

# AINEDSPLANNER: A Workbook to Support Effective Collaboration Between AI Experts and Clients

Dae Hyun Kim  
dhkim16@cs.stanford.edu  
KAIST  
Daejeon, Korea

Hyungyu Shin  
hyungyu.sh@kaist.ac.kr  
KAIST  
Daejeon, Korea

Shakhnozakhon Yadgarova  
yadgarova@kaist.ac.kr  
KAIST  
Daejeon, Korea

Jinho Son  
sjhfam@algorithm labs.co.kr  
Algorithm Labs  
Seoul, Korea

Hariharan Subramonyam  
harihars@stanford.edu  
Stanford University  
Stanford, CA, USA

Juho Kim  
juhokim@kaist.ac.kr  
KAIST  
Daejeon, Korea

## ABSTRACT

Clients often partner with AI experts to develop AI applications tailored to their needs. In these partnerships, careful planning and clear communication are critical, as inaccurate or incomplete specifications can result in misaligned model characteristics, expensive reworks, and potential friction between collaborators. Unfortunately, given the complexity of requirements ranging from functionality, data, and governance, effective guidelines for collaborative specification of requirements in client-AI expert collaborations are missing. In this work, we introduce AINEDSPLANNER, a workbook that AI experts and clients can use to facilitate effective interchange of clear specifications. The workbook is based on (1) an interview of 10 completed AI application project teams, which identifies and characterizes steps in AI application planning and (2) a study with 12 AI experts, which defines a taxonomy of AI experts' information needs and dimensions that affect the information needs. Finally, we demonstrate the workbook's utility with two case studies in real-world settings.

## CCS CONCEPTS

• **Software and its engineering** → **Collaboration in software development**; • **Computing methodologies** → *Artificial intelligence*.

## KEYWORDS

AI application planning, client-AI expert collaboration, information flow, workbook

### ACM Reference Format:

Dae Hyun Kim, Hyungyu Shin, Shakhnozakhon Yadgarova, Jinho Son, Hariharan Subramonyam, and Juho Kim. 2024. AINEDSPLANNER: A Workbook to Support Effective Collaboration Between AI Experts and Clients. In *Designing Interactive Systems Conference (DIS '24), July 1–5, 2024, IT University of Copenhagen, Denmark*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3643834.3661577>

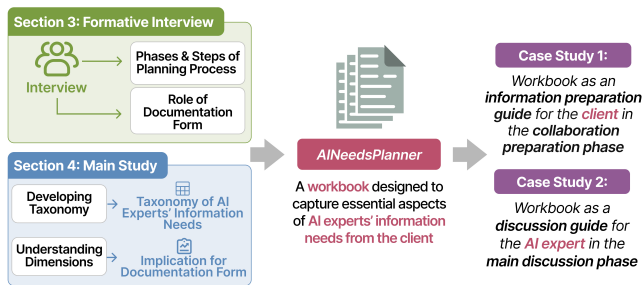
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
*DIS '24, July 1–5, 2024, IT University of Copenhagen, Denmark*  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0583-0/24/07.  
<https://doi.org/10.1145/3643834.3661577>

## 1 INTRODUCTION

Creating AI applications that are human-centered, ethical, and meet regulatory requirements requires effective communication and collaboration between different stakeholders with diverse backgrounds, incentives, and values [11, 64]. Because of the complexity of the AI technology [10, 20, 45, 72], we have seen the separation of those who desire to introduce AI into their workflow and AI experts who have the ability to build AI models. Hence, *client-AI expert collaborations* has become a widespread form of collaboration between stakeholders [56].

Client-AI expert collaborations, which can either take *inter-organization* forms or *inter-team* forms, are characterized by the large amounts of discussion during the planning stage prior to execution by the AI experts [48, 60]. During this discussion, the clients should not only determine and share their goals, expectations, and available resources, but also share additional information, including domain knowledge around the data [66, 67] and domain-specific performance metrics [29, 54], for a successful AI application development. Through a formative interview of 10 successfully completed AI application development projects (Section 3), we further find that the planning stage of AI application development is innately an information flow from the client to the AI expert through three phases; (1) the client initially define their goals and assess their needs and resources before entering collaboration, (2) the client passes on the information with the AI expert while undergoing discussions to iterate on the information, and (3) the AI expert makes execution plans based on the iterated information.

Despite the need to closely involve the clients, many client-AI expert collaborations encounter significant challenges. To begin with, traditional engineering methods have not been designed for sufficiently involving the clients, despite the importance of information flow from the client [55, 61]. To make things worse, the knowledge barriers between the client and the AI expert only add to the challenges [6, 16, 57, 69]; clients not only have difficulty forming a mental model of the information needs for development to start, but also expressing information in forms that can be readily incorporated into the AI application design by the AI experts. The need to closely involve clients and the knowledge barriers between the clients and AI experts exacerbate the challenges in determining the uniquely important aspects of AI application development, such as datasets and success measures [44, 50, 61]. For instance, consider the planning process for an AI-powered cancer treatment



**Figure 1: An overview of the structure of the paper. We use the formative interview and the main study results to inform the AINEEDSPANNER workbook, on which we performed two case studies.**

recommendation tool for doctors. Despite the dataset’s importance, doctors may struggle to figure out the expected data format and quantity for training the AI model. Similarly, AI experts may face challenges in determining available data and relevant fields. Moreover, doctors may have vague goals, making it difficult to express them in quantifiable metrics for AI development. Discussions on these topics will require overcoming the knowledge barrier that entails substantial effort from both stakeholders due to the high degrees of specialization of the two fields.

Based on these existing challenges and the widespread adoption of documentation during the planning stage [48], we introduce the AINEEDSPANNER workbook as a tool for facilitating effective collaboration between AI experts and their clients with deeper client involvement (Figure 1). Based on the needs we uncovered during the formative interview, we design the workbook as a versatile form that can not only be used as a guide for the client as they define their goals and assess their needs and resources prior to discussions but also as a discussion roadmap during the client-AI expert discussion. To define the contents of this workbook, we conducted a study with 12 AI experts to taxonomize and characterize the information needs of the AI experts from their clients. To understand the value of AINEEDSPANNER as an information preparation guide for the client while preparing for collaboration and a discussion guide and checklist, we put the workbook to use in two different case studies.

The contributions of this paper are as follows:

- Identification and characterization of constituent steps of the AI application planning in client-AI expert collaboration;
- A taxonomy of AI experts’ information needs from clients during AI application planning;
- AINEEDSPANNER, a workbook for supporting client-AI collaboration in the planning stage; and
- Two case studies providing insights into the value of AINEEDSPANNER as both an information preparation guide for the client and a discussion guide and checklist.

## 2 RELATED WORK

Our work is grounded in three main areas of prior work: (1) Supporting AI specification, (2) Workflows for building AI applications, and (3) Collaboration for building AI applications.

### 2.1 Supporting AI Specification

Building AI applications is a challenging task that involves a large number of decisions with multiple expertise needed. Since it requires a significant amount of time and effort to build, the process needs careful planning that clearly defines the goal and identifies necessary resources to make sure that the resulting AI application creates value. However, little research has investigated designing specifications for AI application planning. Existing specification templates such as Datasheets for Dataset [22] and Model cards [44] can be used in the planning stage for clearly setting the design goals of AI applications, but the nature of “reporting” as the goal misses important dimensions that need to be considered in the planning stage (e.g., available resources for training models and expected impact of the application to the organization). More importantly, the templates assume some level of AI expertise by using AI-related terminology, which makes it challenging for clients (i.e., AI novices) to fill in the information in the planning stage.

To support lay users in designing AI applications, research introduced guidelines, more hands-on workbooks, and interactive systems. The guidelines [1, 2, 27, 33, 39, 42, 50, 57, 68] inform a set of important considerations with specific examples. When it comes to clients planning AI applications, however, it is challenging to contextualize the guidelines for a specific context because of a lack of AI experiences. People+AI Guidebook [50] provides fill-in-the-blank exercises that scaffold the process across six chapters, using simple languages. However, it is unclear whether they are complete as the planning material for supporting the early stage [73]. In this work, we aim to (1) identify a set of information needed in the planning stage and (2) create a workbook that elicits the information from the clients so that the clients can be well-prepared for building AI applications.

### 2.2 Workflows for Building AI Applications

The process of building AI applications involves multiple stages where some of the important tasks are (1) setting clear goals about what AI should produce and defining success metrics, (2) designing and collecting datasets, and (3) building and evaluating models. Due to the uncertain and complex nature of AI, research showed that the AI application-building process entails unique challenges across the design process [72]. To address such challenges, research introduced interactive systems supporting various stages (e.g., data analysis [4], prototyping [58], and UI building for multiple stakeholders [19]). However, it is not yet investigated how to systematically support planning AI applications, especially for the clients.

Planning is a crucial step for determining the goals and means of achieving the goal [25]. Prior work in various domains shows that planning can significantly increase the performance of organizations and quality of work [12, 52, 53, 62]. For the AI application building process, research also highlighted the planning phase as a separate step [59] and noted the stage as a long preparatory stage [71]. However, it remains unclear what the workflow of the planning is as well as the underlying challenges. As clients are knowledgeable in their domain and contexts, they have the potential to make unique contributions in the planning stage such as clearly defining what AI should do, designing domain-specific target metrics, estimating data availability, and depicting interactions

between the user and AI to make it usable in real-world contexts. But to make such a contribution, clients need to have concrete guidance that elicits such information from the clients. We aim to design a workbook for eliciting such information from the clients so that the workflow of the planning gets clearer and more systematic.

### 2.3 Collaboration for Building AI Applications

As the process of building AI applications is complicated, it involves a collaborative approach between multiple stakeholders (e.g., domain experts, AI experts, and UI/UX designers). However, it is challenging to support such collaboration largely due to the complexity of AI [72]. For example, designers find it challenging to understand the capabilities of AI models and their behaviors [14, 57]. In making decisions, domain experts require information more than prediction results of AI models [7, 38]. To address such challenges, research suggested improved methods for collaboration [57, 59]. For instance, research proposed a process model for collaboration that utilizes user data as probes [59].

