Scilight

Human-Machine Interaction
https://www.sciltp.com/journals/hmi



Article

Structuring Early Design of AI-Infused Information-Intensive Applications

Xiang 'Anthony' Chen

UCLA HCI Research, University of California, Los Angeles, CA 90095, USA; xac@ucla.edu

How To Cite: Chen, X.A. Structuring Early Design of AI-Infused Information-Intensive Applications. *Human-Machine Interaction* **2026**, *1*(1), 5.

Received: 15 June 2026
Revised: 20 August 2026
Accepted: 12 October 2026
Published: 30 October 2026

Abstract: AI holds the promise of augmenting a wide range of knowledge workers' tasks, many of which involve collecting, organizing, interpreting, comparing, transforming, and acting on domain-specific information. Designing such *AI-infused information-intensive applications* (A3I) is difficult in the early, divergent stage because designers must decide what role AI should play, where AI should enter a human information-processing workflow, and how initiative should be divided between people and AI. Existing human-AI guidelines and pattern collections provide important support, but they do not necessarily offer a systematic way to surface and compare these early design possibilities before teams commit to an interaction direction. We contribute A3I: a conceptual framework and workbook for structuring early-stage design exploration of AI-infused information-intensive applications. The framework organizes exploration around two design questions. First, what can an AI system do to help a human process information? We distinguish sense-making (via a hierarchic view and filters) from decision-making (using predictive and/or generative AI). Second, when and how should AI be invoked as a human processes information? We consider AI automation (revealing findings upfront or in due course) and mixed initiative (AI acting reactively and/or proactively). We operationalized the framework as the A3I WORKBOOK and conducted a formative, three-part evaluation: (i) deployment in graduate-level course projects (25 students across nine teams); (ii) expert interviews with three industrial researchers; and (iii) a week-long design exercise with seven UX designers in industry. Findings document how participants interpreted and used the workbook, including reported value in structuring thought and noticing previously implicit design dimensions, as well as interpretive difficulties around hierarchy. We discuss lessons for translating HCI theory into actionable design scaffolds and clarify the limits of our evidence, including the need for future comparative and longitudinal studies.

Keywords: AI-infused systems; information processing; ideation; interaction design

1. Introduction

Classic HCI research models a human as an information processor [1]. Indeed, tasks for knowledge workers often involve processing a significant amount of information, such as analysts gathering intelligence in overhead imagery, biologists detecting wildlife activities in video footages, physicians finding signs of tumors on tissue slides, and designers defining the next fashion trend from a library of existing examples. Zarmer and Chew characterized systems that support such tasks as information-intensive applications [2], often used in specialized and complex domains where the visual organization and interaction is critical to how users perform information-processing tasks. In this paper, we use the term operationally to describe applications whose primary user task involves collecting, organizing, interpreting, comparing, transforming, or acting on information. This category is defined by the task rather than by a particular interface modality: A dashboard with natural-language querying, a conversational tool for synthesizing documents, and a hybrid visual-conversational review browser may all be information-intensive if they help users process substantial information. By contrast, applications centered mainly on physical manipulation,

Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
Publisher's Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

generic chat, or delegated command execution may only partially fit AI-infused information-intensive (dubbed as A3I) applications unless the user's central challenge is still to make sense of and decide from information.

Recent advances in AI (defined by Riedl [3] as "algorithms that perform tasks or behaviors that a person could reasonably deem to require intelligence if a human were to do it"; we use "AI" and "system" interchangeably, although not all features in an AI-infused system are AI-based) has given rise to a new class of AI-infused [4,5] information-intensive (A3I) applications to assist knowledge workers. In this paper, we developed a conceptual framework and workbook for structuring the early ideation stage of creating A3I applications. This early stage is difficult because designers must make several interdependent commitments before there is a stable design to evaluate: what role AI should play, where AI should enter a human information-processing workflow, and how initiative should be divided between people and AI. Rather than claiming that designers have no support for early AI design, we argue more narrowly that existing methods do not necessarily provide a systematic way to surface and compare the particular possibilities formalized by A3I. Designers can use our A3I framework after initial need-finding studies with target users and before converging on a specific interaction direction, where they seek to answer two design questions.

1.1. What Can an AI System Do to Help a User Process Information?

When processing a large amount of information, it is important to first develop the right representation of information chunks [1] (i.e., the process of sense-making [6]) before examining each chunk (i.e., decision-making). We can infuse AI in both sense-making and decision-making phases:

- In the sense-making phase, AI can help structure the information into a hierarchic representation, such as dividing a medical image into nested regions of interest (ROIs). In addition, AI can filter in a subset of information based on domain-specific or user-defined criteria, e.g., prioritizing which subsets of the ROIs a physician should examine first.
- In the decision-making phase, AI can help a user process selected information chunks predictively (e.g., identifying the key sentence that summarizes a paragraph) and/or (the two are not mutually exclusive) generatively (e.g., AI-generated questions about a paragraph to help a reader better understand the text).

1.2. When and How to Invoke AI as a User Processes Information?

Design can enable users to delegate one of the information-processing tasks for AI to automate or to work mixed-initiatively [7] with AI.

