragmatic Approaches in Designing AI Systems. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA*. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3491102.3501943

1 INTRODUCTION

Artificial intelligence (AI) is transforming the way we apply and design digital solutions in everyday tasks. In the public discourse there is an increased focus on mitigating the possibly negative aspects of AI. Ensuring control, including human values in the AI system, or designing AI products [17, 18, 24] are highly desired qualities. However, building a user-centred AI future is only possible with the engagement of responsible interaction designers who can effectively work with AI. As a consequence, it is a task for Human-Computer Interaction (HCI) research to study existing and future design processes of AI-based systems to ensure that such solutions are designed with the user in mind.

Recent work in HCI has suggested that AI can be interpreted as a design material in interaction design [5, 8, 27, 30, 31]. This implies that designers need to effectively leverage AI's specific properties and capitalise on the opportunities it provides. Holmquist [8] compared this need to graphic designers who need to be aware of the properties of the paper, printing press, color blends, and much more to design significant products. There is consensus in the community that we need to build a better understanding of design processes which involve AI and the ways in which designers harness the opportunities provided by AI.

Past research investigated the properties of AI which make it particularly complex in a design process. Such challenges include understanding the functionality [1, 2, 5, 30] and capabilities of AI [5, 8, 28, 32], prototyping and testing [5], creating user value [5, 8], and collaborating with developers and data scientists [5, 30]. Given the multitude of factors which designers need to control, there is a need to empirically study design processes with AI in their current form to identify how they can be improved and how we can train future interaction designers. This paper aims to juxtapose the aspects of designing in AI reported in the literature with front-line experiences of design practitioners.

Current research investigating the influence of AI on the design process is limited to a relatively small and specific sample. Yang et al. [30] interviewed 13 designers who all had extensive experience designing AI-enhanced products, mostly worked at very large companies, and designed products with a vast user base. In this work, we study a different group of designers. We portray the diverse landscape of designers by interviewing 20 designers from companies of various sizes. These designers have different levels of AI competence. In this work, we relate existing frameworks for AI-driven innovation [14, 15, 31] to the experiences reported by our participants. Our analysis resulted in four pragmatic approaches

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9157-3/22/04.
https://doi.org/10.1145/3491102.3501943

which design teams adopt when designing for AI: *a priori*, *post-hoc*, *model-centric* and *competence-centric*.

This paper is structured as follows. We first report on related work to motivate our research questions. We then explain our methodology and sample and report on our findings. Based on these findings, we propose four approaches for how designers work with AI. Finally, in our discussion, we relate these processes to the five human interventions in data science by Muller et al. [14].

Our work makes three key contributions: (1) we provide empirical insights into how designers across a diverse sample help to create AI products; (2) we contribute four approaches adopted by interaction designers working with AI: *a priori*, *post-hoc*, *model-centric* and *competence-centric*; and (3) we relate the approaches in the development of AI products to the five human interventions in AI [14]: discovery, capture, curation, design, and creation to show differences in how design for AI systems is enacted in diverse environments.

2 RELATED WORK

In this section, we reflect on related work through three lenses: (1) challenges and opportunities in designing Human-Centered AI (HCAI); (2) the growing body of work around the challenges of designing for AI; and (3) models and procedures for understanding AI and data-centric work practices.

2.1 Human-Centered AI

HCAI is gaining importance to counteract concerns about biased data sets, discrimination of minorities, or privacy threats. Li and Etchemendy [13] suggested that HCAI's goal is to serve "the collective needs of humanity." Therefore, in contrast to technical AI, HCAI does not only focus on algorithmic performance but also evaluates human performance and satisfaction, values user needs, and ensures control [17, 24]. Shneiderman [19] argues that AI processes must be complemented with traditional HCI methods, such as user requirements analysis, design iterations, and usability testing, to create valuable products and services. They propose 15 recommendations to enable this shift towards HCAI and increase the reliability, safety, and trustworthiness of AI systems on three levels: (1) *team*, responsible for reliable software, (2) *organization*, ensuring a safety culture, and (3) *industry*, taking care of trustworthy certifications from outside the organization. In this context, Shneiderman [18] especially highlights the importance of designers who are responsible for explainable user interfaces. Our work provides an answer to Shneidermann's call for focusing on the designers. We study the difficulties and opportunities experienced by interaction designers working with AI to better understand and improve design processes for HCAI.

2.2 Challenges Designing for AI

Recent work has investigated the difficulties that make the design of AI products uniquely different and challenging. Xu [25] states that current HCI methods are inherently insufficient for designing for AI, because they were created for non-intelligent systems. Instead, they suggest enhanced HCI methods, tailored explicitly for AI solutions. Prioritizing the use of machine intelligence functions over the usual focus on visual and interactive design has, for example, already shown to lead to more intuitive UIs and thus better UX [26].

To investigate the practice and challenges of experienced UX designers working on AI products, Yang et al. [30] conducted an interview study with 13 designers. They found that these designers did not think that extensive AI knowledge was necessary to design valuable products; much more critical was the focus on collaboration and a data-centric culture. Moreover, they found that the designers employed a "divide-and-conquer" approach when designing AI products, focusing on the design while barely engaging in technical aspects such as model construction or data collection. As a result, they did not have to adapt their design processes or acquire extensive AI knowledge. However, this strategy also led to protracted and iterative design processes [31]. While these findings accurately describe the experiences of those 13 designers, nine of whom worked at very large companies with more than 10,000 employees, they are only representative of a specific group of designers. In Yang et al. [31]'s study, nine interviewees designed products used by more than one billion users and they were experienced in AI. Our work aims to extend these findings by recruiting a more diverse set of designers from all company sizes and with different AI experience levels to portray the wider landscape of designers.

The challenges surrounding the design of AI products are diverse and can be subdivided into several distinct problem spaces. One such challenge is the designers' knowledge of the capabilities of AI. Designers are often unaware of what AI can and cannot do, which makes it challenging to come up with innovative ideas [5, 9, 29, 33]. Close cooperation or co-designing with developers and data scientists can resolve this issue [7, 22]. However, such collaboration is often difficult, because developers and data scientists can be a limited resource in the company [30], or communication is inherently difficult because both parties lack a common language and a shared workflow [9, 29, 30].

Designers often struggle to create AI products that add value for users while also considering ethical aspects such as who is responsible for potential mistakes by the system [5, 8]. Prototyping AI products is another critical challenge because the prototypes have to resemble the complexity and uncertainty of AI models, while often requiring a large amount of data [23]. This makes the prototyping process slower and more difficult [5]. As a result, designers often resort to Wizard of Oz methods for prototyping and testing [4, 10, 16, 20]. Another possible remedy is model-informed prototyping as proposed by Subramonyam et al. [21], which enables designers to directly work with AI features during prototyping.