However, little research has investigated the collaboration between clients and AI experts for planning AI applications. Research showed that communications between different stakeholders entail challenges in establishing common grounds [40]. Boundary objects can resolve the challenge [8], but such boundary objects for the planning AI application have not been designed yet. Similar to prior research introducing various artifacts as boundary objects (e.g., checklist [5], action plan [49], data documentation [28], AI explanations [3], and UI prototype [58]), we design a workbook as a boundary object for the planning stage by understanding information that needs to be shared between clients and AI experts as well as underlying challenges during the planning stage.

## 3 FORMATIVE INTERVIEW: UNDERSTANDING THE PLANNING PROCESS BETWEEN AI EXPERTS & CLIENTS

To understand the process of AI application planning between an AI expert and their clients as well as how documentation forms could fit in to further facilitate this process, we performed a semi-structured interview of clients and AI experts who have participated in successfully completed AI application development projects. In the formative interview, we specifically sought to find answers to the following research questions:

**[RQ1]** What steps constitute the AI application planning process between an AI expert and their client and what are the key characteristics of these steps?

**[RQ2]** How can a form of documentation fit into the AI application planning process in client-AI expert collaboration to further facilitate the process?

### 3.1 Participants

To obtain detailed perspectives from the two stakeholders of interest, we recruited (1) *AI experts* and (2) *clients* of recent AI application projects that have concluded successfully (self-reported), a restriction put to capture the good planning practices. We required that the participant took part in the planning phase of the AI application project.

**Table 1: An overview of AI application projects we surveyed.**

Project	Project Structure	Project Domain	Participant Role
IO1	Inter-Organization	Education	Client, AI Expert
IO2	Inter-Organization	Food, Manufacturing	Client, AI Expert
IO3	Inter-Organization	Pharmaceutical	Client, AI Expert
IO4	Inter-Organization	Robotics, Cooking	Client
IT1	Inter-Team	Education	Client
IT2	Inter-Team	Commerce, Video	Client
IT3	Inter-Team	Medical	AI Expert
IT4	Inter-Team	Language	AI Expert
IT5	Inter-Team	Finance, Banking	AI Expert
IT6	Inter-Team	Medical, Pharmaceutical	AI Expert

We focused our interviews on *inter-organizational projects* (IO1-4) as well as *inter-team projects* within the same organization (IT1-IT6). We excluded inner-team projects or projects motivated by the AI experts themselves as they display no AI expert-client relationship. The AI experts were typically AI engineers, researchers, or technical leads of the AI companies (e.g., CTO, team leader), while the clients ranged from employees in client teams or companies to management roles at client teams or companies (e.g., project manager/owner, CEO).


We recruited the participants through the authors' personal and academic connections and then employed snowball sampling [23] via referrals. While recruiting, we focused on capturing diverse project domains by limiting ourselves from including more than two projects from the same domain. In addition, to respect any secrecy regulations of the participants' organizations, we informed the participants about the contents of the interview beforehand. While we recorded the interviews by default, we also offered the option of not recording if it goes against any of the policies of the participant's organization.

### 3.2 Procedure

We performed the interviews on Zoom [74] to abide by various COVID-related regulations and to allow ourselves to reach out to participants outside of our immediate geographical vicinity.

At the beginning of the interview, we introduced ourselves and provided an overview of the contents of the interview. We then asked for general information about the participant's organization and the domain of the AI application development project we would discuss during the interview. Next, we conducted the main interview, where we asked about the procedures and the contents of planning. While asking them to describe the process of planning, we asked about the stakeholders involved in the process and their roles, the information flow between the stakeholders, and various characteristics of the process including the challenges and the remedies. Finally, we allowed the participants to freely share their thoughts and ask us any questions they had throughout the interview.

For analysis, we first extracted each step in the process mentioned for each of the projects in Miro [43]. We then used an affinity diagram [35, 36] on these steps to come up with steps that commonly appear throughout multiple projects. After identifying the common steps, we referred back to the original interviews to further group these steps into phases and to understand the relative ordering of these steps.



**Figure 2: The three phases of the AI application planning process in client-AI expert collaboration and the comprising steps of each of the phases.**

### 3.3 Assessing RQ1: The Process of Planning

Through our analysis, we identified a total of 11 steps within the planning and grouped them into 3 phases: (1) the collaboration preparation phase, (2) the main discussion phase, and (3) the execution preparation phase (Figure 2).

**3.3.1 Phase 1. Collaboration Preparation Phase.** Before the collaboration with an AI expert begins, planning begins on the client-side. A client typically (1) defines the goal they wish to achieve with the AI application, (2) confirms whether the AI application would make 'financial sense', (3) defines any specific needs for the model, and (4) performs initial preparations before entering a collaboration with AI expert.

**Defining Goal of the AI Application.** One of the most important tasks that the client must perform during this phase is to define the goal of the AI application. In most if not all cases, the goal includes the problem that the client wishes to solve with the application, the target users of the application, and how the application is to be used, although the level of specificity may differ. Although not common, the client may occasionally possess concrete target measures for success (IT1).

**Assessment of Financial Feasibility.** Before seeking AI teams and kick-starting the AI application development, the client needs to confirm what budget they have available for data collection, development, and deployment and whether introducing the AI application would pay off, especially if the client is part of a for-profit organization (IO2-4, IT1-6). While the budget for development and deployment is usually the person in managerial roles, the considerations that go into determining profitability vary. For smaller organizations (e.g., startups), the profitability assessment is often based on the likelihood of securing more investment by distinguishing themselves from their competitors (IT2, 3), government funding for startups applying AI technology in niche domains (IO3, IT2, 6), or a more direct suggestion by the investors to introduce AI into their workflow (IO4).

**Defining Specific Needs.** Some of the clients have very specific needs about various components of AI application development. The most common needs that clients have are a timeline or a deadline by which they need the application developed by (IO1, IT1, 6) or specific features to be included in the resulting AI application (e.g., dashboard visualization for IO2, personalization for IO4, explainability for IT5). There are occasionally more specific needs about the core AI model that would power the AI application, especially

an update to a novel AI model (e.g., a specific deep learning-based natural language processing framework for IT3, large language models for IT4, deep learning for IT6).

**Preparation for Collaboration.** Before starting collaboration with AI experts, some clients perform additional preparation beforehand to streamline the communication and to be able to have more concrete discussions. Clients occasionally precollect data and preprocess the data into a form that they can share with AI experts in accordance with corporate and regional restrictions (IO1, 3). Compliance with corporate restrictions in terms of data sharing is more of an issue of inter-organization collaborations than of inter-team collaborations. Other than data-related preparations, one of the clients went as far as building an evaluation framework before starting collaborations with AI experts (IT1).

**Notes on Phase 1.** While it would be ideal for the client to have been able to clearly define their goals and needs prior to starting collaboration with AI experts, it is usually unrealistic. In particular, clients' lack of AI knowledge and the uncertain nature of AI uncertainties at the time of planning usually limit forming concrete and correct expectations about the AI model that would power the application (IO1), the characteristics of the errors the underlying AI model may give (IT1), and the data quantity and format required to train the underlying AI model (IO4, IT2). Hence, while the first phase serves an important purpose in specifying the client's initial goal and expectations, these initial goals and expectations are concretized and calibrated through collaboration with AI experts in later phases.

**3.3.2 Phase 2. Main Discussion Phase.** The second phase is where the client's goals, situational information, and expectations formulated in Phase 1 are passed to AI experts so that the AI experts can have sufficient information based on which they can come up with initial approaches for execution. A notable characteristic of this phase is that while the general direction of information flow is from the clients to AI experts, significant amounts of discussions occur between clients and AI experts, allowing further concretization of the various components determined by the client in Phase 1.

**AI Expert Matching.** The first step of the main communication phase involves seeking AI experts who can best fulfill their AI application development needs. The main criteria for determining the fit of the AI expert in the AI application development is their understanding of and experience with the domain of the desired AI application.

For inter-organization projects, the client selects the AI expert best fit for developing their AI application by going through a bidding process (IO1, 2) or by looking at the AI expert's portfolio to gauge prior experience in the domain (IO3, 4). For inter-team projects within the same organization, the client team simply asks the team with AI expertise, which naturally possesses both an understanding and experience in the domain, to develop an AI application for them.

**Client Understanding.** While the core information passed from the clients to the AI experts is focused on the goal, available resources, and client expectations, AI experts often need to request additional information from the client about the domain and their workflow to make it across the knowledge barrier. Doing so allows

AI experts to carry on discussions with the client about the data and technical feasibility and sketch out potential technical approaches. Hence, the clients provide descriptions of their domain and their workflow to the AI experts (IO1, 2, 4, IT6).

However, despite the efforts to reduce the knowledge barriers and provide context, many collaboration efforts still end up experiencing some challenges in terms of communication and need to invest additional time and effort to work around the barrier (IO1-4, IT5, 6).

**Data Discussion.** Data is among the most important pieces when it comes to AI application development not only because it further informs the input and the output, but also because data quality is crucial in determining the performance of the AI model trained on the data.

If the client already has a collected dataset, they would pass the dataset over to the AI expert and discuss whether the data can be worked with or additional data or preprocessing is required (IO3). In the cases in which the client has not collected data, AI experts and the client arrive at a plan for dataset collection. They do so by matching what the AI experts believe they need based on the given goal and what the client can collect and provide to the AI expert (IO1, IT2). Furthermore, if the label is missing from the data, the client and the AI experts also discuss data labeling methods (IO3).

The client also shares with the AI expert their domain insights on the features in the data that are likely more relevant than others (IO1, 2, IT5).

**Technical Feasibility Discussion.** Based on the information provided by the client, the AI expert must gauge the feasibility of the goals set by the client given the available resources and expectations. It may occasionally be the case that the client's goals and expectations exceed what is possible with the available resources. In these cases, the AI experts need to calibrate their client's goals and expectations to something feasible with the given resources.

Because the clients usually lack AI expertise to perform technical feasibility assessments, they rely on AI experts for the decision (IO2, IT3, 5, 6)

**3.3.3 Phase 3. Execution Preparations.** The execution preparation phase is characterized as a stage in which the AI experts play the central role, although the client may be involved to some degree.