- If AI alone can automate most of the task, we can present AI's findings upfront or in due course. While it seems straightforward to show AI's findings upfront, past research [8,9] suggests that revealing AI's findings in due course (e.g., on-demand or with a wait time) reduces overreliance on AI at the cost of an increased cognitive load and a decreased user experience.
- If AI still needs users to work together mixed-initiatively, we can explore a spectrum of design choices between AI behaving reactively (i.e., acting only upon a user's request) and proactively (i.e., able to anticipate a user's needs and act on its own).

In this paper, we formalize the A3I framework with theoretical foundations and existing design examples. We further developed the framework into the A3I WORKBOOK with more detailed instructional language that one can follow to ideate and explore a design of A3I applications (available at https://hci.ucla.edu/a3i and in the supplementary material). We then conducted a formative, multi-context assessment of the workbook as one implementation of the framework, focusing on understandability, applicability, friction points, transferability across contexts, and opportunities for refinement.

First, we deployed an early version of the workbook in a graduate-level class on human-AI interaction where 25 students in nine project teams used the workbook to create a high-fidelity prototype of an AI information assistant from scratch. We found that each project team went through several stages of using the workbook to apply the framework to their project, from understanding the fundamental concepts, to interpreting the concepts in the context of a project, and to adapting or extending the prompts for a project's specific workflow. Project teams reported using the workbook to broaden, deepen, structure, and reconsider their ideas, while their workbooks and interview walkthroughs documented conceptual ambiguities that informed revisions.

Second, to understand how the workbook can support professional designers, we conducted in-depth interviews with 3 current or former industrial researchers who had research experiences on AI-infused systems in a professional setting. The expert interviews surfaced a set of issues different from the aforementioned deployment study, such as the needs to help designers decide whether the workbook is relevant to their projects, to use non-academic and designer-friendly language, and to reduce overall workload due to professionals' busyness.

Finally, after synthesizing findings from the above two studies that iteratively improved the workbook, we recruited 7 UX designers in a North American software company to conduct a week-long design exercise. Participants used the A3I WORKBOOK to guide the design of a product that infused AI to help e-commerce users process a large number customer reviews to make purchasing decisions. While participants differed in how they expected to use the A3I WORKBOOK in their work, they reported that the workbook's main value lay in supporting a structured thought process and helping them discover design dimensions they had not explicitly considered. At the same time, participants mainly struggled to understand and respond to the hierarchy-related step, revealing a broader translation problem between theoretically meaningful framework concepts and practitioners' existing vocabulary around information architecture, interface hierarchy, and workflow structure. Our findings also surface the trade-offs of keeping the workbook lightweight and providing detailed instructions when needed.

1.3. How A3I Differs from and Complements Existing Guidance

As shown in Table 1, HAX, PAIR, Amershi et al., and pattern libraries help teams plan, instantiate, or assess human–AI interactions at different stages. In contrast, A3I has a narrower role—supporting the early stage of surfacing alternatives for AI role, involvement, and initiative in information-intensive applications, before a team commits to a direction.

Table 1. How A3I complements representative human–AI design guidance.

Resource	Design Stage	Scope	Typical Output
A3I	Early divergent ideation after initial need finding	Information-intensive applications; AI role, timing, and initiative	Alternative workflows, sketches, AI-role descriptions, and assumptions
Microsoft HAX Toolkit [10]	Requirements, design, and prototyping	Broad user-facing AI; interaction requirements and common failure modes	Critiques, detection of failure, and mitigation plans
Google PAIR Guidebook [11]	Across product development	Broad human-centered AI product practice	Methods, worksheets, workshop outputs, and design patterns
Amershi et al. [4]	Interaction planning and evaluation across use over time	Broad human–AI interaction	Guideline checklist and assessment criteria
Interaction pattern libraries [12]	Requirements and early prototyping after a design problem is identified	Recurring interaction problems	Candidate solution patterns and examples

The main contribution of this paper is a theoretically grounded conceptual framework that organizes the design space of AI support for information-intensive work. The framework's two design questions are not research questions to be answered by the studies; rather, they are the organizing questions designers use to explore alternatives. The paper further contributes (i) the A3I WORKBOOK, which operationalizes the framework into an early-stage ideation process; (ii) formative empirical findings about how students, industrial researchers, and UX professionals interpreted and used the workbook; and (iii) broader lessons about translating HCI theory into actionable design scaffolds.

In the remainder of this paper, we first overview the A3I framework, then provide a detailed walkthrough of its two organizing design questions while reviewing related work. Next, we present findings from deploying the A3I WORKBOOK in an academic environment, followed by expert interviews and a week-long task-based study in an industrial setting. Finally, we close by discussing key issues, limitations, and future work.

2. A3I Overview

Before delving into the details of the A3I framework, we first overview its intended use, construction process, and relationship to the workbook.

2.1. Using A3I as an Early-Stage Thinking Aid

A3I is intended to support divergence before commitment. After a team has a preliminary understanding of users and their information-processing task, the framework prompts the team to enumerate alternatives: what information needs to be organized, filtered, interpreted, or transformed; which of these activities AI might support; and whether AI should act upfront, in due course, reactively, or proactively. The accompanying workbook operationalizes these prompts as a step-by-step structure so designers can make assumptions explicit, compare plausible options, and reduce the chance of overlooking a major design dimension. This structure is not meant to replace brainstorming, user research, prototyping, or established human-AI guidelines. It is also not a lookup catalog, an end-to-end development process, a

feasibility validator, or evidence that a resulting design will be better than one produced by another method. Instead, A3I helps teams produce intermediate design artifacts, such as opportunity statements, sketches, workflow alternatives, AI behavior descriptions, and assumptions that can be examined in later design and engineering work. Table 2 defines the terms used throughout the framework and workbook.