Reflecting on these challenges, Yang [27] constructs an AI complexity framework that attributes these hurdles to two root causes: "evolving capabilities whose limits are difficult for designers to grasp, and complex, adaptive interactions that resist simulation." They argue that systems that share these characteristics arrive at the limits of conventional HCI methods and thus call for new methods and tools. Our work studies how these challenges and, possibly, solutions, are used in design practice in a diverse sample of interaction designers working in different domains.

2.3 Models and Perspectives around AI and Data-Centric Work

AI is rapidly changing the development and interaction of systems across all domains. In order to describe these changes, understand the complex activities involved in the design process, and provide frameworks for research and practice, the HCI field has developed models and procedures involved in AI systems design. Olsson and Väänänen [15] proposed to look at AI design practice through *four perspectives*: *product, people, principles, and process*. The *product* perspective concerns designing collaborative and proactive agents. *People* asks to take into account *for whom* systems are designed and *who* is impacted by them, i.e. secondary and tertiary users. *Principles* links to a set of fundamental ethical considerations, including explainability, autonomy, and fairness. Finally, the *process* perspective closely relates to our goal of better understanding design processes for AI systems. While Olsson and Väänänen argued that AI design should follow the basic stages of human-centered design, they stressed that this process was impacted by new challenges, including the need to recognize risks earlier and to develop an evolving understanding around the definition of 'done'. In our work, we investigate how these perspective translate to the lived experience of interaction designers working with AI on an everyday basis.

Yang et al. [31] constructed a framework showing how the two main challenges of designing for AI, namely uncertainty and output complexity, affect the design process. They mapped these two challenges on the translation process between the technological capabilities and the user experience. Muller et al. [14] explored how data science workers work with data through an interview study with 21 professionals from IBM. Based on their findings, they mapped five human interventions that help to understand and formalize complex activities around data practices: Discovery, Capture, Curation, Design, and Creation. *Discovery* relates to *data as given*, for example in the form of client-provided or public datasets. *Capture* is referred to tasks around *data integration*, and *data selection and substitution*. *Curation* relates to activities involving *data-cleaning, converting meta-data, and data-alignment*. *Design* covers a wide range of activities, including *imputing missing data, validating data, engineering features*, and *appropriation of external tools or data*. Finally, the fifth intervention, *Creation*, concerns the crafting of data. Our research revealed that approaches employed by designers involved in the development of AI systems closely align to those five human interventions. Therefore, we adapted this model and demonstrated its applicability in the context of designers working with data and models for AI systems.

2.4 Summary and Research Questions

As AI increasingly raises concerns related to biased data sets, discrimination, or privacy, HCAI becomes more important [11]. Designers assume a unique role of ensuring the embedding of human aspects in AI systems. Designing AI products raises new challenges over those already discussed in the literature. Research has also found that designers barely engage in the technical aspects of AI projects, which leads to lengthy and repetitive processes. It is, however, unknown how these findings expand to the diverse landscape of designers from different organizational sizes and with diverse experience levels. To this end, this paper applies existing models of designing AI systems in a new context. We provide an understanding of how designers adjust their processes to accommodate AI. We also reveal possibilities for the optimization of the AI design process. Our research was guided by the following two research questions:

RQ1: How does AI change interaction design processes? Research suggests that how and if the design process changes depends on the model's complexity level. While for low-level systems with static output the standard HCI methods suffice, completely new tools and techniques are needed for level four systems that constantly change and deliver unpredictable output [31]. We investigated whether this holds for the everyday practice of a diverse sample of designers by examining what tools and methods they leverage and which specific activities they perform in the context of AI projects.

RQ2: How do interaction designers engage in the development of AI systems? Research found that designers often do not engage in the technical aspects of AI projects [30]. This may lead to lengthy and repetitive processes and difficulties in using the capabilities and output complexities of AI systems to their full extent [31]. We investigated whether this assumption holds for a diverse sample of designers and studied how they engaged in the development of AI systems.

3 METHOD

To achieve a comprehensive overview of designers' practices, we conducted semi-structured interviews with 20 designers. We used a retrospective research method to allow the participants to discuss completed projects from beginning to end, covering the entire design process. Previous studies on the practice of designers working with AI have used surveys and interviews [5, 30, 32], which allowed for retrospective data gathering. Since design practice is difficult to capture in numbers or short textual descriptions, the type of data commonly gathered from a survey, we decided to conduct an in-depth interview study. Interviews are often used to characterize practice and acquire deeper insights associated with the matter [12]. In addition to the interviews, we used a pre-interview survey to gather background information on the projects in which the designers were involved and demographic data.

3.1 Interview protocol

Before we started an interview, we provided our interviewees with a short introduction and explanation of the study, followed by practical interview information, such as time and confidentiality. After they had signed the consent form, we began the actual interview. We centered the narrative of the interviews around an example AI project selected by the interviewee before the interview. We first asked them to briefly describe their selected project. Then we asked how the project was initiated including at which point they got involved, the process they followed and specific activities they performed, and any further challenges encountered during the project. We prepared nine additional topics, which we discussed either when participants mentioned them or later as specific questions we initiated. The topics included data collection, AI knowledge, familiarization to AI as design material, design methods, prototyping, mental models, system changes over time, personalization, and the black-box nature of AI. We provide the complete interview guide as supplementary material.

3.2 Participants

We recruited the participants through four channels: the AI x Design Slack community, LinkedIn, our network, and snowball sampling. We identified potential participants on LinkedIn by mentions of AI and design on their profiles. To be included in our sample, participants had to match our inclusion criteria which were: 1) must self-identify as a designer, 2) must have experience working on products that contain or are powered by AI technologies, and 3) must have at least one (nearly) completed project or design case in the field of AI.

Our sample illustrates the diverse landscape of designers involved in designing for AI. We interviewed 20 participants between 23 and 50 years old (M = 31.7, SD = 7.45). The experience levels of our participants in designing AI-enhanced products differed, see Table 1. Most (13) of the participants had a design background, three held a degree in computer science, one in law, two in engineering, and one did not hold a formal degree. The participants worked at companies of various sizes, with six coming from companies with fewer than 100 employees, three with 100-1,000 employees, three with 1,000-10,000 employees, and eight from large corporations with more than 10,000 employees.