**Technical Approach Planning.** Based on the information obtained and discussed with the client in Phase 2, the AI experts sketch out ways to meet the provided goals and the client's expectations, while being subject to the various restrictions set by the available resources.

The technical approach planning often involves sandbox testing along the way, especially if prior attempts to apply the planned AI model on the specific problem domain are sparse (IT1-5). After defining a potential technical approach, AI experts build small prototypes of the AI model often with parts of the data (e.g., minimum viable product) and verify the prototype to check whether they can achieve the proposed goal of the client and whether it meets the client's expectations. When the success targets are relatively clear, the AI expert can perform this on their own (IT1, 2, 4); when they are unclear, verification of the sandbox testing results occasionally involves the client (IT3, 5).

**Dataset Collection & Labeling.** Clients or AI experts collect data and label to be used during the execution stage if they have not been collected by this phase (IO1, 3, IT1, 4).

### 3.4 Assessing RQ2: Documentation Form in Facilitating the Planning Process

As previously noticed by Piorski et al. [48], many of the projects we surveyed already employed some documentation during the planning stage (IO1-4, IT1, 2). Yet, the projects either utilized documentation in the absence of a predefined format (IT1) or documentation with only the high-level sectioning (IO1-4, IT2). The task of determining the exact contents with which to fill in the documentation follows the discretion of two stakeholders based on their tacit knowledge of the required discussions acquired through years of experience. A standardized documentation form would help centralize and surface the tacit knowledge that is required through the process. The centralized documentation form would not only help lower the entry barrier to client-AI expert collaboration, but also help the two stakeholders form expectations when initiating client-AI expert with a stakeholder they have not worked with, which is especially common for inter-organization collaborations (IO2-4).

We observed two major ways that the clients and AI experts were employing documentation in planning. The first use case of documentation was centered around the client during the collaboration preparation phase as they determine what information to consider (IO1-2), and during the main discussion phase, they initially pass on that information to the AI expert (IO1-2). The second use case of documentation was centered around the AI expert during the main discussion phase as they take notes of the meetings for future references (IT1, 2) and guide discussions with their clients (IO3), and during the execution preparation phase as they formulate a technical approach and communicate the back to the client asynchronously (IO1-2, 4).

Based on these use cases, we deduce that a documentation form that clearly outlines the information needs of the AI experts as they draft a technical approach can be useful for both use cases leading up to the main discussion phase. In the first use case centered around the client, a detailed documentation form would help clients understand what expectations they should set and assessments they should make prior and guide them through the process, while also serving as a form that they can send to the AI expert without much additional effort. In the second use case centered around the AI expert, a detailed form can not only help the AI expert organize the discussion, but can also help them ensure that all key information has been discussed with the client.

## 4 MAIN STUDY: UNDERSTANDING & TAXONOMIZING THE INFORMATION NEEDS OF AI EXPERTS FROM CLIENTS

While we have established a need for a documentation form outlining the detailed information needs of AI experts from their clients through the formative interview in Section 3, the steps identified from the interview only provide some higher-level insights about the contents of the documentation form. Hence, to arrive at a more detailed list of information needs to inform the contents of the

documentation form, we conducted a study on 12 AI experts asking for their information needs on AI application development problems collected from the real world. We form a taxonomy of the information needs of AI experts from their clients through affinity diagramming [35, 36], which would form the basis of the contents of our documentation form in Section 5.

We performed the study with two research questions:

**[RQ1]** What is the taxonomy of the information needs from the AI expert to the client during the planning stage?

**[RQ2]** What affects these information needs?

The study protocol we used is included in the supplemental material.

## 4.1 Participants

For the study, we recruited AI experts who have prior experience participating in AI application development with clients so that we can realistically simulate the information needs of the AI experts. Because of our focus on the planning stage, we required that the participants have participated in the planning stage, directly interacting with their clients.

To obtain diverse perspectives, we recruited the participants through two channels: the authors' industry connections and referrals, and Upwork [63]. For participants we recruited through our industry connections and referrals, we targeted AI engineers in AI solutions companies or non-AI companies since the AI engineers in these companies tend to have experience directly interacting with their clients during planning. For participants we recruited through Upwork, we carefully reviewed their portfolios and prior experience published on their information pages to ensure that they have prior experience in building AI applications. Through the recruitment process, we recruited a total of 12 participants (Table 2), from various countries, including the US, Egypt, India, South Korea, and Spain.

The study lasted 60-75 minutes and we compensated the participants with an equivalent of 50 USD or the rate at the hourly rate posted on Upwork, up to 50 USD per hour.

## 4.2 Procedure

To separate out and examine the AI expert's information needs from the client, we recruited AI experts and had them go through simulated asynchronous discussions based on a compiled list of AI application ideas from the real world. We had them organize their AI application development approaches for the execution stage by filling in an *execution preparation document* that we compiled based on a model card format [44], while noting any questions they need to ask the hypothetical client. Based on the questions, we utilized affinity diagramming [35, 36] to arrive at a taxonomy of AI experts' information needs from the client.

**4.2.1 AI Application Ideas List Compilation.** For a realistic setup for the participants of the study, we generated a list of AI application ideas based on the AI applications mentioned in the interviews in Section 3, exploratory discussions with various clients and AI experts prior to the interview, and pilot studies. Because the purpose of the study is to elicit information needs from the AI experts, we carefully limited the information provided was in the format of: “[AI

application description] ([purpose of the AI application])” (Table 2 right column).

**4.2.2 Execution Preparation Document Generation.** As a way to clearly identify the reasons why the AI expert would have specific information needs and to avoid the AI expert accidentally leaving out certain execution-related information during the limited duration of the study, we built an *execution preparation document*. The execution preparation document is designed to capture the list of information that is directly involved in the AI application development. It is designed so that filling in the document would require roughly the same amount of information as sketching out a technical approach in the execution preparation phase. Specifically, the execution document includes not only information that is derived directly from the client's needs and expectations (e.g., target users, use cases) but also decisions and deductions made by the AI team based on the provided information (e.g., AI model to use).

We generated this document based on *model cards*, document forms used to communicate key information about an AI application to broad audiences, because model cards are already widely adopted across various AI expert communities [17, 24, 34]. Of the model card formats, we started with the model card format presented by Mitchell et al. [44], as it has been reviewed through a peer-review process and not tuned to specific communities. However, a model card is aimed at disseminating an AI application after its development is complete, while the execution preparation document is aimed at organizing information required for execution preparation. To amend these differences, we added missing fields related to the execution preparation phase to the model card and removed those that were irrelevant. To identify the missing fields, we reviewed the contents mentioned in the interview and marked any contents that the interviewee mentioned as relevant to the execution preparation phase. Next, we asked three AI experts within our organizations to fill out the model card based on the AI application development project they are participating in, while noting down any missing fields or irrelevant fields. Some notable fields marked as missing and added in the final execution preparation document were about the timeline and the regulations and policies around data and AI use. On the other hand, some fields marked as extraneous and removed from the final execution preparation document were model version and software license.

We ported the execution preparation document to Google Docs [13] for the study so that we can observe the participant's edits as they perform the task and so that we can use the comment feature for collecting their information needs.

**4.2.3 Study Execution.** Since the participants of the study were from various geographical locations, we performed the study on Zoom [74] and Google Meet [41], depending on participant preference.

After introducing the research team and the purpose of the study, we began by asking about the participant's background: amount of experience with AI and the project domains of some of their noteworthy AI application development projects.

Based on our findings from the interview that a client seeks an AI expert who is likely to have prior experience and expertise in their domain, we matched each participant with projects from the

**Table 2: An overview of the AI experts who participated in our study and the AI application ideas they selected for the study with the respective application domains in brackets.**

Part.	AI Exp.	Selected AI Application Idea
P1	3 yrs	[Manufacturing / Logistics] Predicting demand for manufactured food products at various store locations (for planning manufacturing and shipping of goods)
P2	7 yrs	[Finance / Banking] Classify loan applications into those that should be approved/declined, or need further review (to assist and speed up the loan application review process)
P3	6 yrs	[Education / HR] Predict student dropout at universities (to proactively provide support to those most vulnerable)
P4	4 yrs	[Manufacturing / Logistics] Predicting demand for manufactured food products at various store locations (for planning manufacturing and shipping of goods)
P5	4 yrs	[Art / Cartoon] Generate detailed sketches of cartoon cuts in the artist's style based on the artist's rough sketch (to help artists with the exploration of the design space)
P6	2 yrs	[Education / Recommendation] Providing personalized recommendations of courses (to help learners continue their path of learning on an online educational platform)
P7	4 yrs	[Manufacturing / Quality Control] Detecting defective pixels in newly manufactured televisions (to automate the process of items with defects)
P8	1 yr	[Productivity / Customer Support] Building an automated phone answering system that can respond to various consumer questions (to help reduce the burden of smaller companies that cannot keep its own customer care team)
P9	5 yrs	[Productivity / Customer Support] Building an automated phone answering system that can respond to various consumer questions (to help reduce the burden of smaller companies that cannot keep its own customer care team)
P10	6 yrs	[Rendering / Images] Manipulate the light source (direction, color, brightness, etc.) in an image (so that artists can remaster their images without needing to recapture images)
P11	6 yrs	[Military / Surveillance] Identifying key military activities and targets from surveillance images (to help generals quickly and accurately make decisions)
P12	10 yrs	[Art / Cartoon] Generate detailed sketches of cartoon cuts in the artist's style based on the artist's rough sketch (to help artists with the exploration of the design space)

AI application ideas list that they are the most familiar with. Specifically, we asked the participants to choose up to three project ideas from the AI application ideas list based on (1) their familiarity with the project domain and (2) the machine learning tasks potentially involved in approaching the project (e.g., object detection, question answering). If the participant stated that their familiarity with a certain project idea greatly exceeds that with others, we assigned them to the project given that the same project was selected by at most one other participant. Otherwise, we selected the project within their top 3 selection that no other participant had selected.