Table 2. Nomenclature used in the A3I framework and workbook. Definitions are operational within this paper.

Term	Operational Definition
A3I application	An AI-infused application whose primary user task is collecting, organizing, interpreting, comparing, transforming, or acting on substantial domain-specific information.
A3I framework/workbook	The conceptual organization of the design space/its operationalization as prompts, examples, and sketching spaces.
Information chunk	A meaningful unit of information that the user can process as a single entity for the task at hand.
Sense-making/decision-making	Organizing and focusing information to establish a useful representation/interpreting a selected chunk to reach a judgment, choice, or insight.
Hierarchy/filter	Organizing information across levels or groups/reducing the information requiring attention according to explicit or inferred criteria.
Predictive/generative support	AI support that helps narrow toward a judgment, classification, or ranking/support that produces, retrieves, or synthesizes information to expand or transform the available alternatives.
Automation/mixed initiative	AI performs most of an information-processing activity/human and AI contribute through interleaved actions.
Upfront/in due course	Automated findings are revealed before the user's own action and assessment/after some user engagement or on demand.
Reactive / proactive	AI acts following an explicit user action / anticipates a need and initiates an action.

2.2. Constructing the Framework

We constructed A3I through a top-down synthesis of information processing, sense-making, overreliance, and mixed-initiative interaction and a bottom-up coding of the 34 system papers summarized in Appendix A. The first author coded each artifact's information operation, type of AI support, workflow, and invocation behavior, then iteratively merged recurring codes into higher-level concepts. Table A1 shows how the resulting concepts map to individual systems, while Table A2 describes the corpus.

We retained two necessary questions because the coding repeatedly separated what support AI provides from when and how that support enters the workflow. We do not treat these questions as orthogonal, or expect them to follow a linear sequence. They interact in use: a team may revisit invocation after changing AI's role, and sense-making and decision-making may repeat or reverse order, as when T2 proposed filtering AI-generated findings. Figure 1 therefore presents an expository path through the questions rather than a fixed process model.

The key concepts foreground information-specific operations and interaction design choices that transfer across domains and modalities. Hierarchy is one recurring organization strategy and an ideation prompt, not a universal prerequisite or an exhaustive account of information organization; filtering foregrounds selective attention. Predictive and generative support distinguish narrowing toward a judgment from expanding or transforming available information, while automation timing and the reactive–proactive continuum expose choices about reliance, effort, control, and anticipation. We did not organize the framework by UI component, model family, modality, or application domain because those categories describe implementations or contexts rather than the cross-cutting design alternatives A3I is intended to surface.

The conceptual framework is the paper's primary contribution. The workbook is its operationalization into prompts, examples, and sketching spaces, and the studies evaluate how that operationalization was understood and used, rather than whether the framework produces objectively better designs.

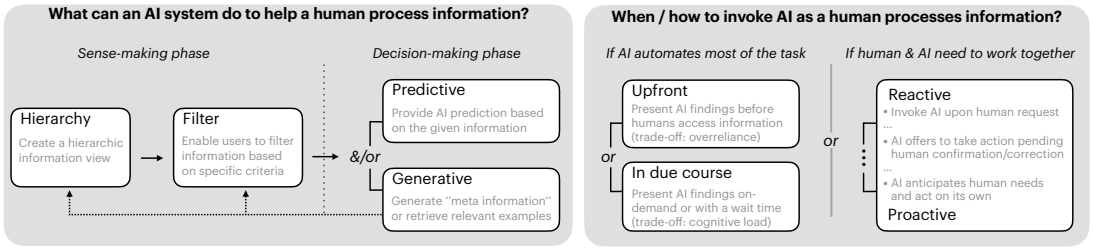


Figure 1. Overview of A3I, a conceptual framework and workbook for structuring the early ideation stage of designing AI-infused information-intensive applications [2]. Complementary to guidelines and pattern libraries that help teams evaluate or instantiate existing design directions, A3I prompts designers to explore what AI can do to help a knowledge worker process information—from sense-making to decision-making—and when and how to invoke AI, including when to reveal findings from AI automation and how mixed-initiative AI can act reactively or proactively.

3. A3I Part I: What Can an AI System Do to Help a Human Process Information?

First, to construct an AI-infused workflow, we should define AI's role in the two phases of a human's processing of information: the sense-making and decision-making phases.

3.1. The Sense-Making Phase in AI-Infused Information Processing

While most AI models are intended to help humans reach a decision given some information, we argue that an AI-infused workflow should start from supporting humans to make sense of the information.

3.1.1. Theoretical Foundation

Sense-making emerged as an important concept in early research on information foraging [13]. Russell et al.defined sense-making as "the process of searching for a representation and encoding data in that representation to answer task-specific questions" [6]. We consider sense-making a preparatory step to get information in the right form that lends itself well to subsequent decision-making. Sense-making is also one of the main objectives in designing information visualization, as Shneiderman concisely summarized in his mantra: "Overview first, zoom and filter, then details-on-demand" [14], which inspired the development of seminal techniques, such as Tree-Maps [15] and zoomable user interfaces [16].