3.3 Data Analysis

We recorded and transcribed 22.36 hours of audio during the 20 interviews (M = 70.6min, SD = 10.2min). We used thematic analysis and Atlas.ti to analyze our qualitative data [3]. Thereby, three researchers first independently open-coded two randomly selected interviews. Afterwards, we met to discuss and compare these initial codes, and following this initial discussion the remaining interviews were divided among the three researchers for coding. In the third iteration, we formed categories of related codes and constructed themes. These themes were repeatedly reviewed by comparing and reviewing the coded extracts and comparing them to extracts from other themes. The full Atlas.ti code group report is available as supplementary material. In the end, this process led to the following four themes: Models and Data, Process, Team, and Translation.

4 FINDINGS

We present findings from our study using the four themes Data and Models, Process, Team, and Translation. **Data and Models** are unique to AI-powered products and describe the properties of the models and data used in the design process and designers' responsibility and involvement in the data collection and model construction. The theme **Process** describes how the design process changes when AI products are designed, how these products are tested, and at what stage designers are involved. An interdisciplinary **Team** is of special importance in the context of AI projects. Designers frequently need to consult AI experts for the technical feasibility of their ideas or ask how a specific algorithm or model works to be able to design an interface for it. On the other hand, the AI experts often need the designers' perspective to identify new opportunities for the applications of AI and not lose focus on the real-world value of the product. Designers often take a unique role in the team: they advocate for the users. This might take the form of **Translation**—translating the technical aspects for the users so that they can understand and use the product. Often, translation becomes mediation, where the designers need to negotiate between the user needs and the needs of the engineers.

4.1 Data and Models

Our participants extensively discussed the challenges around data collection and preparation, the different roles they adopted in these processes, and the impact of data quality and fidelity on AI models. We present findings related to four key aspects that informed this major theme: data collection; exploration; fidelity; and model and AI interaction.

Several designers stressed that they were for the first time confronted with tasks and issues around data collection when they got involved in an AI project. Some designers were involved directly with in-situ data collection. P6 was involved in the development of an inventory organization app and took pictures of objects that the application was supposed to classify. Here, the designer needed to have an understanding of the approximate number of photos needed, as well as requirements around sample diversity (e.g. perspective and illumination). In contrast, P1 described a project for which the designers were themselves creating the data as part of the data collection process. They were tasked with creating a service that digitized sketched user interface wireframes and making them interactive. The designers first sketched various prototypes themselves before taking pictures of their designs. Here again, the designers had to develop an understanding of model training requirements, in particular related to sketch and picture diversity. Other participants advised that for some of their projects their "data was given", either as part of public data sets or data provided by the client. In this context, P5 described the challenge of "narrowing the data" rather than creating or capturing it:

> I participated in narrowing the data that is about to add to the model.—P5

Finally, we observed that designers who were not taking an active part in any of those data collection strategies might be involved in the coordination, as another interviewee expressed:

> I do not manage the data myself, but I am in constant alignment with the people who do it.—P19

Designers often help in different ways to explore or interpret the collected data. To ensure good data quality and to report on quality assessments, some designers assume responsibilities associated with roles of data engineers or data scientist:

> One of the roles I see for myself often, when making data visualizations, is I try to sort of expose things like noise in the data, or outlying data points that maybe are indicative of some kind of error in the data sets. [...] Is the data that we are looking at actually correct?—P8

Our participants discussed practices and challenges around data fidelity. Several designers highlighted that they created and used fake data at some stage in the product development cycle. They did this instead of relying on data collection when the cost for

Table 1: Demographic and background data about the participants in our study including educational degree, their role in the project, their experience designing AI-enhanced products, the size of the organization, and the example project they chose to discuss. Our goal was to recruit a diverse sample of designers.

Education Degree	Role (job title)	Exp	Org Size	Example Project
Design, HCI	Product Designer	1-2 yrs	<100	Image recognition for UI tool suite
Design, AI	Human-AI Researcher	2-4 yrs	1,000-10,000	Co-learning autonomous agent
Design	Experience Designer	<1 yr	<100	Fan experience for the visually impaired
Design	Design Manager	1-2 yrs	100-1,000	Recommender for streaming platform
Law	Product Designer	2-4 yrs	>10,000	Prediction for programmatic advertising
Design	Experience Designer	1-2 yrs	100-1,000	Inventory counting application
Design	Experience Designer	2-4 yrs	>10,000	Image making experiences
Design	Data Designer	2-4 yrs	>10,000	Intelligent sleep health monitoring
Computer Science	Data Designer	2-4 yrs	>10,000	Intelligent sleep health monitoring
Design, Business	UX Designer	1-2 yrs	>10,000	Enterprise search
Mechanical Eng.	Product Designer	1-2 yrs	100-1,000	Detection of human animal interaction
Design	UX Designer	8-10 yrs	1,000-10,000	Topic model generator
Engineering	Service Designer	4-6 yrs	>10,000	Assistant for process optimization
None	UX Designer	1-2 yrs	<100	Topic model generator
Computer Science	AI Designer	2-4 yrs	>10,000	Recommender for streaming platform
Computer Science	UX Designer	2-4 yrs	1,000-10,000	Machine learning platform
Design	Service Designer	2-4 yrs	<100	Intelligent text summarization
Design	Product Designer	1-2 yrs	<100	Intelligent coaching chat bot
Design, HCI	Design Manager	1-2 yrs	<100	Diary for patients with respiratory issues
Design	Product Designer	1-2 yrs	>10,000	Prediction for financial markets

collecting real data was too great or when they wanted to quickly test the feasibility of ideas:

> Sometimes you do not need real data to work on. You could totally fake up the data and start working and testing your model.—P11

In this context, several designers emphasized that the experience of building systems using fake data was similar to using real data:

> There was no real data. This was all just design. But I kinda set up the prototype so it felt like you were adding data, you were adding things, and moving through the steps.—P12

The participants further discussed trade-offs and usefulness of fake and real data extensively. In particular, the interview with P11 demonstrates this well. While emphasizing the value in using fake data to test models early on, the designer further demonstrated, related to a concrete project, that using mock data was not always feasible.