After deciding on the AI application project, we asked the participant to imagine a scenario in which a client has approached them to develop an AI application for them, and pause the scenario right after the client has only had a chance to state what the AI application would be about and the purpose of the AI application as given in the description. As a method of collecting information needs of the AI experts, we asked the participant to fill in the execution preparation document while noting anything they either expect to hear from the client upon unpause the scene or questions they would ask the clients as Google Doc [13] comments on the fields that require that information. When filling in the study preparation document, we first asked the participants to skim through the fields to familiarize themselves with the contents of the document and then fill it in in the order they wish to. We also instructed our participants to assume that their client has AI knowledge so that the participant would not limit themselves from listing out information needs due to their assumptions on what their client can answer.

Once the participant finished filling in the execution preparation document, we first checked whether they could imagine an initial execution plan if they were given answers to all the information needs they listed. Then, we asked them how the information needs would have differed if the client did not have sufficient knowledge about AI or if the AI application project was instead in a domain that they were not familiar with. Finally, we asked them to freely express any other thoughts they had about the study.

**4.2.4 Result Analysis & Taxonomization.** For analysis of the results, we first extracted a total of 227 instances of information needs from the Google Doc comments left on the execution preparation document. Two of the authors put the instances of information needs into Miro [43] and used an affinity diagram [35, 36] to come up with the lowest-level classes of information needs. From these

lowest-level classes of information needs, we repeatedly used affinity diagramming to arrive at higher levels of classes until no longer being able to further group similar classes of information needs.

### 4.3 RQ1: Information Flow Taxonomy

As a result of our analysis, we obtained a taxonomy of the information needs of AI experts from their clients (Figure 3). The taxonomy includes a total of 71 nodes and has a maximum depth of 4. We provide a brief overview of the top-level nodes of the taxonomy in the order of the number of instances of information needs corresponding to each top-level node. Please refer to Figure 3 for further details.

- **Dataset (88 instances; 39%)** includes information about the data format (input & output), availability of the data and the collection method, domain insights into the data, and regulations and restrictions around data use; this information influences various parts of the AI expert's decisions, including the AI model to use, need for preprocessing or data augmentation, metrics, and feasibility assessment.
- **Model Needs (39 instances; 17%)** includes information about the client's needs in selecting the AI model to use (e.g., explainability, generalizability, runtime) as well as regulations on AI use; this information directly influences the AI expert's decision around the choice of the AI model to utilize.
- **Defining Success (30 instances; 13%)** includes information about the metrics for success and the threshold to meet on the metric for the application development to be satisfactory to the client; this information directly affects the AI metrics and target thresholds that the AI expert would use during development.
- **Budgets, Resources & Framework (22 instances; 10%)** includes information on the available resources (e.g., GPU, API) and budgets available for development and deployment; this information serves as the limiting factor for the AI expert as they determine the AI model to use.
- **Understanding the Current Situation (17 instances; 7%)** includes information on the client's workflow and difficulties as well as related work and the client's prior attempts; this information helps the AI expert gain the required background knowledge and the information on related work and client's prior attempts in particular inform potential target thresholds.
- **Client-AI Agreement (14 instances; 6%)** includes information about the timeline, deliverables, and functionality requirements

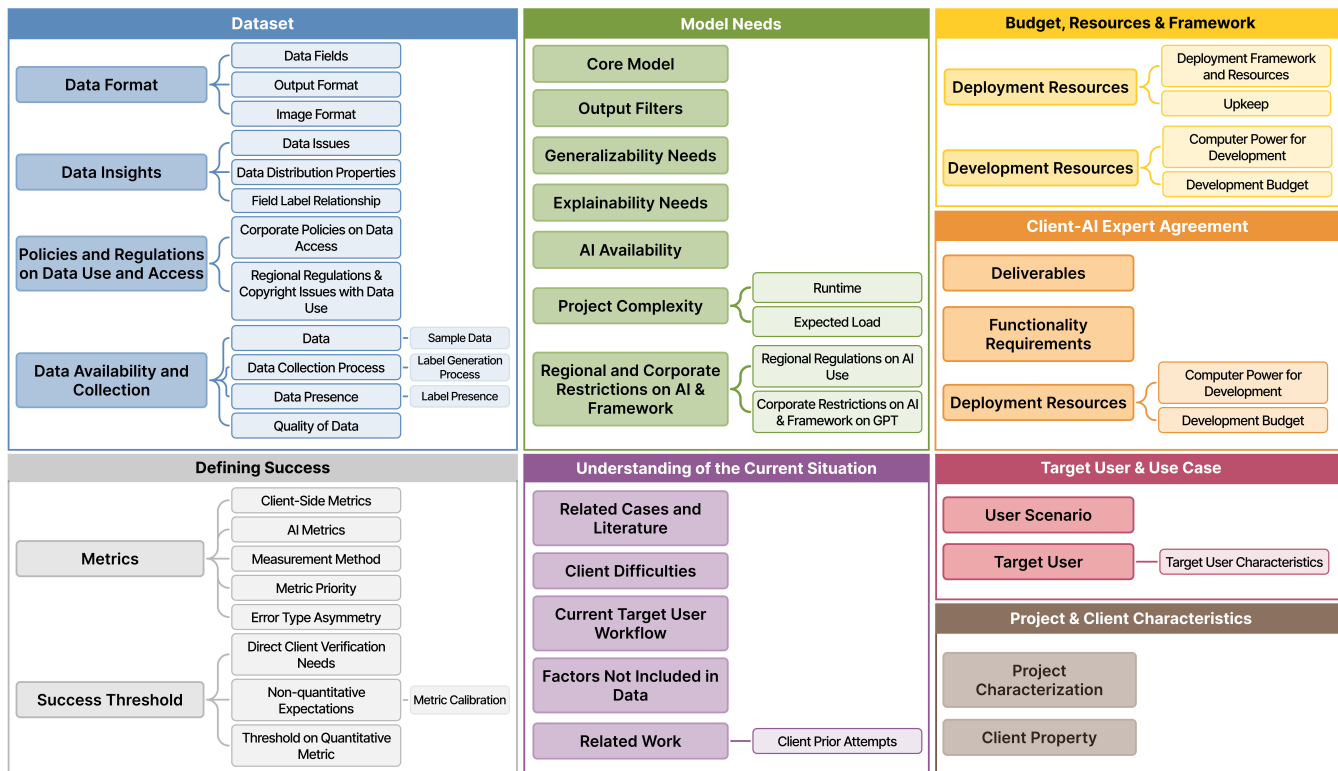


Figure 3: A taxonomy of the information needs of an AI expert from their client as they sketch out their technical approach.

that the resulting AI application must meet; this information informs the AI expert about what terms they would expect to be bound to as a part of the client-AI expert collaboration.

- **Target User & Use Case (10 instances; 4%)** includes information about who would be using the resulting AI application and the usage scenarios; this information serves as background knowledge for the AI expert as they reason about various decisions, including those about the input and the output.
- **Project & Client Characteristics (7 instances; 3%)** includes information about the nature of the project (e.g., research prototype, global deployment, etc.) and the philosophy of the client company; this information mostly serves as a guide for the AI expert in forming expectations about the collaboration style.

#### 4.4 RQ2: Dimensions Affecting Information Needs

We found four dimensions affecting the information needs of AI experts from their clients: (1) AI knowledge level of clients, (2) domain knowledge level of AI experts, (3) input data type, and (4) sensitivity of data handled in the domain. We discuss how the information needs depend on each of these dimensions and the implications each of the dependencies has on the documentation form design.

**4.4.1 AI Knowledge Level of Clients.** In the study, we asked the participants to assume that they possess AI expertise in order to elicit information needs from the participants without holding back

questions based on the assumption that the client will not be able to answer them. While it is true that many of the clients do possess some AI knowledge or learn about AI along the way, clients who do not possess AI knowledge are common (noted separately by P2, 7, 10-12). In general, the participants stated that as AI experts, they would be much more likely to make the more technical decisions in defining success (AI metrics for development progression, threshold on the AI metric) on their own instead of asking the clients about the technical details (P5, 7, 9, 12). In addition, they also stated that they would give *directives* about the data requirements (e.g., format, quantity, collection methods) instead of expecting that the clients would be able to make the right decisions on their own (P2, 5, 7, 12). On the contrary, participants were more willing to invite the clients to determine the tasks that are typically considered the tasks of the AI expert, such as modeling methods, fine-tuning methods, or even AI-related coding (P3, 5), although such client intervention in the AI expert's traditional role is not always welcome (P11).

**Implications for Documentation Form Design.** A generalizable documentation form should be usable by clients with varying degrees of AI knowledge. The documentation form should allow more control of the technical decisions for clients who possess AI knowledge, while not requiring the same information from clients who do not. One way to achieve this goal would be to include an optional input field about the technical decisions so that clients who have AI knowledge can express their voice, while clients without sufficient AI knowledge can simply skip the input field.



**4.4.2 Domain Knowledge Level of AI Experts.** While clients tend to choose AI experts who are more likely knowledgeable about the problem domain, some AI experts admitted that they had instances in which they had to work with clients in a problem domain they were unfamiliar with (P7, 10). We believe that this is because of the existence of a long tail when it comes to the space of all possible AI application domains (e.g., P7’s prior experience with tattoos). To understand how information needs would change in these cases, we asked the hypothetical question about what would change about the information needs if the participant were instead given a problem in an unfamiliar domain. Participants stated that they would inquire much more heavily about understanding the target user workflow (P1, 4, 6, 7, 9, 10) and insights into the data and the data fields (P1, 5, 8, 11). Yet, many participants said that they would also seek information sources other than the clients for additional information (P2, 7, 9-12).

**Implications for Documentation Form Design.** A generalizable documentation form, when filled out, should contain sufficient domain information for the AI experts to comprehend. Because the amount of information needs increases with a decrease in an AI expert’s domain expertise, the documentation form should be made for the AI experts with the least domain expertise in order to be usable by any AI expert. Hence, the documentation form should be inclusive of the information needs of AI experts without domain expertise.

**4.4.3 Input Data Type.** The data the AI application takes as input can take many different forms (e.g., text, images, tabular data). We observed that there are information needs specific to certain data formats. For instance, if the input is of an image format, some of the commonly mentioned information needs were around the dimensions and resolution of the image (P7, 11, 12) and the number of color channels in the dataset (P11, 12). In comparison, if the input is of a tabular format, information needs focus more on the fields of the data (P1, 3, 4).