Two common design "ingredients" emerge from sense-making research: hierarchy and filter. As pointed out by Simon [17] and Meadows [18], it is unsurprising that most complex systems are structured hierarchically because hierarchy is the ubiquitous approach of dealing with complexity in the natural world, e.g., cells in biology and elements in physics. To support sense-making, it is useful to present a hierarchic view of the information while allowing users to filter in a subset of information based on specific criteria.

The importance of sense-making was debated [19] in the late 90's as a school of work on interface agents aimed to automate information processing tasks while dispensing with human efforts, including performing the sense-making phase. However, almost three decades later, many AI models still cannot autonomously make decision entirely without humans, especially in high-stake scenarios such as medical diagnoses. Sensing-making provides humans with an opportunity to perform a "sanity check" of the information that might lead to detection of abnormal AI behaviors.

3.1.2. Design Examples

When involving AI in the sense-making phase, we found systems that employed two approaches of using hierarchy. First, when involving AI to help a human deal with a large number of information items (or "chunks" [1]), prior work employed a hierarchic view to present such items. For example, to support the task of finding relevant papers, Apolo uses AI to help a user to organize a large network of papers into a hierarchy, i.e., selecting and grouping papers by multiple levels of topics [20]. Analogously, Xiao et al.'s system allows a user to group similar photos while AI predicts which photo can best represent each group [21]. Second, AI can help a human make sense of a single piece information of complex nature, (e.g., a high-resolution medical image with miniscule features). Such information is often presented in a very large data view [22] and a system can employ a hierarchic structure to help the user navigate within that information item. For example, for BVI users where even processing a normal picture becomes challenging, ImageExplorer creates a hierarchy of AI-recognized objects contained in an image, thus guiding a BVI user to navigate and explore these objects [23]. In the medical domain, both CheXplain [24] and xPath [25] employ a hierarchic UI layout that allows physicians to navigate AI diagnoses based on high-resolution medical images across multiple levels of details.

Even with a hierarchic view, users still might find it overwhelming to make sense of all the information presented; thus a natural next-step is employing AI to filter in a smaller subset of information for the user to focus on. Some prior work provides explicit filtering controls, e.g., FashionQ allows a user to filter attributes used by AI to compare and find similar examples [26] and Podium filters in important attributes inferred from a user's manual ranking of data rows [27]. Other systems employ filters implicitly (similar to recommender systems). For example, when interacting with generative AI, as a user sketches a concept [28], directly manipulates a design [29] or specifies requirements or constraints [30], AI presents the best-matching results while implicitly filtering out (a large number of) others that do not meet the user's intents. Some systems implicitly infer user-desired criteria for filtering, such as AI-enabled fact-checking that filters news sources by relevance to the search queries [31] and filtering AI-generated design recommendations based on learning a user's goal from their manipulation of parameter-controlling sliders [32].

3.2. The Decision-Making Phase in AI-Infused Information Processing

In processing large amounts of information, one often needs to make a series of decisions to reduce and transform the given information into choices (e.g., whether the patient has a benign or malignant tumor) or insights (e.g., what specifically one can learn from reading an article). A3I treats predictive and generative as types of user-facing support rather than mutually exclusive types of AI systems. A single model or application may provide both across different workflow states.

3.2.1. Theoretical Foundations

Using AI prediction to assist decision making can be traced back to Simon's "bounded rationality" theory [33], that humans have limited cognitive capacities and have to settle with less than optimal decisions ("satisficing"). AI, since its early days, have held the promises of overcoming the human limits in decision-making by computing larger amounts of information faster [34], such as expert systems that assisted diagnostic decision-making by accessing a large database of symptoms and evaluated the patient's case by rules [35]. However, early research also realized AI's limited usefulness if all it could do was making a choice amongst a number of plausible diagnoses. Specifically, Blois visualized the diagnostic process as a funnel that starts with a wide opening of gathering facts and proposing hypotheses, followed by a narrowing-down process of gradually converging to a final diagnosis [36]. Blois argued that AI at the time was limited because it was only good at the narrow end of the funnel. Therefore, to better assist the holistic decision-making process, AI needs to go beyond presenting humans a prediction (convergent thinking [37]); rather, AI can also help a human by retrieving or generating information relevant and useful to decision-making (divergent thinking [37]).

We use predictive support for interactions that help users narrow toward a judgment, classification, or ranking, and generative support for interactions that produce, retrieve, or synthesize information to expand or transform the available alternatives. These interaction-level roles are broader than a model-family distinction: the same foundation model can rank candidate information in one step and generate an explanation in another.

3.2.2. Design Examples

To assist decision-making, AI prediction has become a well-established approach in human-AI collaboration research, from deciding the diagnosis of a medical image [38] to deciding the accessibility of a video segment [39] and to deciding the thematic label of qualitative data [40]. Meanwhile, building off of recent advances, generative AI has started to play a useful role in assisting humans' decision-making process.

The primary direction is employing AI to generate "meta information" that can help a user's decision-making based on the original information, such as generating questions to promote interactive parent-children storytelling [41], generating analogy between papers to help users understand broader literature [42], generating tagclouds to help designers iterate and communicate their work via a digital mood board [43], and generating question-answer pairs to help a reader decide what to focus before reading an article [44].