> Another project that I am working on right now, for that one you cannot work on fake data at all because the whole concept is from real data. It's impossible.—P11

Our participants described the direct link between data and model fidelity, as well as the resulting implications. Some designers emphasized the importance of creating mock AI without underlying models to build an interaction-based communication channel with other team members. The time at which the model is introduced into the development cycle, as well as its fidelity, further impacts both visual and interaction design. In this context, P10 described an experience of a strongly evolving model, stressing:

> It would be almost irresponsible to lock a certain design in as the one that must be followed.—P10

This notion of *locking in* a design is interesting, as it relates to a design process that is driven by the interface and user experience insights. In contrast, an example from P16 provides a contrasting perspective around data and model driven design processes:

> Then engineering also brings up quite often some edge cases. We are like, well, we are gonna forecast ten metrics, and then how are we going to display our charts? And, you know, so that is where I think we came up with various layouts.—P16

In general, we see that the designers' workflows around core interaction tasks and data-centric responsibilities is becoming more diverse in AI projects. P4 stressed this in the context of *discovery* and *delivery*:

> One quarter we are doing more discovery and in the next part we are more dedicated to deliver something. And this cadence happens all the time. So you have to understand this to work well. It is not always discovery time, it is not always delivery time. You have to balance.—P4

4.2 Process

This theme describes how the designers experienced the changes in their roles, responsibilities and work routines when using AI.

Most of the participants expressed clear opinions about how AI changed design processes. Few designers like P7 expressed that just because they were using machine learning, principal challenges and responsibilities around interaction design remained similar. Another interviewee argued that designing AI products was just like designing every other product:

> Now, why can't we look at that and go, there is nothing different here. It's still software, it's just a lot of advanced mathematics, or whatever you want from that.—P13

Yet most interviewees described substantial changes in the design process of AI applications. Some participants reported on preconceptualized or existing systems where AI was later integrated in order to transform the user experience through new advanced features or to solve very specific problems. P12 stressed that such an integration of limited AI features might have been the result of attempts to automate actions that users previously had to complete manually.

Our participants discussed when they were involved in the projects. The timing of designer involvement often determined the format of the design process. If an AI model serves as the starting point, designers tend to get involved too late, after a product or feature has already been developed to a certain degree. That means they can no longer bring in their specific expertise to guarantee that users will accept the feature. P12 echoed this:

> (I) was introduced too late to the process. That is a huge thing I know UX designers face. And there was one project where I was introduced to fix things after the engineering team had already implemented a big chunk of the product—P12

Most participants, instead, emphasized the need to involve designers early on, even in projects with a pre-existing AI system. In this context, P10 highlighted that designers should be consulted about the design of application programming interfaces (APIs):

> In a perfect world, design comes before engineering so we can guide the APIs and the other builds, what to build right.—P10

Further, we found that the time at which designers get involved impacts the range of activities that they can guide or impact during the design and development process. In projects built on existing AI, for example, their work focuses on the interaction design. In processes where AI is to be added to existing applications, designers participate in various activities around data collection and preparation, as outlined in the theme DATA AND MODELS. Here, they also adapt and apply common design processes to support the development process at various stages. For example, P4 reported adopting common card sorting techniques to metadata. P9 further described applying wizard of oz methods to compensate for models that were not yet ready:

> We want to quickly get some signal to the field, [...] even if it's these days not yet possible by a machine. We try to still make our interventions in a way that we expect that later on the algorithms will be available or the development team will just be able to develop such a system.—P9

4.3 Team

Our participants stressed that team diversity and interdisciplinary efforts represent key pillars for many successful AI projects. For example, when developers create new models and features, they ask designers for input and visualizations. P16 talked about concrete experiences:

> They want to bounce their ideas around with me and say, oh, you know, we have this technical piece that we are working on in an algorithm. And we think that it can be used here. So can you maybe spend half a day and put together some quick wireframes to capture the flow?—P16

In this context, P18 went even further and addressed the intricate interdependence:

> I always say that it's teamwork. I cannot get my work without my other two colleagues. [...] My work is normally always dependent on theirs and theirs dependent on mine.—P18

The study participants extensively discussed the challenges of communication between developers and designers. Most participants stressed that designers needed a basic understanding of AI and its abilities and limitations in order for this relationship to work. In this context, P12 emphasized that AI knowledge had helped them in communicating with team members and taking part in the decision making process. P15 noted that the designer at the very least wants to be able to have an idea of the kind of technology that is suitable to address a real-world problem best. At the same time, most participants indicated that they had little to no knowledge about AI when starting on their current projects. Most participants reported learning the basics either through online courses or by learning on the fly, i.e. by continuously consulting the AI experts on the team for clarifications or feasibility checks, as P11 stated:

> I was having a lot of meetings internally with our product team, with our dev people to learn about limitations that we might have. Things we could do, things we could not do.—P11

Yet, we also note that several designers stressed that they could not profit from in-depth knowledge around AI development. Few designers reported challenges in demonstrating their value and establishing their expertise within more technical project teams:

> I am thinking about, more about how people accept the designer working on machine learning projects. It's not always easy. Some people still see the designers as UI designers. The designers have nothing to do with machine learning.—P15

4.4 Translation

This theme relates to the designers' tasks and responsibilities around communication and translation of different perspectives, backgrounds, and expectations within the project teams. While designers traditionally play an important role at the interface between users and developers, their role is changed when designing with AI.

The interviewees described how it was often the job of the designers to mediate between the AI experts and non-AI experts in the team. Even though this translation within the team was essential, it was perceived as challenging and required additional work:

> I stayed over at work, and I have asked them to explain to me. Sometimes I have to admit I did not understand. They were also sometimes using neural networks and stuff that they sometimes explained to me like a black hole even. It was complex, even for them.—P1

The need to explain constraints involved in the AI and its models was discussed as a novel and challenging task, possibly unrelated to traditional tasks of interaction designers. P14 provided an example that depicts the conflict between user expectations and model limitations:

> We had to find a way to try to simplify what (the product) does in a way that was easily understandable. And they don't have a great mental model of all the machine learning going on when it takes fifteen minutes to get an output. But they're used to doing a Google search and getting immediate results. That was something that people didn't like. They didn't understand why it would take so long.—P14

The *buzz* around AI that has been pushed by media reports and fictional material has resulted in many people overestimating its capabilities, necessitating additional translation efforts between the developers and both customer and end users. P11 talked about their frustration dealing with unrealistic expectations, stressing that users had "absurd and unrealistic wishes". Finally, turning those unfeasible expectations into features that add value, represents a major effort in translation

> It's very hard to find this match between users and AI. AI is powerful, but in some cases, it's over-marketed [...] The biggest factor here is to find the right use case for making AI work in a meaningful way and a valuable way.—P5

4.5 Design Approaches in Interaction Design for AI systems

Based on findings from the interviews, we classified the design decisions and processes in which the design teams engaged into four approaches. We used the term *approaches* to encapsulate both intentional actions, similar to strategies, as well as actions that are at least partly driven by the socio-technical frameworks of the designers. These approaches represent archetypes of how design teams deal with the challenge of including a data model as part of a user-centered design process of interactive technologies. Figure 1 shows the four approaches: *a priori*, *post-hoc*, *model-centric* and *competence-centric*. While the first two approaches focus more on temporal aspects, the second two are more process-orientated. We describe the approaches in terms of how the system's interface interrelates with the underlying AI model. This enables us to discuss the influence of including AI in the design process and showcases the differences between how system design is understood within teams. In the following, we describe the approaches in detail.