**Implications for Documentation Form Design.** Unlike the knowledge levels discussed in the previous sections, the input data type is a categorical variable. Thus, it is necessary to identify possible data types the input data can take and understand what can be asked in common across multiple input types and what needs to be asked differently for each input type. Input data can be multimodal and include multiple data types (e.g., image question-answering, data tables containing images), but attempting to deal with each data type in separation could lead to exponential growth in the cases that the documentation form needs to cover. Therefore, it would be more tractable to break down each multimodal input type into each modality and deal with them in isolation. Furthermore, there exists a long tail of input data types (e.g., 3D mesh data) and may be a place for adaptations for specific domains instead of trying to deal with every possible input data type; domains dealing with less common input data types can extend the documentation form based on their needs.

**4.4.4 Sensitivity of Data Handled in the Domain.** Another dimension affecting information needs is the sensitivity of the data being handled in the domain. While information needs on client policies on data sharing hold for all domains, information needs on

<b>I. Overview</b>	<b>a.</b> AI Application Summary <b>b.</b> Target User	<b>V. Budgets, Resources &amp; Framework</b>	<b>a.</b> Development Resources <b>b.</b> Deployment Resources & Framework
<b>II. Dataset</b>	<b>a.</b> Data Availability & Collection <b>b.</b> Dataset Format <b>c.</b> Policies and Regulations on Data Use and Access	<b>VI. Understanding the Client Situation</b>	<b>a.</b> Target User Workflow <b>b.</b> Related Work <b>c.</b> Non-data Field Factors
<b>III. Defining Success</b>	<b>a.</b> Metrics <b>b.</b> Success Threshold	<b>VII. Client-AI Expert Agreement Terms</b>	<b>a.</b> Timeline <b>b.</b> Deliverables <b>c.</b> AI Application Feature
<b>IV. Model Needs</b>	<b>a.</b> Core Model <b>b.</b> Regional & Corporate Restrictions on AI & Framework <b>c.</b> Output Filtering Needs	<b>VIII. Other Notes</b>	

(a) Chapters & subchapters.

Please fill in the following statement that summarizes what your AI application would do

I intend to build an AI application that [application description] so that [target user] [end goal].

- I intend to build an AI application that [predicts that supply of beef in 6 months] so that [the government employees] [can make informed decisions about policies for stabilizing the price]
- I intend to build an AI application for [generating personalized recipes based on dietary or religious restrictions, preferences, etc. of the consumers] so that [automated kitchen owners where robots do the cooking] [can provide personalized dining experience].

...

(b) Example question.

**Figure 4: An outline of the workbook contents and a sample of the workbook content style.**

regulations are more pronounced in some domains than others. In our study, we observed stronger information needs when the AI expert needs to deal with educational data (P3, P6), military data (P11, mentioned by P10 in their previous experience), copyrighted artwork (P5, 12), medical data (mentioned by P11 in their previous experience), and personal finance data (P2).

**Implications for Documentation Form Design.** Since there are strictly more information needs if the data being handled in the domain is sensitive, the documentation format can by default ask about the regulations and copyright issues regarding the data.

## 5 THE AINEDSPLANNER WORKBOOK: DESIGN AND CASE STUDIES

Based on the results of the formative interview (Section 3) and the main study (Section 4), we built the AINEDSPLANNER workbook that outlines the information needs of the client to help support effective collaboration between AI experts and clients. In this section, we describe the design of the workbook as well as the two case studies we conducted to understand the values of the workbook in the use cases we identified through the interview in Section 3.

## 5.1 Workbook Design

In designing the workbook, we had the goal of supporting the two use cases we discovered through the formative interview:

**G1.** The workbook should serve as an effective guide for clients as they prepare information for collaboration.

**G2.** The workbook should serve as an effective discussion guide and checklist for AI experts as they elicit information from their clients.

We selected Google Docs [13] as the platform for hosting the workbook because it not only includes an outline of the chapters and subchapters that helps with easy navigation, but also allows easy collaborative edits and asynchronous communication through comments. However, the design of the workbook itself does not depend on the platform and can be ported to other platforms depending on user needs.

We followed the 5-step process outlined below to build the workbook:

**Step 1. Choosing a model documentation form.** We first reviewed various structured documentation forms (e.g., questionnaires, workbooks) within and outside of the AI application development domain to find a model format on which we can base the design of our documentation form. After rounds of internal discussions and tests on the information communication capacities of the different formats, we decided that the *workbook* format utilized by the Google People + AI guidebook [50] (example chapter worksheet: [51]), a guidebook that introduces best practices and examples for designing with AI, is a good model documentation form to start with; it is a documentation form also related to AI application development that has already been adopted in actual worksites [73].

**Step 2. Structuring the workbook into chapters & subchapters.** The decision on the holistic structure of the workbook came from the model workbook. We first took the notion of chapters into our workbook by translating the top-level nodes in the information flow taxonomy into chapters. Because of the size and complexity of our taxonomy, we further subdivided each of the chapters into subchapters for navigability, which resulted in a total of 8 chapters and 19 subchapters (Figure 4a).

**Step 3. Generating question contents.** Afterward, we generated the question contents by looking at the taxonomy as well as each instance of information needs we used to build the taxonomy. Whenever possible, we extracted the wordings that the AI experts used to express their information needs to the clients during the study in Section 4 to inform the exact wordings of the questions. For the style of the questions, we also adopted the style of the questions used in the workbook: fill-in tables, fill-in-the-blanks, and free-form questions.

**Step 4. Constructing examples.** Along the lines of G1, to further help clients, who may not have much AI expertise, comprehend and perform tasks around the workbook, we added examples or additional hints about filling it in using lighter text (e.g., Figure 4b). We constructed the examples by going through the AI application ideas list used in the study in Section 4 and trying to answer the questions ourselves and selecting the most representative samples.

**Table 3: A summary of the AI application development projects covered in the two case studies.**

Project	Project Domain	AI Expert Exp.	Client Domain Exp.
Case Study 1: Worksheet as a Client Guide			
CS1A	Medical Imaging	3 yrs	6 yrs
CS1B	Marketing / Consumer	3 yrs	5 yrs
CS1C	Oncology / Hematology	11 yrs	40 yrs
CS1D	Digital Fashion / App Design	4 yrs	3 yrs
CS1E	Medical Robotics	4 yrs	2 yrs
Case Study 2: Worksheet as a Reflection Guide for AI Experts			
CS2	Telecommunication	8 yrs	Undisclosed

**Step 5. Iterating on wording and examples.** We reflected on the suggestions for changes in wording and examples from three test runs with domain experts within our organizations.

We include the complete AINEEDSPANNER workbook in the supplemental materials.

## 5.2 Case Studies: Application in the Real World

To gain initial insights into whether we met our goals (G1, G2) for supporting the use cases identified in Section 3, we performed case studies based on the two use cases.

In addition to the two use cases, the AINEEDSPANNER workbook has been adopted by AlgorithmLabs, an AI solutions company located in South Korea. However, due to the sensitivity of the information as well as the legal issues around non-disclosure agreements, we are unable to publicly share further detail.

**5.2.1 Case Study 1: Workbook as a Client Guide.** To understand the value of the workbook as a guide for the client as they determine information in the collaboration preparation phase, we specifically recruited five domain experts with clear needs for AI application development to fill in AINEEDSPANNER over one week. We shared the filled workbooks with AI experts for initial feedback to gain insights about how this leads into the following main discussion phase.

**Participants.** In order to avoid the restrictions arising from the non-disclosure agreements between AI solutions companies and their clients, we recruited the domain experts ourselves and acted as a mediator between them and the AI solutions companies. We reached out to five domain experts who have clear needs to make use of AI applications through referrals through our professional and personal connections (Table 3 Block 1). The five domain experts were in different domains and had varying levels of domain expertise. Specifically, while the domain expert recruited for CS1A, CS1C, and CS1D had very little or no prior experience with AI, the domain expert recruited for CS1B and CS1E reported that they possess some prior experience with AI.

**Procedure.** We first met with each of the domain experts to explain the process, either on the phone or in person. During this meeting, we asked briefly about their AI application development needs, their experience in the domain, and their experience with AI. We asked them to assume that they would be working with an AI expert for the development and gave them the AINEEDSPANNER workbook that will be sent to the AI expert to begin an initial discussion. After answering any questions from the domain experts, we asked them

to fill in the workbook and return it to us with a completed short survey about their experience within one week. Once we received the completed workbooks from the domain experts, we passed them to an AI expert recruited from an AI solutions company for review and asked for comments.

Because the AI application projects are to be developed, the domain experts asked us to keep the details of the plan private for publication.

**Results.** All domain experts were able to fill in the workbook without much difficulty, including the domain experts with little or no experience in AI (CS1A, CS1C, CS1D), with the domain expert in CS1A stating that “*I was concerned that filling in the form would require knowledge of AI, but that was not the case.*” This result suggests that our workbook may be able to support even the clients without much AI knowledge. In particular, the domain experts of CS1A, CS1C, and CS1E specifically pointed out that the examples were especially helpful in knowing what to fill in for each question.

Moreover, the domain experts indicated that the workbook successfully guided them as they considered and gathered information about the AI application (CS1C, CS1D, CS1E). For example, the domain expert in CS1E stated that the workbook breaks down the various aspects of AI application planning beyond what they would have thought about on their own. In addition, the domain expert in CS1C reported that the questions of the workbook helped them actively seek specific information from a collaborator in a different institute regarding their successful introduction of a similar AI application. They reported that this helped them further develop the details of their AI application idea. In fact, the AI expert for CS1C stated that the detailed information provided in the workbook is well beyond what they would expect to be able to elicit from less prepared clients with multiple extended discussion sessions.

Based on the AI expert feedback, while the filled workbooks from client-side preparation had varying levels of imperfections, the AI experts were able to comprehend the key ideas and identify components that require further discussion. For example, in CS1B, the AI expert was able to pinpoint that they will need to obtain further discussions on the data label format and the purpose of the AI application, and that some calibrations will be needed about the domain expert’s model needs. As another example, in CS1D, the AI expert expressed specific needs to further understand the current workflow in the domain as well as discuss the quantitative target measurements.