Further, prior work also employed AI to retrieve additional information to inform decision-making, such as helping fashion designers determine a subset of relevant examples from a large database [26], presenting comparable AI diagnoses to help physicians contextualize a model's analysis of an X-ray image [24], and enhancing Content-Based Image Retrieval to let pathologists find histopathological images relevant to a patient's case [45].

Predictive and generative support can therefore be combined, e.g., generating question–answer pairs about product reviews and ranking which pairs are most relevant to a purchase decision. Recent agentic systems make this combination especially visible: one agent can retrieve memories, synthesize reflections, form plans, and initiate actions, moving among generative and predictive support as well as reactive and proactive behavior within one workflow [46]. A3I focuses on these user-facing roles and invocation choices rather than assigning an entire contemporary system to one AI category.

4. A3I Part II: When and How to Invoke AI as a Human Processes Information?

After exploring the role of AI through the sense-making and decision-making phases, a parallel question to consider is when and how to "insert" AI into the human's workflow. Generally speaking, this can be thought of as a turn-taking question: when should human or AI take over the information processing task from the other and how to initiate such turn-taking? We first consider a simpler case (AI Automation) where AI is intended to automate most of the information processing task (e.g., AI making a diagnosis given a medical image) and there would be very little turn-taking. Then, we consider the other cases (Mixed Initiative) where it would require a fair amount of turn-taking between AI and human to accomplish the information processing task (e.g., AI guides a physician to navigate and examine a high-resolution medical image).

4.1. AI Automation: Revealing AI Findings While Upfront vs.in Due Course

In the case where AI is expected to automate most of the information processing task, the design goal is to reveal AI's findings while preventing the human from overrelying [9] on AI automation (i.e., accepting erroneous results with a lack of their own judgment).

4.1.1. Theoretical Foundations

While it might seem straightforward to present AI findings upfront to the human. various past work has shown that this approach will likely lead to overreliance on AI [8,9,47], which is defined as "users accepting incorrect AI recommendations" [9]. To prevent overreliance on AI automation, Passi and Vorvoreanu reviewed a comprehensive set of mitigation techniques during initial interactions with AI (e.g., educating users during onboarding) to regular use (e.g., providing real-time feedback). Amongst these overreliance mitigation techniques, the category most related to when/how to invoke AI is the use of cognitive forcing functions [8], e.g., making AI available on-demand or simply adding a wait time before showing AI's findings. Another solution is to offer "simple explainable AI", e.g., explanations and uncertainty, which, however, did not prevent overreliance as effectively as cognitive forcing functions. Prior work also noted that preventing overreliance comes at a cost of increasing the user's cognitive load and time consumption for completing a task [8].

We lay out two design choices when presenting AI's automation of humans' work: (i) showing AI findings upfront before humans access information with a risk of overreliance; and (ii) showing AI findings in due course, including making them available on demand or with a wait time, which mitigates overreliance but adds cognitive load and processing time.

4.1.2. Design Examples

Many AI-infused information processing systems tend to reveal AI findings upfront, such as automatically identifying salient features of assessment to perform patient-specific rehabilitation analysis for therapists [48] and automatically recognizing and grouping objects in an image for BVI users to explore [23]. Yang et al.proposed and studied an interesting idea that embeds AI-generated prognostic information "unremarkably" into physicians' slides used in the decision making process, which solves the problem that physicians are "unlikely to encounter or to actively engage with" a decision support tool [38]. This technique addresses the opposite problem of underreliance where a human disregard or reject correct AI findings. The Impetus system [49] is an example that integrates both "upfront" and "in due course" designs when involving AI in pathologists' diagnostic process. The system would present AI's prediction up-front if the model shows a high confidence whereas, with a medium confidence, Impetus would only show AI's findings upon the pathologist fills out a diagnosis form only available after a manual examination.

4.2. Mixed Initiative: Designing AI to Take Turns Upon Human Requests or on Its Own

When AI is not yet capable of automating a task, the design should invoke AI mixed-initiatively—that is, combining both AI's and human's contributions to accomplish the task.

4.2.1. Theoretical Foundations

In Horvitz's original vision of mixed-initiative interaction [7], AI would engage itself with humans both reactively (e.g., waiting for user confirmation of a tentative schedule) and proactively (e.g., prompting the user with a potential need to schedule a meeting based on the reading of an email). The vision of proactive AI dated back to the debate between direct manipulation and interface agents [19] where the latter envisions a future of "building machinery for sensing a user's activity and taking automated actions". Horvitz argued that the use of agents might suffer from a lack of considering "costs and benefits of automated action" and "poor timing of action". Horvitz

then demonstrated how to analyze the utility (i.e., benefits) of taking an automated action (or not) given correctly inferring a user's goal (or not), yielding a a threshold of probability AI can use to determine whether it should proactively take certain actions ([7], Table 1 and Figure 4). It is also possible to go beyond a binary "action or not" but consider reactive-proactive AI behavior as a continuum. Both Horvitz's original paper and its later extension into guidelines [4] advocate for "scoping"—varying and degrading the precision of AI's actions based on the cost, benefit, and uncertainty analyses.

Building on this literature, A3I lays out three "key points" on this reactive-proactive continuum to illustrate where AI can "chime in" as a user processes information: (i) Only acting upon a user request (reactive); (ii) Offering to take action pending user confirmation/correction (a mix of reactive and proactive); and (iii) Taking an action without a user request (proactive) while allowing for undoing or modifying the action.