4.5.1 A Priori. As we saw in Data and Models, many designers were confronted with data and models for the first time when AI was introduced to the project. This led to entirely new challenges, uncertainties and processes. Further, as seen in team, some designers struggled with understanding the capabilities of AI or even felt that deeper AI knowledge would distract them from their core tasks. As a consequence, some design teams opted for not engaging with AI whatsoever and treated the AI functionality as *a priori* correct and functioning according to plan. The model could have been considered ready both when it was implemented and tested or when no further engineering work was possible due to company policy. An alternative scenario was assuming that the specification of an unimplemented model was fully correct and perfect functionality was assumed through an intermediate artifact—a model which was considered to be ready. An advantage of this approach was that the design team felt independent of the AI specialists and could focus on enacting their usual, unchanged design process. However, adopting this approach also introduced additional vulnerability as most models are imperfect and change over time. A design process where the model is assumed to be static may introduce problems where the model does not produce ideal results. Another risk is that the finalised product lacks real user value and thus real-world applicability because the designers were involved too late (cf. process).

Example Project. Designers working with this approach repeatedly reported being involved too late to sufficiently influence the project's success or even to fix an unsuccessful product already on the market. P14 was employed to evaluate and redesign a topic model generator that had already been released but was not being very well received by the users. The problematic aspects were mainly the result of the developers neglecting the non-technical perspective, as we observed in teams. One issue was, for example, its long execution time. The designer reported being overwhelmed by the complexities of the project and spending weeks coming to understand the project and constraints of the model. Model-wise, the designer was mainly focused on the in- and outputs because the model was already implemented, making any further impact impossible.

4.5.2 Post-hoc. The *post-hoc* approach is, figuratively, a mirror image of the *a priori* approach. Like the former, this approach is primarily adopted by teams who desire a clear separation between model and interface development. The teams adopting this approach usually limited their translation activities to the minimum. Instead, a clear separation between the design team and the data team was in place, as we observed in teams. Providing the data team with a detailed specification of what the AI in the system should do was the ultimate goal of the design team. Design iterations were conducted until an accurate enough specification of the required model was developed. Once the interface was finalised, the design team defined the inputs of the model and specified what kind of inference was required for the system to offer a user-friendly experience. This approach has the advantage of giving the design team a high degree of independence. On the other hand, it ignores possible design constraints posed by the technical implementation of a model.

Example Project. P3 reports about a project where they were tasked with creating a fan experience for visually impaired people.

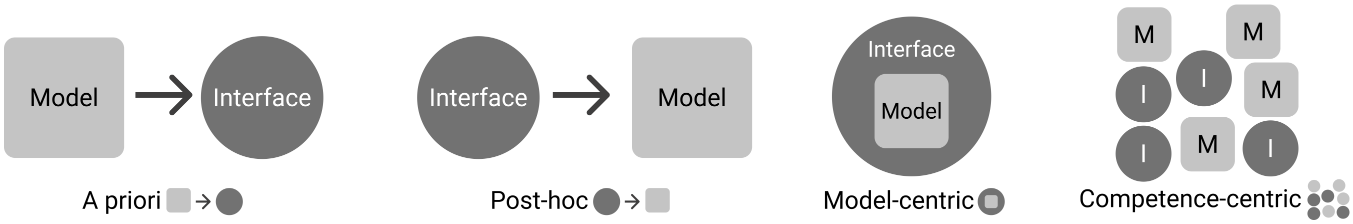

Figure 1: Four approaches for designing for and with AI. The approaches describe how the design teams in our study included models in their design process and how they adapted their methods to include (or avoid) designing a data model.

The designers first developed the design using traditional techniques. They hardly deviated from their usual process by employing regular and known strategies for requirements analysis and ideation such as brainstorming sessions or participatory design. After the design was finished, designers tasked the AI experts with the realization of the feature. The designers were not involved in any of the technical aspects of the project, such as data collection or model construction. In fact, the participant was unsure whether the final product included the designed functionality.

4.5.3 Model-Centric. Our results show that teams that adopted this approach used the data model not only as a key element of the systems but also as a vehicle for driving the design process forward. In contrast to the previous two approaches where the data and model were treated as separate from the design, the designers were now directly involved in technical aspects. Thus, this approach required TEAMS where the AI experts were eager to share their expertise with the interface designers. In PROCESS, we observed how the integrated team developed the model together, considering its limitations and possible training sets. To allow for iterative interface development, the teams built intermediate models with limited functionality. This enabled early prototyping and synchronisation (MODELS AND DATA). Further, as the design team developed an understanding of what was possible given the AI solutions available in the project, they could develop *fake AI* tools for prototyping in anticipation of trained models later in the process. This approach allows for coherent coordination between the design and data teams and facilitates developing the model and the interface in parallel. It does, however, come at a cost. In the TRANSLATION theme, we observed that teams adopting this approach spent extra effort and resources on coordination and translation. Further, this approach is highly dependent on the team members' skills in sharing knowledge.

Example Project. P13 reported on a project where the *model-centric* approach was applied. They developed an intelligent assistant for optimizing chemical processes and operations. The design and model were constructed together with a steady and clear focus on teamwork. There was consensus that the combination of diverse expertise would lead to a successful outcome. Thus, the designers were involved in almost every step of the data collection and model construction. In the data collection, they could bring in their extensive knowledge about human decision-making, as described in MODELS AND DATA. The designers were even responsible for the decision of what model to use. They decided to use an optimizer instead of a concrete prediction to keep the decision-making power with the users. The designers were so involved in the technical aspects of the project that they even participated in testing and executing the model. Artifacts used in the design process, such as paper prototypes, were developed and evaluated together to guide the model development.