These results suggest that the value of the workbook arises not from the fact that it is able to draw a complete AI application from the client, but from the fact that it guides clients to think about and develop AI application ideas during the collaboration preparation phase so that the AI experts can concretely reason about the plan and recognize additional discussion points. This could reduce the amount of back-and-forth communication and hence pave the way for efficient communication between clients and AI experts during the main discussion phase.

**5.2.2 Case Study 2: Workbook as a Communication & Reflection Guide for AI Experts.** The second case study is aimed at understanding the value of the workbook as a guide for the AI experts as they communicate with the client in the main discussion phase.

**Participants & Procedure.** For the second case study, we provided our workbook to an AI engineer at an AI solutions company and asked them to utilize the workbook while performing planning in their next AI application development project in any way they wish to (Table 3 Block 2). After the planning was complete, the AI engineer reported to us about their experience.

Both the client company and the AI engineer agreed to the AI engineer sharing their experience for academic publication as long as the contents of the project and the identities of the stakeholders are concealed.

**Results.** The AI engineer reported to us that they used the workbook to guide communication with the client and record the discussions.

The key value of the workbook for the AI engineer was as a tool for ensuring completeness of the communication. The AI expert specifically noted that they appreciated the completeness of the document, and especially the section on budgets, resources & framework. This is because in their prior experience, they forgot to discuss budgets and resources during the planning phase and this caused confusion later on when they discovered that the client did not have sufficient resources to deploy the model.

## 6 DISCUSSION

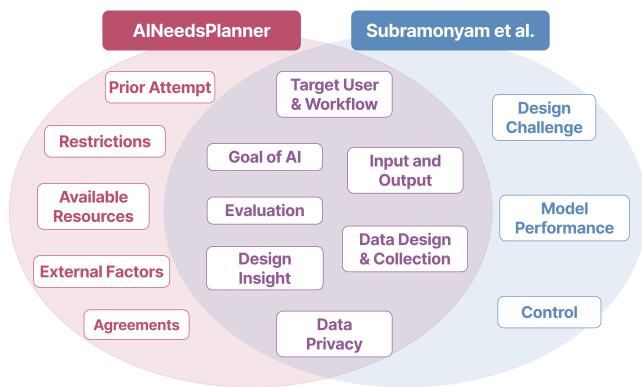
The two case studies illustrate the potential of using AINEDSPANNER for various purposes in the clients-AI experts collaboration process. In this section, we discuss the broader utility of the AINEDSPANNER workbook, incorporating more design attributes and stakeholders in the workbook by comparing it with a prior taxonomy and a potential for systematic support.

### 6.1 Broader Utility of the Workbook for Clients and AI Experts

Our case studies showed that AINEDSPANNER was able to elicit key information for building AI applications from clients, give them an awareness of the information they need to prepare before meeting with AI experts and help AI experts reflect on the information needed from the clients. AINEDSPANNER offers an easy and effective planning experience by employing the “Recognition Rather Than Recall” strategy [46], a useful design pattern for interface design. The ease of use of the workbook could encourage more active participation of the clients in the application-building process, contributing to better contextualized AI applications [55].

It also illustrates the potential for AINEDSPANNER as a learning material for the clients to understand what needs to be prepared for building AI applications. The workbook not only provides essential questions but also gives contextualized instructions (or even lectures) about the corresponding AI concepts for each section so that the clients can answer the questions in an informed way, reducing the technical knowledge gap between the clients and AI experts. On the flip side, AI experts can also benefit from having learning material for understanding the clients’ domain contexts, reducing the domain-specific knowledge gap between the two.

Additionally, AINEDSPANNER served as a boundary object facilitating shared understanding between the client and AI experts by clearly scoping, organizing, and retaining the key information needed in the planning stage. An interesting note from CS2 was that



**Figure 5: A Venn diagram showing the commonalities and differences between the information covered by AINEEDSPANNER and the information covered by Subramonyam et al. [57].**

the workbook can promote more structured discussion with clearer sub-stages in the planning. For example, some of the sections in the workbook (e.g., metrics, success thresholds) need a discussion between the two, but other sections (e.g., detailed dataset information, available budgets, and resources) mostly require information from the clients. Therefore, a possible planning structure would be first making a clear consensus about the information needed from the clients and then having a discussion for making decisions that require both stakeholders' inputs.

## 6.2 Extending Design Attributes and Stakeholders in the Workbook

AINEEDSPANNER focused on the planning stage between clients and AI experts, identifying key information needed in the planning stage. The workbook can be extended by incorporating more stakeholders as well as design attributes of AI applications.

To understand the difference between AINEEDSPANNER and known design attributes of AI applications, we compared the AINEEDSPANNER workbook with a prior taxonomy of Human-AI interaction guidelines [57] by listing dimensions of each taxonomy and grouping similar dimensions. Figure 5 shows a Venn diagram presenting the relationship between the dimensions of each taxonomy. Out of 15 dimensions of the prior taxonomy, AINEEDSPANNER covers 9 dimensions (4 training data, 2 AI models, 2 AI-powered user interface, and 1 human mental models categories from the prior taxonomy) whereas 6 dimensions (2 human mental models, 2 AI-powered user interface, and 2 AI models categories) were not the core focus of the workbook. We found three themes of the 6 dimensions (Design Challenge, Control, and Model performance) and hypothesized possible reasons for not having them as the core focus in AINEEDSPANNER.

Design challenge discusses how to shape users' mental model and interaction model, which is one of the important considerations in designing usable AI applications [1, 50]. However, such design aspects were not brought up during the planning stage as the focus was mostly on the technical side of AI applications rather than user experiences. Depending on how important it is to design usable

AI applications in the project, such design aspects can be covered in the planning potentially by incorporating designers as another stakeholder and prompting the other stakeholders to consider such design aspects through the workbook.

The control aspect was not emphasized in the planning process as well. A possible reason could be a lack of AI experiences for the clients; clients might not have needs for controlling AI as they would not expect inconsistent and imperfect behavior of AI [15]. Since it is difficult to imagine AI behavior in the planning stage [30, 72], prototyping AI behavior [58] even in the planning stage could facilitate the discussion of AI control.

Finally, AINEEDSPANNER does not focus on how to ensure model performance for diverse use cases. A possible reason would be around the initial scope of the target AI application; clients would be interested in building AI applications for their specific contexts and use cases first for feasibility checking, which can be expanded later once the application is successfully deployed and used.

The planning workbook can incorporate these important design attributes as well, depending on the characteristics of clients' needs and contexts. It illustrates the need for extending the workbook that can be adapted to various stakeholder combinations. For example, having designers in the planning stage, as well as clients and AI experts, may need additional sets of questions, which helps create a more robust plan considering both technical and user experience in the early stage. The managerial role is another important stakeholder who has the power of decision-making as well as budget management, which might need more sets of detailed practical questions. As AINEEDSPANNER can be a good starting point, it can be an effective approach to start planning with AINEEDSPANNER in such diverse settings, come up with questions that are important in such specific context, and design another version of the workbook that can be more specialized in the contexts.

## 6.3 Towards Designing Interactive Support for the Planning Process

The AI application planning through a collaboration between clients and AI experts is a complicated process, which brings multiple challenges, such as eliciting information from clients who are not knowledgeable in AI, reducing the knowledge gap between clients and AI experts for the client contexts, and AI, and creating plans that align with the clients' AI needs. AINEEDSPANNER addresses such challenges as a boundary object that facilitates clear communications on the same page. AINEEDSPANNER is a static documentation format, but we can go beyond that by designing interactive systems that provide a more adaptive planning experience, considering the contexts of planning (e.g., types of AI applications, participating stakeholders, and knowledge gaps between stakeholders).

As we noted in the prior discussion, the design attributes that need to be considered in the planning depend on the context of planning, such as client domain context and participating stakeholders. Based on important design attributes to be considered, the system can adaptively organize the questions, leveraging questions in AINEEDSPANNER. For example, designing AI applications for a high-stake scenario (e.g., medical domain [7]) might need to put more emphasis on the explainability and controllability of AI with

designers in the collaboration process. Then, the system might ask questions about existing practices in verifying prediction results.

Eliciting information from the clients can also benefit from having interactive systems. Research showed that an automated chatbot can elicit quality responses compared to an online survey, while providing engaging experiences [70]. Asking follow-up questions [21] can be effective in getting detailed and well-aligned responses for the question. Scaffolding the responses via autocompletion by sharing client contexts (e.g., suggesting task-specific measures for the types of AI applications and suggesting potentially relevant features to the label when describing data formats) can reduce clients' efforts in the knowledge elicitation process.

As a prototype is commonly used as a boundary object in software development [32], supporting quick prototyping of AI applications based on the planning contents has the potential to significantly help stakeholders be on the same page. Automated techniques such as AutoML [26] and UI prototyping techniques [58] can be employed, but it needs more investigations to create such prototypes based on the planning contents. Another promising technique to have a concrete artifact could be finding relevant projects in the wild (e.g., from Papers With Code [9]) to assess the feasibility of the current plan.

## 7 LIMITATIONS AND FUTURE WORK

**Conducting more in-depth case studies in real world.** Although our two case studies suggest that AINEDSPANNER does support client-AI expert collaboration, the coverage of project domains was limited and for CS2, we only had limited access to just the AI expert's experience due to a non-disclosure agreement between the AI expert and the client. Hence, we believe that more in-depth studies observing both sides of the client-AI expert collaboration in more cases would help us further our understanding about the value of our workbook. While AINEDSPANNER is a workbook for generic AI applications and we have identified dependencies of information needs specific on certain data types (Section 4.4.3, we believe that the in-depth studies could lead to insights that can lead to more specialized workbooks tailored to specific application domains or application types. Future work can also study put our workbook to use in other types of collaborations, such as intra-team collaborations or collaborations with additional collaborators (e.g., UI/UX designers, managerial roles) and find ways to generalize our workbook.