4.2.2. Design Examples

Most systems to-date tend to invoke AI that only reacts to users' explicit commands or inputs. One key difference in proactive AI is the ability to anticipate a user's needs without the user realizing or expressing such needs. In other words, the user might not expect AI to act under certain circumstances yet AI would appropriately take the initiative to get involved and contribute to the task.

There have been only a few examples in A3I applications where AI acts both reactively and proactively. For instance, in Datatone, as a user queries a database, AI can proactively detect potential ambiguity, provide appropriate widgets for the user to clarify the query, and then reactively show resultant data points [50]. AI can also proactively act in the peripheral (similar to the notion of Calm Technology [51]), such as how ImageSense generates tagclouds proactively while the user interacts with images on the digital mood board [43]. As another example, NaviPath is a tool that involves AI throughout a pathologist's navigation of a high-resolution histopathological image [52]: a pathologist can toggle whether AI would proactively recommend regions of interest as they enter a new area of the image. Visual Captions provides users three levels of AI proactivity when retrieving images to illustrate a video chat: Auto-display (AI autonomously adds images), Auto-suggest (AI proactively suggests images), and On-demand-suggest (AI suggests images when prompted by the user) [53].

5. Deployment in Student Course Projects

Across the three studies, our objective was to conduct a formative, multi-context assessment of whether the A3I framework and its workbook operationalization were understandable, applicable, and useful as an early-stage thinking aid for designing an AI-infused information-intensive application. We started with the first study in an academic environment before transitioning subsequent studies to an industrial setting. This evaluation design matches the maturity of the artifact: before comparing A3I against other ideation methods, we first needed to understand how people interpret the framework, where the workbook creates friction, and what revisions are needed for practical use. Accordingly, the studies do not establish that A3I outperforms unstructured brainstorming, other frameworks, generic workbooks, or no-workbook conditions. Specifically, we aimed to investigate the following formative evaluation questions:

- Does the workbook provide clear and understandable instructions for designing an A3I application?
- How do participants use the workbook in their design iterations?
- Do the participants find the workbook useful in helping them with the design tasks?

5.1. Translating the Framework into a Workbook

We went through a few iterations to translate the conceptual framework to an instructional workbook. Our initial version simply consisted of two set of bullet points corresponding to the two high-level questions in parallel as well as the individual "blocks" under each question. However, this version appeared more like a checklist than a step-by-step guide, which led us to a change into a "single-thread" structure. The new version follows the sense-making → decision-making thread: first designers explore what AI can do in both of these phases and then the workbook prompts them think about how and when to invoke AI. Specifically, the workbook consists of five parts—Instructions, System Overview, Sense-making/Hierarchy, Sense-making/Filter, and Decision-making, each of which prompts a designer to answer multiple open-ended questions with design artifacts (e.g., low-fidelity sketches). In response to later feedback, we now propose adding a preliminary "Data and Feasibility Check" between System Overview and the main framework prompts. This step would ask designers to state what information AI requires, where it comes from, whether it is available, and what privacy, quality, bias, fairness, or governance constraints may shape the design. We position this as a lightweight assumption-surfacing step rather than a gate that rules out ideas too early. Note that we divide the sense-making phase into two sections so that each section has about equal lengths.

To help designers understand the relationship between sense-making and decision-making, the workbook introduces the term "information chunk": sense-making involves hierarchically rearranging and filtering information chunks while decision-making is instead focused on processing individual information chunks.

Overall, we used plainer language to explain each ideation step and provided references to examples along the way to assist designers' connecting concepts to actual design ideas. Further, we added an introduction on the first page to quickly help designers understand the gist of this workbook and to determine its applicability to their projects.

5.2. Participants

The study was IRB-approved and conducted in the research team's institute in a graduate-level HCI course over the period of 10 weeks. All 25 students opted-in to participate in the study. The students formed nine project teams (T1-T9): one consisted of a solo member, one had two members, one had four members, and the remaining six had three members each. These participants were aged 21 to 35; 22 were male and three female; all majored in Electrical Engineering and Computer Science. When asked how they described their level of knowledge (no knowledge, passing knowledge, knowledgeable, or expert) in HCI prior to taking the class, eight students described themselves as knowledgeable, 15 had passing knowledge, and two had no knowledge. When asked about the same question about AI, four students described themselves as experts, 10 were knowledgeable, and 11 had passing knowledge.

5.3. Tasks & Procedure

We taught the A3I framework after three weeks of lectures on fundamental HCI concepts and methods. We then introduced student participants to the A3I WORKBOOK as part of their course project kick-off. Participants then spent the remaining seven weeks on a course project given the following description: "Create an AI Information Assistant to help the user process complex information, for example, reading long text, recalling memorable moments from your phone's photos, manage your overwhelming projects and assignments, or helping you cue the best music to boost productivity".

The project's structure consists of four iterations from low, to medium, to high-fidelity, and to the final proof-of-concept functional prototype. Participants were instructed to use the workbook (provided as a Google Doc) as a template to iterate their initial design (in low-fidelity sketches) towards the second iteration (i.e., to create a medium-fidelity prototype). Figure 2 shows excerpts from T6's workbook. The final prototypes of three teams illustrate the range of resulting designs: one used AI to guide a user through a recipe while cooking, filtering its key steps and presenting them at multiple levels of detail upon verbal request; one let a student describe a high-level learning objective, from which AI assembled a selection of courses; and one helped a programmer understand the scikit-learn library through a hierarchic view of its code files and AI-generated explanations of individual functions.