4.5.4 Competence-Centric. The final approach was used by teams where the design and data teams featured members with diverse competence. In MODELS AND DATA, we observed how understanding the different roles of the team members led to helping each other define tasks. Thus, users would collectively decide when to train and build models or when to use simulated input. This resulted in a divergent design PROCESS where multiple intermediate artifacts were developed at the same time. The multitude of ideas and prototypes contributed to an exploratory approach to developing the model and interface. Further, working on more atomic artifacts contributed to more TRANSLATION and team members understanding the priorities of their colleagues. Adopting this approach emphasises the individual skills and creativity of the team members. However, the divergent nature of the process may make it difficult to control the temporal aspect of the process, especially as divergent design iterations are often not synchronised. This approach appears to be most suitable for exploratory projects where the exact nature of the AI system is to be defined.

Example Project. P10 employed a *competence-centric* approach to develop an enterprise search for a large corporation. Even though the design and model were developed in parallel, the process was not as interwoven as in the *model-centric* approach. Instead, the designers and AI specialists focused on their specific expertise and repeatedly synced to ensure the individual developments matched, which could be time-consuming. The designer was once again responsible for bringing in human expertise. This was, however, not achieved by developing the model together with the AI specialist but rather by testing the finished model (or selected parts of it) and reporting problems. Since the teams were disconnected, the artifacts used in the process were also divergent and required TRANSLATION. The design team, for example, constructed their own test data set to be able to evaluate whether the model matched their user interface.

4.5.5 Summary. Depending on the capabilities of the team and constraints of the project, designers employ different approaches when designing AI products. The *a-priori* approach assumes the model to be finished before the start of the UI development, so the challenge for designers is fulfilling the requirements of the model. In contrast, in the *post-hoc* approach, the UI design is completed before the model, implying that the model is built to fit the requirements of the design. In both approaches, the design team is not required to deeply engage in the data aspects of the project. In contrast, in the *model-centric* approach, the model is at the core of the project and the designers are strongly involved in nearly all technical aspects. The *competence-centric* approach is similar in the way that it also requires designers to engage with the data, albeit never in a leading. The competence-centric approach leverages a clear focus on the heterogeneous expertise within the team.

Finally, we note that there is an apparent conceptual divide between the four approaches. The first two approaches, *a-priori* and *post-hoc* are focused on temporal aspects, that are heavily impacted by techno-economical constraints. This implies that these approaches are primarily motivated by the infrastructure already in place within an organization and, possibly, how technical processes are organized. In contrast, *model-centric* and *competence-centric* approaches are process-oriented. Thus, they are primarily determined by the social dynamics within the development team and the particular combination of skills present when developing a given system.

5 DISCUSSION

The findings of our interviews show that the designers adopted diverse pragmatic solutions to effectively include AI development in their design processes. Contrary to past work, our results paint a picture of designing with AI which is more diverse and requires more makeshift solutions. Our approaches provide, on the one hand, useful lenses through which the design processes can be analyzed in retrospect. On the other, they can serve as a guideline to decide for a certain approach or an adaptation of organizational frameworks tailored to desired outcomes or constraints of a project.

In order to contextualise our findings within past work, we relate the four approaches which we observed to the five human interventions in data science practice as postulated by Mueller et al. [14], as shown in Figure 2. As predicted by Mueller et al., we identified elements of all five interventions in our findings. In the *a priori* and *post-hoc* approaches, some of the human interventions are decoupled from the interaction design process and designers do not take part in them. Thus, discovery, capture and curation are delegated to a different team. In our results, teams using those approaches did not report participating in those interventions. In the *model-centric* stategy, we observed that there was intensive translation work at the boundary between curation and design. Here, design teams used artifacts which helped them communicate what the model could (potentially) deliver and how it could be improved. This took the form of developing fake AI or partially functional models. These intermediate artifacts allowed the entire team to be part of all five interventions. Finally, in the *competence-centric* approach, the teams developed numerous divergent artifacts. This implied that multiple interventions happened simultaneously. This way, different members of the design teams could participate in different interventions at different times.

5.1 From Perfect Processes to Pragmatic Approaches

In contrast to previous work that found that AI design processes are characterized by a clear separation of technology and design [31], our findings suggest diverging levels of designer involvement in the data collection and model construction. Our four approaches show these different levels. When either the *a priori* or *post-hoc* approach is applied, the model and design are independent, which is in line with insights from previous studies [30]. Our results show that a low involvement in the model can be a result of specific company culture, a strategic decision or a result of the designers being uncomfortable in dealing with the technicalities of AI. However, in contrast to previous research, our interviewees also reported occasions of great involvement in the technical aspects, where the designers even took part in the data collection and model testing. As seen in the *competence-centric* or *model-centric* approaches, the designers are consciously involved in the technical aspects. In contrast to previous work that stated that designers did not feel a need for extensive AI knowledge to be able to effectively collaborate with the technical team [30], our interviewees reported that greater knowledge was an important building block for effective collaboration. In the *model-centric* approach, designers and AI specialists develop the model and product together. This approach results in the biggest involvement of designers in the technical process as they are asked to bring in their expertise to every stage of the model development. When the *competence-centric* approach is employed, designers and AI specialists develop the product in parallel while each team is focused on their specific expertise. Thus, the involvement of the designers takes a different form.

Looking at the approaches from a designer involvement point of view, it can be observed that adapting one of the approaches can result in a certain level of involvement. Alternatively, the approach could be a pragmatic product of the level of involvement which is comfortable for the interaction designers. In the design processes studied in this paper, a common requirement of all approaches was that designers needed to develop an understanding of the data quality and diversity needed to train models. This was achieved with differing levels of success. Our work suggests that there is a need to further study the organisational processes which lead to determining the approaches in order to understand why particular design scenarios are enacted. Further, our results show that while past models of AI-based design processes are useful in understanding how the intellectual value is built through AI, further work is needed to see how design teams can reach a stage where these processes are in place in an optimal form.

5.2 Towards Improved Processes for Designing AI systems

The approaches presented in this work can serve as archetypes for understanding the pros and cons of different process choices when designing for AI. For example, the *competence-centric* approach is especially useful for exploratory projects where time is not an important factor since its divergent nature fosters creativity, which may come at a time cost. The *a-priori* and *post-hoc* approaches might be a direct response to the project's constraints when the design team is included in the project to develop a product after a viable model is built. Some teams may not have the capacity to foster the close cooperation needed in the *model-centric* or *competence-centric* approaches. While future design processes will most likely use a blend of the approaches identified in this paper, our insights can