**Guiding assessments of risks and unintended consequences.** Assessing the risks and unintended consequences of AI applications has been widely recognized [65]. Some of the AI experts in the main study hinted at the need to understand the risks and consequences. However, they considered them as components of other information needs, not as an independent information need on their own. For example, P10 expressed information needs on potential harms of certain types of errors, but as a part of understanding and determining the metrics and the target values; P11 expressed information needs on the 'seriousness' of the consequences of the project while trying to understand nature of the project. Based on these results, we include a question about error type asymmetry within the section on defining success and a question about consequences of malfunction within the section on project and client

characteristics. Yet, we believe that the importance of discussions about risks and unintended consequences of AI applications will continue to rise with rapidly growing AI capabilities. Hence, we suggest that future work further explores the information needs around risks and unintended consequences to form an independent section with additional emphasis. We believe that this will steer discussions between AI experts and their clients towards responsible, ethical, and safe usage of AI.

### **Supporting input formats for the workbook beyond plain text.**

The current form of AINEDSPANNER only allows the users to express themselves using plain text. While text is an expressive medium that can capture most if not all intents, it is not necessarily the most efficient medium for expressing, communicating, and recording information. For example, while the workflow in the domain can be expressed in text, formats like a flow chart may be more effective in communicating the information. To further support this, the AI expert in CS2 of our user studies further supported the idea and suggested that although they could fill in the workbook in the current format, they would have had an easier time filling the workbook if it had supported free-form sketches or diagrams. Based on these observations, future work could either integrate tools that can accept input beyond plain text (e.g., Figma [18], Miro [43]) or conversely attempt to build our workbook into these tools.

**Adaptations to the shifting AI landscape.** The field of AI has been moving rapidly in recent years with frequent introduction of novel models and APIs (e.g., generative models). Because it is a recent invention, we had only one sample of *successfully concluded* projects around generative models in our interview in Section 3 (IT4). On the contrary, many AI experts from the interview in Section 4 shared a desire to utilize generative models not only as the core model (P5, 6, 8, 12), but also as a way to augment and simulate data (P7, 10, 11), showing the rapid pace of adoption of this new technology. In addition, governments around the globe including the US [31] and the EU [47] have been moving diligently towards introducing regulations and directives on AI, especially in recent years. While we believe that the overall structure and contents of planning would remain less affected by the development of new models, we admit that our understanding and the workbook may need adaptation over time with the shifting AI landscape that will certainly continue. For example, we hypothesize that information needs on the data may decrease as LLMs have enabled few-shot and zero-shot learning approaches [37]. We leave monitoring the effects of the endlessly shifting AI landscape on planning as well as making the required adaptations to the workbook as future work.

## 8 CONCLUSION

Client-AI expert collaborations in AI application development are commonplace, and the planning stage is a key stage of this collaboration in which the client collects information about their goals, expectations, and resources to the AI expert for them to sketch out a technical approach to the development. We examine the planning stage in AI application development in detail through an interview surveying 10 successfully concluded AI application development projects. Through the interview study, we identify the steps of the planning stage along its three phases and deduce how a documentation form outlining the information needs of the AI expert from the

client can facilitate the collaboration. Next, we built a taxonomy of the AI experts' information needs and uncovered the characteristics of the AI development project that affect the information needs through an experiment on 12 AI experts. Based on our findings, we introduce AINEEDSPLANNER, a workbook designed to support effective collaboration between AI experts and clients. Through two case studies, we find that the workbook not only can serve as an information collection and assessment guide for clients as they prepare for collaboration, but also as a discussion guide for AI experts that helps AI experts ensure the completeness of the information flow from the clients. Although we verify the values of the workbook as it stands, the workbook will require continuous adaptation to the quickly changing AI landscape to provide maximal value to the client-AI expert collaboration in AI application planning.

## ACKNOWLEDGMENTS

The authors thank Yoonsu Kim and Nicole Lee for their early contributions during the formative interview and initial explorations of the workbook format. The authors also thank the AI engineers at AlgorithmLabs and the members of KAIST Interaction Lab for their constructive feedback throughout the work and piloting the formative interview and the main study. This work is supported by the G-CORE Research Project grant at KAIST. Hariharan Subramonyam is supported by NSF award IIS-2302701.

## REFERENCES

- [1] 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* (Glasgow, UK). ACM, New York, USA, 1–13. <https://doi.org/10.1145/3290605.3300233>
- [2] Apple. 2022. Human Interface Guidelines - Human Interface Guidelines - Design - Apple Developer. Retrieved January 20, 2024 from <https://developer.apple.com/design/human-interface-guidelines/guidelines/overview/>
- [3] Amid Ayobi, Katarzyna Stawarz, Dmitri Katz, Paul Marshall, Taku Yamagata, Raúl Santos-Rodríguez, Peter Flach, and Aisling Ann O'Kane. 2021. Machine Learning Explanations as Boundary Objects: How AI Researchers Explain and Non-Experts Perceive Machine Learning. In *2021 Joint ACM Conference on Intelligent User Interfaces Workshops* (College Station, USA). CEUR Workshop Proceedings, College Station, USA, 9 pages. <http://ceur-ws.org/Vol-2903/IUI21WS-TESS-3.pdf>
- [4] Alex Bäuerle, Ángel Alexander Cabrera, Fred Hohman, Megan Maher, David Koski, Xavier Suau, Titus Barik, and Dominik Moritz. 2022. Symphony: Composing Interactive Interfaces for Machine Learning. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA). ACM, New York, USA, Article 210, 14 pages. <https://doi.org/10.1145/3491102.3502102>
- [5] Aditya Bharadwaj, Pao Siangliulue, Adam Marcus, and Kurt Luther. 2019. Critter: Augmenting Creative Work with Dynamic Checklists, Automated Quality Assurance, and Contextual Reviewer Feedback. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, UK). ACM, New York, USA, 1–12. <https://doi.org/10.1145/3290605.3300769>
- [6] Veronika Bogina, Alan Hartman, Tsvi Kuflik, and Avital Shulner-Tal. 2021. Educating Software and AI Stakeholders about Algorithmic Fairness, Accountability, Transparency and Ethics. *International Journal of Artificial Intelligence in Education* 32 (2021), 1–26. <https://doi.org/10.1007/s40593021002480>
- [7] Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 3 (2019), 1–24. <https://doi.org/10.1145/3359206>
- [8] Paul R. Carlile. 2002. A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development. *Organization Science* 13, 4 (2002), 442–455. <https://doi.org/10.1287/orsc.13.4.442.2953>
- [9] Papers With Code. 2024. Retrieved January 20, 2024 from <https://paperswithcode.com/>
- [10] Thomas H Davenport, Rajeev Ronanki, et al. 2018. Artificial Intelligence for the Real World. *Harvard Business Review* 96, 1 (2018), 108–116.
- [11] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond" Add Diverse Stakeholders and Stir". *arXiv preprint arXiv:2111.01122* 1 (2021), 7 pages.
- [12] Frédéric Delmar and Scott Shane. 2003. Does Business Planning Facilitate the Development of New Ventures? *Strategic Management Journal* 24, 12 (2003), 1165–1185. <https://doi.org/10.1002/smj.349>
- [13] Google Docs. 2024. Retrieved January 20, 2024 from <https://docs.google.com/>
- [14] 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* (Denver, USA). ACM, New York, USA, 278–288. <https://doi.org/10.1145/3025453.3025739>
- [15] Mary T Dzindolet, Scott A Peterson, Regina A Pomranky, Linda G Pierce, and Hall P Beck. 2003. The Role of Trust in Automation Reliance. *International journal of human-computer studies* 58 (2003), 697–718. [https://doi.org/10.1016/S1071-5819\(03\)00038-7](https://doi.org/10.1016/S1071-5819(03)00038-7)
- [16] Vladimir Estivill-Castro, Eugene Gilmore, and René Hexel. 2022. Constructing Explainable Classifiers from the Start—Enabling Human-in-the-Loop Machine Learning. *Information* 13, 10 (2022), 464. <https://doi.org/10.3390/info13100464>
- [17] Hugging Face. 2024. Model Cards. Retrieved January 20, 2024 from <https://huggingface.co/docs/hub/en/model-cards>
- [18] Figma. 2024. Retrieved January 20, 2024 from <https://www.figma.com/>
- [19] Jules Françoise, Baptiste Caramiaux, and Téo Sanchez. 2021. Marcelle: Composing Interactive Machine Learning Workflows and Interfaces. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA). ACM, New York, USA, 39–53. <https://doi.org/10.1145/3472749.3474734>
- [20] Gartner. 2019. Gartner Predicts The Future Of AI Technologies. Retrieved January 20, 2024 from <https://www.gartner.com/smarterwithgartner/gartner-predicts-the-future-of-ai-technologies>
- [21] Yubin Ge, Ziang Xiao, Jana Diesner, Hen Ji, Karrie Karahalios, and Hari Sundaram. 2023. What Should I Ask: A Knowledge-Driven Approach for Follow-up Questions Generation in Conversational Surveys. In *Proceedings of the 37th Pacific Asia Conference on Language, Information and Computation* (Hong Kong, China). ACL, Stroudsburg, USA, 113–124. <https://aclanthology.org/2023.paclic-1.12>
- [22] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. *Communications of the ACM* 64, 12 (2021), 86–92. <https://doi.org/10.1145/3458723>
- [23] Leo A. Goodman. 1961. Snowball Sampling. *The Annals of Mathematical Statistics* 32, 1 (1961), 148–170. <https://doi.org/10.1214/aoms/1177705148>
- [24] Google. 2024. Model Cards. Retrieved January 20, 2024 from <https://modelcards.withgoogle.com/about>
- [25] John R Hayes and Jane Gradwohl Nash. 2013. On the Nature of Planning in Writing. *The Science of Writing: Theories, Methods, Individual Differences and Applications* 1 (2013), 29.
- [26] Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. AutoML: A Survey of the State-of-the-Art. *Knowledge-Based Systems* 212 (2021), 106622. <https://doi.org/10.1016/j.knosys.2020.106622>
- [27] Jeffrey Heer. 2019. Agency Plus Automation: Designing Artificial Intelligence into Interactive Systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850. <https://doi.org/10.1073/pnas.1807184115>
- [28] Amy K. Heger, Liz B. Marquis, Mihaela Vorvoreanu, Hanna Wallach, and Jennifer Wortman Vaughan. 2022. Understanding Machine Learning Practitioners' Data Documentation Perceptions, Needs, Challenges, and Desiderata. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2, Article 340 (2022), 29 pages. <https://doi.org/10.1145/3555760>
- [29] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, 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* (Glasgow, UK). ACM, New York, USA, 1–16. <https://doi.org/10.1145/3290605.3300830>
- [30] Matthew K. Hong, Adam Fourney, Derek DeBellis, and Saleema Amershi. 2021. Planning for Natural Language Failures with the AI Playbook. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan). ACM, New York, USA, Article 386, 11 pages. <https://doi.org/10.1145/3411764.3445735>
- [31] The White House. 2023. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. Retrieved January 20, 2024 from <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- [32] Thomas L. Huber, Maik A.E. Winkler, Jens Dibbern, and Carol V. Brown. 2020. The Use of Prototypes to Bridge Knowledge Boundaries in Agile Software Development. *Information Systems Journal* 30, 2 (2020), 270–294. <https://doi.org/10.1111/isj.12261>
- [33] IBM. 2021. Infuse IBM Artificial Intelligence Field Guide. Retrieved January 20, 2024 from <https://ac-gm-static-files-server.lahgrqm5xee.au-syd.codeengine>.