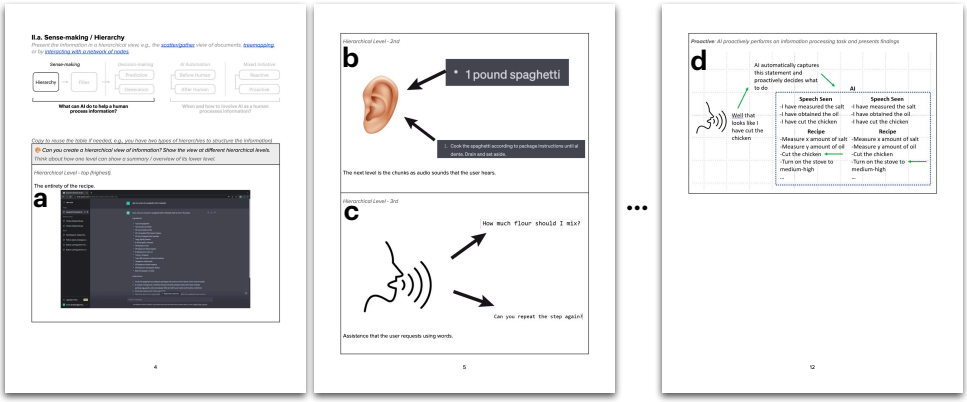


Figure 2. Excerpts from the workbook by T6 who were designing a AI-infused system that enables a user to follow a recipe while cooking. T6 identified a hierarchic way of presenting recipe information: from a top-level overview (**a**), to a second-level break-down of each steps (**b**), and to a third level where a user can ask for more details within each step (**c**). Further, T6 came up with a proactive mixed-initiative design (**d**) where AI infers what information a user needs based on their speech. For example, as a user says "Well that looks like I have cut the chicken", AI identifies the next step from the recipe and proactively informs the user.

The first author taught and graded the course, recruited and consented participants, and conducted the team interviews. Participation and interview responses did not affect grades; each team's project was graded by students from other teams. All nine teams received the same version of the workbook, which was not changed during this deployment. We nevertheless recognize that the instructor–researcher relationship may have shaped participation or favorable responses, as discussed under Evaluation Limitations.

We conducted an in-person interview with each team after their use of the workbook. During the interview, we walked through the workbook and asked them whether each step was clear and understandable, whether they had difficulty interpreting and responding to each step's prompts, and what design changes they made, if any, after following the workbook.

5.4. Analysis

We conducted a thematic analysis [54] of the qualitative data collected from interviewing the participants, performed by two members of the research team in two sessions within one week. In the first session, we reviewed the data bottom-up and summarized two levels of themes; in the second session, we reviewed each theme, rephrased a few for better clarity, identified and resorted ill-fitting data points, and merged some pairs of themes when necessary.

5.5. Findings

5.5.1. How Project Teams Used the Workbook

We found that each project team (T) often experienced multiple stages of using the framework for their project, from understanding the fundamental concepts, to interpreting the framework in the specific context of a project, to adapting or extending the framework for a project's workflow.

Understanding the Fundamental Concepts in the Framework

There was some initial confusion about the framework stemmed from the wording of "AI Automation" vs."Mixed Initiative". In response, we revised the wording in the workbook, replacing "AI Automation" with "If AI automates most of the task" and "Mixed Initiative" with "If human and AI need to work together" (We still use the "AI Automation" and "Mixed Initiative" in this paper for brevity purposes.). Further, we added examples in the workbook to show that mixed-initiative systems have to have human's input in order to complete a task.

Another interesting question raised by T7 was how to decide between "AI Automation" and "Mixed Initiative"—should it depend on how good the AI model is (able to automate the tasks) or whether we would like to involve humans in the loop? To address this, we clarified in the workbook that, as a general rule of thumb, if the AI model is unable to automate a task, the design should enable mixed-initiative interaction between human and AI; otherwise, both options are possible and worth exploration.

Based on teams' feedback, we further clarified a few other concepts in the revised workbook.

- "Decision-making" can be an integral part of a larger task rather than the entirety of a task (T4) and should be made by human and AI together rather than either alone (T7).
- "Filter" made the most sense when explained as "rules that reduce the information a user has to process" (T5); however, reducing (less relevant) information does not contradict the next step of using generative AI, which means providing user with more relevant information (T7).
- "Hierarchy" is not a prerequisite of "filter", which can happen to information presented at any hierarchical view or in a non-hierarchial way (T4).
- "Information chunk" is meant to be processed by the user, not AI (T1). This is an important distinction because AI can process much lower-level information chunks than humans. Thus A3I UI should present information chunks that can be processed by both human and AI.

Interpreting the Framework in the Specific Context of a Project

Building on their understanding of the basic concepts, project teams often went through their process of interpreting such concepts in the specific contexts of their projects.

For "hierarchy", only one project team (T6) reported to have difficulties interpreting this concept as they were designing a primarily audio-based interface of AI instructing a user to follow a recipe while cooking. After discussions, they considered a secondary visual display of a hierarchic view of the recipe's steps as the "back-up" when the user has trouble understanding the audio instructions.