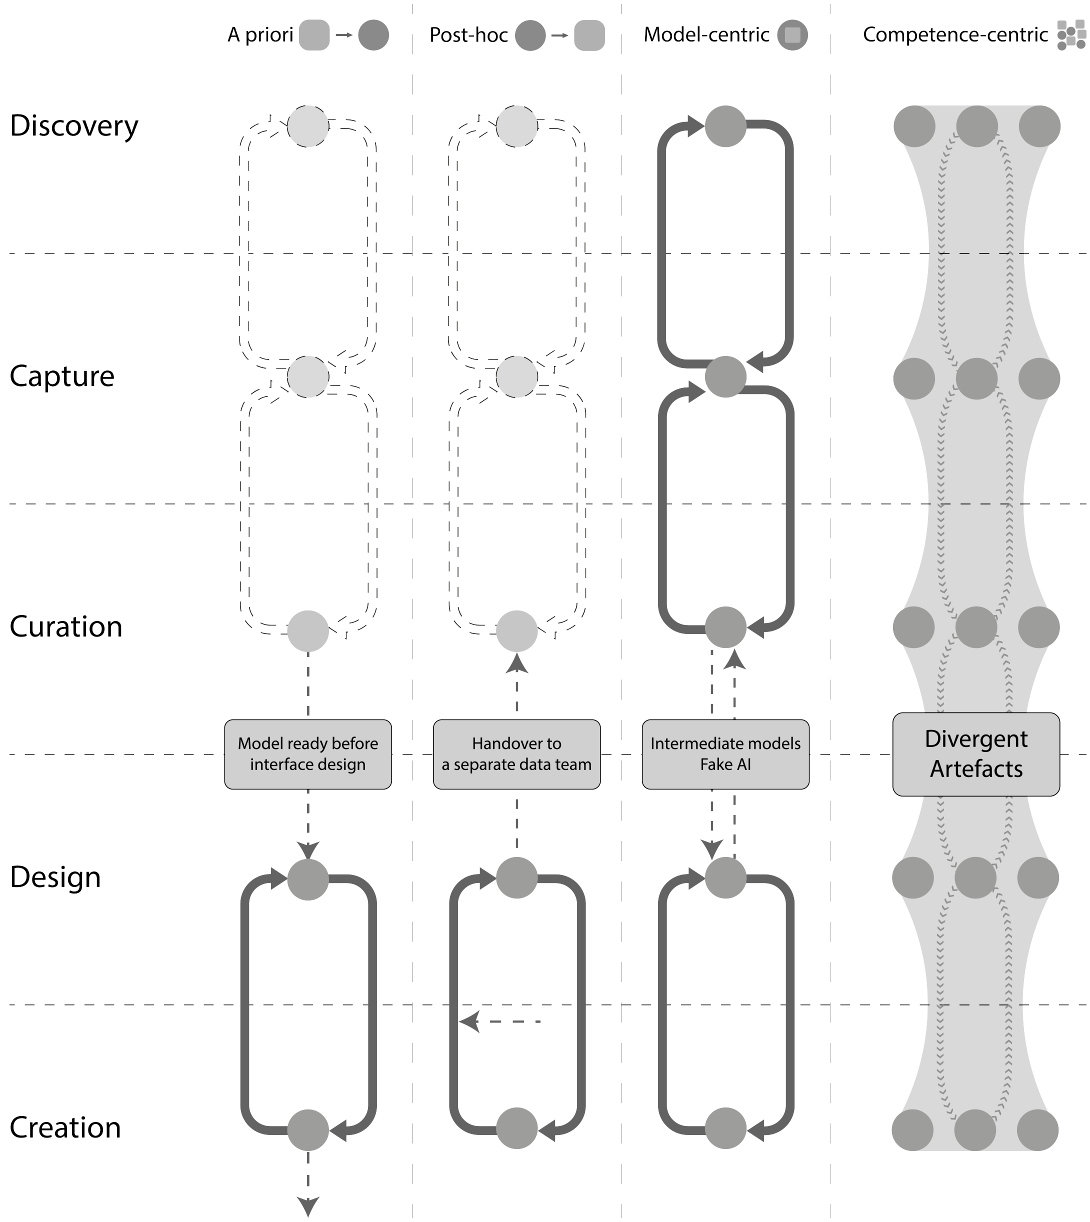

Figure 2: The four approaches for designing with AI identified in our study related to the five human interventions in data science by Mueller et al. [14]. As our approaches identify continual processes, they need to be visually presented in a certain order, yet the intervention can happen in any order and their timing is primarily determined by the needs of the design process.

help understand the pros and cons of the approaches and tailor particular processes to a particular task and design team.

We note that the approaches identified in our work are purely empirical—many of the design cases described in this paper describe often emergent practices rather than elaborate processes with extensive planning. Thus, our account can be used by future researchers to chart if and how the nature of designing with AI will change in the future. Further, while future design activities can choose to explicitly adopt one of the approaches, we anticipate that it might be effective to use different approaches at different stages of the design. While we did not observe such shifts, we anticipate that this should be possible with the development of designing with AI. For instance, design processes could initially adopt a *competence-centric* approach to foster divergence and then shift to a *model-centric* approach as the technical product matures. These approach choices could mimic how using different tools in the design process allows for different levels of convergent thinking [6] and thus allows for better control of the design process.

Finally, a key finding of our work is the fact that most of the processes we studied were grassroots efforts determined by particular team compositions and organizational cultures. This indicates that there is future research needed in how we can explicitly organize such processes. More importantly, we need to understand how we can meaningfully include AI in interaction design education so that processes such as the ones described in this paper require less effort. The approaches presented in this paper contribute to the discussion on how we can create effective design teams for AI that support interdisciplinary collaboration.

5.3 Limitations

We recognise that our inquiry is prone to certain limitations. While we did aim to recruit a diverse sample of designers, there certainly is more diversity in AI design teams. Future studies should investigate approaches in AI design teams across different organisational cultures and geographical locations. Moreover, our analytical lens treated design teams as single entities which traversed design processes. Further insights can be achieved by studying different roles of designers between teams and investigating individual accounts of multiple designers from a single team. Finally, we note that we classified the design processes reported in our data into a finite set of four approaches. While these approaches present analytical archetypes of the stories contained in the interviews, they also present an inherently holistic view of the approaches. Thus, the approaches do not fully represent the temporally unstable nature of the design processes.

6 CONCLUSION

In this paper, we studied accounts of designing with AI by 20 interaction designers. In contrast to past work, we studied a diverse group with different levels of AI competence, working in organisations of different sizes. Through thematic analysis, we identified four themes which described the designers' experience of designing AI systems: Data and Models, Process, Team, and Translation. Based on the themes, we identified four design approaches (*a priori*, *post-hoc*, *model-centric* and *competence-centric*), which teams adopted to adjust their design process to include AI. We then related the approaches to Mueller et al.'s five human interventions in data science to show how the design processes can vary in levels of designer involvement. We hope that our work inspires further research on how AI design is enacted by practitioners.