- appdomain.cloud/cloud/architecture/files/ai-analytics-fg.pdf
- [34] Kaggle. 2024. Model Cards. Retrieved January 20, 2024 from <https://www.kaggle.com/code/var0101/model-cards>
- [35] Jiro Kawakita. 1982. *The Original KJ Method*. Tokyo: Kawakita Research Institute.
- [36] Jiro Kawakita. 1986. *The KJ Method: Let Chaos Tell*. Chuokoron-sha. (in Japanese).
- [37] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large Language Models are Zero-Shot Reasoners. *Advances in Neural Information Processing Systems* 35 (2022), 22199–22213.
- [38] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, USA). ACM, New York, USA, 1–15. <https://doi.org/10.1145/3313831.3376590>
- [39] Josh Lovejoy. 2019. Human-Centered AI Cheat-Sheet. Retrieved January 20, 2024 from <https://uxdesign.cc/human-centered-ai-cheat-sheet-1da130ba1bab>
- [40] Yaoli Mao, Dakuo Wang, Michael Muller, Kush R. Varshney, Ioana Baldini, Casey Dugan, and Aleksandra Mojsilović. 2019. How Data Scientists Work Together With Domain Experts in Scientific Collaborations: To Find The Right Answer Or To Ask The Right Question? *Proceedings of the ACM on Human-Computer Interaction* 3 (2019), 1–23. <https://doi.org/10.1145/3361118>
- [41] Google Meet. 2024. Retrieved January 20, 2024 from <https://meet.google.com/>
- [42] Microsoft. 2022. Microsoft Responsible AI Standard, v2. Retrieved January 20, 2024 from <https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-Responsible-AI-Standard-v2-General-Requirements-3.pdf>
- [43] Miro. 2024. Retrieved January 20, 2024 from <https://miro.com/>
- [44] 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* (Atlanta, USA). ACM, New York, USA, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [45] Kichan Nam, Christopher S Dutt, Prakash Chathoth, Abdelkader Daghigh, and M Sajid Khan. 2021. The Adoption of Artificial Intelligence and Robotics in the Hotel Industry: Prospects and Challenges. *Electronic Markets* 31 (2021), 553–574.
- [46] Jakob Nielsen. 1994. Usability Heuristics for User Interface Design. Retrieved January 20, 2024 from <https://www.nngroup.com/articles/ten-usability-heuristics/>
- [47] European Parliament. 2023. EU AI Act: First Regulation on Artificial Intelligence. Retrieved January 20, 2024 from <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- [48] David Piorkowski, Soya Park, April Yi Wang, Dakuo Wang, Michael Muller, and Felix Portnoy. 2021. How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1, Article 131 (2021), 25 pages. <https://doi.org/10.1145/3449205>
- [49] Sajjadur Rahman, Pao Siangliulue, and Adam Marcus. 2020. MixTAPE: Mixed-initiative Team Action Plan Creation Through Semi-structured Notes, Automatic Task Generation, and Task Classification. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2, Article 169 (2020), 26 pages. <https://doi.org/10.1145/3415240>
- [50] Google People + AI Research. 2019. Google People + AI Guidebook. Retrieved January 20, 2024 from <https://pair.withgoogle.com/guidebook/>
- [51] Google People + AI Research. 2019. User Needs + Defining Success Chapter Worksheet. Retrieved January 20, 2024 from <https://pair.withgoogle.com/worksheet/user-needs.pdf>
- [52] Bruce Saddler, Susan Moran, Steve Graham, and Karen R Harris. 2004. Preventing Writing Difficulties: The Effects of Planning Strategy Instruction on the Writing Performance of Struggling Writers. *Exceptionality* 12, 1 (2004), 3–17. [https://doi.org/10.1207/s15327035ex1201\\_2](https://doi.org/10.1207/s15327035ex1201_2)
- [53] Dan E Schendel and Charles W Hofer. 1979. *A New View of Business Policy and Planning*. Little, Brown & Company, Boston, USA. [https://cdn.strategicmanagement.net/conferences/houston/overview/overview\\_rightColumn/book-download2/s3file.pdf](https://cdn.strategicmanagement.net/conferences/houston/overview/overview_rightColumn/book-download2/s3file.pdf)
- [54] Zheyuan Ryan Shi, Claire Wang, and Fei Fang. 2020. Artificial Intelligence for Social Good: A Survey. *arXiv preprint arXiv:2001.01818* 1 (2020), 78 pages.
- [55] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation Is not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (Arlington, USA). ACM, New York, USA, Article 1, 6 pages. <https://doi.org/10.1145/3551624.3555285>
- [56] Statista. 2023. Artificial Intelligence - Worldwide. Retrieved January 20, 2024 from <https://www.statista.com/outlook/tmo/artificial-intelligence/worldwide>
- [57] Hariharan Subramonyam, Jane Im, Colleen Seifert, and Eytan Adar. 2022. Solving Separation-of-Concerns Problems in Collaborative Design of Human-AI Systems through Leaky Abstractions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, USA) (CHI '22). ACM, New York, USA, Article 481, 21 pages. <https://doi.org/10.1145/3491102.3517537>
- [58] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. ProtoAI: Model-Informed Prototyping for AI-Powered Interfaces. In *Proceedings of the 26th International Conference on Intelligent User Interfaces* (College Station, USA). ACM, New York, USA, 48–58. <https://doi.org/10.1145/3397481.3450640>
- [59] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. Towards A Process Model for Co-Creating AI Experiences. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA). ACM, New York, USA, 1529–1543. <https://doi.org/10.1145/3461778.3462012>
- [60] Hariharan Subramonyam, Colleen Seifert, and MI Eytan Adar. 2021. How Can Human-Centered Design Shape Data-Centric AI?. In *Proceedings of NeurIPS Data-Centric AI Workshop*. Curran Associates, Inc., Red Hook, USA, 3 pages. [https://haridecoded.com/resources/AIX\\_NeurIPS\\_2021.pdf](https://haridecoded.com/resources/AIX_NeurIPS_2021.pdf)
- [61] Mei Tan, Hansol Lee, Dakuo Wang, and Hari Subramonyam. 2024. Is a Seat at the Table Enough? Engaging Teachers and Students in Dataset Specification for ML in Education. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1, Article 81 (2024), 32 pages. <https://doi.org/10.1145/3637358>
- [62] Rui A. Alves Teresa Limpo. 2018. Effects of Planning Strategies on Writing Dynamics and Final Texts. *Acta Psychologica* 188 (2018), 97–109. <https://doi.org/10.1016/j.actpsy.2018.06.001>
- [63] Upwork. 2024. Retrieved January 20, 2024 from <https://www.upwork.com/>
- [64] Philip van Allen. 2018. Prototyping Ways of Prototyping AI. *Interactions* 25, 6 (2018), 46–51.
- [65] Ibo van de Poel. 2023. *AI, Control and Unintended Consequences: The Need for Meta-Values*. Springer International Publishing, Cham, Switzerland, 117–129. [https://doi.org/10.1007/978-3-031-25233-4\\_9](https://doi.org/10.1007/978-3-031-25233-4_9)
- [66] Michael Veale, Max van Kleek, and Reuben Binns. 2018. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada). ACM, New York, USA, 1–14. <https://doi.org/10.1145/3173574.3174014>
- [67] Janet Vertesi and Paul Dourish. 2011. The value of data: considering the context of production in data economies. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work* (Hangzhou, China). ACM, New York, USA, 533–542. <https://doi.org/10.1145/1958824.1958906>
- [68] Dakuo Wang, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi, and Qianying Wang. 2020. From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, USA). ACM, New York, USA, 1–6. <https://doi.org/10.1145/3334480.3381069>
- [69] Marian G. Williams and Vivienne Begg. 1993. Translation Between Software Designers and Users. *Commun. ACM* 36, 6 (1993), 102–103. <https://doi.org/10.1145/153571.214831>
- [70] Ziang Xiao, Michelle X. Zhou, Q. Vera Liao, Gloria Mark, Changyan Chi, Wenxi Chen, and Huahai Yang. 2020. Tell Me About Yourself: Using an AI-Powered Chatbot to Conduct Conversational Surveys with Open-ended Questions. *ACM Transactions on Computer-Human Interaction* 27, 3, Article 15 (2020), 37 pages. <https://doi.org/10.1145/3381804>
- [71] 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). ACM, New York, USA, 585–596. <https://doi.org/10.1145/3196709.3196730>
- [72] 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* (Honolulu, USA). ACM, New York, USA, 1–13. <https://doi.org/10.1145/3313831.3376301>
- [73] Nur Yildirim, Mahima Pushkarna, Nitesh Goyal, Martin Wattenberg, and Fernanda Viégas. 2023. Investigating How Practitioners Use Human-AI Guidelines: A Case Study on the People + AI Guidebook. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). ACM, New York, USA, Article 356, 13 pages. <https://doi.org/10.1145/3544548.3580900>
- [74] Zoom. 2024. Retrieved January 20, 2024 from <https://zoom.us/>

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009