ACKNOWLEDGMENTS

This work has been funded by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A. The authors of this work take full responsibilities for its content.

REFERENCES

[1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3173574.3174156

[2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300233

[3] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI Research: Going Behind the Scenes*. Morgan & Claypool Publishers, 51–60. https://doi.org/10.2200/S00706ED1V01Y201602HCI034

[4] Justin Cranshaw, Emad Elwany, Todd Newman, Rafal Kocielnik, Bowen Yu, Sandeep Soni, Jaime Teevan, and Andrés Monroy-Hernández. 2017. Calendar.Help: Designing a Workflow-Based Scheduling Agent with Humans in the Loop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2382–2393. https://doi.org/10.1145/3025453.3025780

[5] 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*. Association for Computing Machinery, New York, NY, USA, 278–288. https://doi.org/10.1145/3025453.3025739

[6] Jonas Frich, Midas Nouwens, Kim Halskov, and Peter Dalsgaard. 2021. How Digital Tools Impact Convergent and Divergent Thinking in Design Ideation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 431, 11 pages. https://doi.org/10.1145/3411764.3445062

[7] Fabien Girardin and Neal Lathia. 2017. When User Experience Designers Partner with Data Scientists. https://www.aaai.org/ocs/index.php/SSS/SSS17/paper/view/15364

[8] Lars Erik Holmquist. 2017. Intelligence on Tap: Artificial Intelligence as a New Design Material. *Interactions* 24, 4 (June 2017), 28–33. https://doi.org/10.1145/3085571

[9] 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* (Glasgow, Scotland Uk) *(CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3290607.3299059

[10] 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* (San Diego, California, USA) *(UIST '00)*. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/354401.354406

[11] Jeff Larson, Julia Angwin, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

[12] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research Methods in Human-Computer Interaction* (2 ed.). Morgan Kaufmann, Cambridge, MA. https://www.safaribooksonline.com/library/view/research-methods-in/9780128093436/

[13] Fei-Fei Li and John Etchemendy. 2021. Welcome to the Stanford Institute for Human-Centered Artificial Intelligence. https://hai.stanford.edu/about/welcome

[14] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q. Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3290605.3300356

[15] Thomas Olsson and Kaisa Väänänen. 2021. How Does AI Challenge Design Practice? *Interactions* 28, 4 (June 2021), 62–64. https://doi.org/10.1145/3467479

[16] Laurel D. Riek. 2012. Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *J. Hum.-Robot Interact.* 1, 1 (July 2012), 119–136. https://doi.org/10.5898/JHRI.1.1.Riek

[17] Filippo Santoni de Sio and Jeroen van den Hoven. 2018. Meaningful Human Control over Autonomous Systems: A Philosophical Account. *Frontiers in Robotics and AI* 5 (2018), 15. https://doi.org/10.3389/frobt.2018.00015

[18] Ben Shneiderman. 2020. Bridging the Gap Between Ethics and Practice: Guidelines for Reliable, Safe, and Trustworthy Human-Centered AI Systems. *ACM Trans. Interact. Intell. Syst.* 10, 4, Article 26 (Oct. 2020), 31 pages. https://doi.org/10.1145/3419764

[19] B. Shneiderman. 2020. Design Lessons From AI's Two Grand Goals: Human Emulation and Useful Applications. *IEEE Transactions on Technology and Society* 1 (2020), 73–82.

[20] Lisa Stifelman, Adam Elman, and Anne Sullivan. 2013. Designing Natural Speech Interactions for the Living Room. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (Paris, France) *(CHI EA '13)*. Association for Computing Machinery, New York, NY, USA, 1215–1220. https://doi.org/10.1145/2468356.2468574

[21] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. ProtoAI: Model-Informed Prototyping for AI-Powered Interfaces. In *26th International Conference*

on Intelligent User Interfaces (College Station, TX, USA) *(IUI '21)*. Association for Computing Machinery, New York, NY, USA, 48–58. https://doi.org/10.1145/3397481.3450640

[22] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. Towards A Process Model for Co-Creating AI Experiences. In *Designing Interactive Systems Conference 2021*. Association for Computing Machinery, New York, NY, USA, 1529–1543. https://doi.org/10.1145/3461778.3462012

[23] Philip van Allen. 2018. Prototyping Ways of Prototyping AI. *Interactions* 25, 6 (Oct. 2018), 46–51. https://doi.org/10.1145/3274566

[24] Dirk Wrede, Tino Stegen, and Johann-Matthias Graf von der Schulenburg. 2020. Affirmative and silent cyber coverage in traditional insurance policies: Qualitative content analysis of selected insurance products from the German insurance market. *The Geneva Papers on Risk and Insurance - Issues and Practice* 45, 4 (01 Oct 2020), 657–689. https://doi.org/10.1057/s41288-020-00183-6

[25] Wei Xu. 2019. Toward Human-Centered AI: A Perspective from Human-Computer Interaction. *Interactions* 26, 4 (June 2019), 42–46. https://doi.org/10.1145/3328485

[26] Wei Xu, Dov Furie, Manjunath Mahabhaleshwar, Bala Suresh, and Hardev Chouhan. 2019. Applications of an interaction, process, integration and intelligence (IPII) design approach for ergonomics solutions. *Ergonomics* 62, 7 (2019), 954–980. https://doi.org/10.1080/00140139.2019.1588996 PMID: 30836051.

[27] Qian Yang. 2020. *Profiling Artificial Intelligence as a Material for User Experience Design.* dissertation. Carnegie Mellon University.

[28] Qian Yang, Nikola Banovic, and John Zimmerman. 2018. Mapping Machine Learning Advances from HCI Research to Reveal Starting Places for Design Innovation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3173574.3173704

[29] 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*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300415

[30] 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* (Hong Kong, China) *(DIS '18)*. Association for Computing Machinery, New York, NY, USA, 585–596. https://doi.org/10.1145/3196709.3196730

[31] 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*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376301

[32] Qian Yang, Jina Suh, Nan-Chen Chen, and Gonzalo Ramos. 2018. Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) *(DIS '18)*. Association for Computing Machinery, New York, NY, USA, 573–584. https://doi.org/10.1145/3196709.3196729

[33] Qian Yang, John Zimmerman, Aaron Steinfeld, and Anthony Tomasic. 2016. Planning Adaptive Mobile Experiences When Wireframing. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) *(DIS '16)*. Association for Computing Machinery, New York, NY, USA, 565–576. https://doi.org/10.1145/2901790.2901858