"Filter" seems to be the most universally interpretable concept across various kinds of projects. One project team (T5) working on generative design further interpreted filtering as processing a user's natural language input to keep only the tokens most relevant to the design attributes of the generated furniture.

Some project teams struggled to interpret "decision-making" in their projects, such as finding cocktail recipes (T1) and following cooking instructions (T6). Through discussions, both teams eventually were able to define what decision-making means for their users, e.g., deciding which cocktail they like best (T1) and deciding what to do next in cooking process (T6).

Adapting or Extending the Framework for the Project's Workflow

Once project teams established their interpretation of the framework, some further adapted the framework to better fit their needs.

Some teams challenged the framework's assumed workflow. For example, T2 asked whether it was possible for decision-making to happen before filtering, e.g., filtering AI-generated findings. This suggestion led to our revision of the framework to add optional arrows directing back from the decision-making to the sense-making stage. T5 (working on AI-generated furniture design) wanted to change AI's mixed-initiative behavior over time as AI initially would need more user input (human-initiative) and later, with more knowledge about the user, can transition to more AI-initiative.

Teams also suggested changes to the current framework to better support their design activities. One common suggestion is to encourage more high-level thinking, e.g., letting the project team speculate early on how AI can help the user (T4), prompting the project team to think about how to invoke AI from the beginning (T4), and providing more examples to overview the flow of information in a hierarchical view (T7). Further, T3 suggested adding cognitive forcing function [8] in the framework while T5 suggested specialized versions of the workbook tailored to different categories of applications.

5.5.2. How Project Teams Perceived the Workbook's Usefulness

Overall, project teams described the workbook as mostly applicable to their projects and perceived the process as constructive for their design iterations in several ways.

Project Teams Used the Workbook to Help Them Think Broader

One team felt the workbook forced them to "look at things we didn't consider ... likely to miss" (T4). Some thought of the process of going through the workbook augmented the user-centered design approach—"gaining different perspectives" (T9), interacting with the system like a user rather than a system developer. T3 initially struggled with applying the sense-making part to their project that employed AI to help fashion designers find existing examples or generate new examples. Via this thought process, they found it "interesting about the sensing-making and hierarchy part because that—[even though it is] outside the scope of this project–would be a nice logical extension" (T3).

Teams also found the workbook helpful for generating broader ideas, useful "at the brainstorming process; got a lot of ideas from going through the process" (T2) and allowing them to "work together and show each other's ideas using this workbook as the collaborative medium" (T2). T6 initially tried to frame their system to fit the workbook but soon realized that they should use it to brainstorm new ideas beyond their current design.

Project Teams Used the Workbook to Help Them Think Deeper

The workbook probed project teams to think of more details. T4 used the workbook to think about when to use cognitive forcing functions in the workflow. Both T6 and T8 used the workbook to break down their system workflow into multiple components, which allowed them to think about how the system "operates from different perspectives" (T6) and was "good for low-level high-granularity thinking" (T8). T8 also pointed out how the workbook's question about filtering at different hierarchic levels prompted them to think about more nuanced designs rather than a one-size-fits-all filtering technique. T7 considered different parts of the workbook address separate issues: the sense-making part helped them iterate the UI design while the AI automation vs.mixed initiative part prompted them to think about "AI design" (which AI to use and when).

Project Teams Used the Workbook to Help Them Think More Structurally

Teams found the workbook useful to help unify and organize their design ideas, e.g., "talking through ideas into a cohesive plan" (T2) and "categorizing what we already designed" (T4). As T3 put it: "not sure it changes the system design but helps us organize the parts of the system and what it is doing" (T3).

Some teams used the workbook to clarify and walk through their system design, e.g., "clarifying the tasks the system is performing and how" (T3) and "to think about the overall procedure especially for the filter and mixed-initiative parts" (T7). Some felt the workbook "provides a vocabulary to articulate design thinking and decision of involving AI" (T9).

Project Teams Used the Workbook to Help Them "Think Again"

"Think again", as articulated in Adam Grant's book [55], refers to the act of challenging one's established belief and being open to changing one's mind based on new evidence or perspectives. Two teams mentioned "think again" moments while going through the workbook.

T4 was designing a system where AI helps a user understand the dependencies of a function in a large codebase. Their focus was on creating a hierarchic view based on which to visualize dependencies as a connected graph; thus they initially did not consider using plain text to explain dependencies because this approach would be "too trivial". Going through the decision-making step of the workbook, they started to rethink how adding plain text explanation as a generative approach might also be useful at this step—as long as it could strengthen the user's understanding.

T7 was creating an AI-assisted writing tool where AI reactively provides suggested text or relevant references in response to a user's writing command. The workbook prompted them to think again whether such AI behavior can be proactive, e.g., showing suggestions without a user's request. At the end, they stood by a reactive approach to let the user stay in control.

5.6. Summary

Our findings encompass how course project teams interpreted the high-level concepts and used the A3I WORKBOOK, including whether and how they perceived the workbook as supporting their design process. We found that each project team often experienced multiple stages of employing the workbook for their project, from understanding the fundamental concepts, to interpreting these concepts in the specific context of a project, and to adapting or extending the workbook for a project's workflow. Meanwhile, interview accounts and workbook walkthroughs suggested that teams used the workbook to broaden, deepen, structure, and reconsider their ideas [55].

Beyond the usability of this particular workbook, the deployment study suggests two broader lessons for design-support tools. First, design scaffolds are most useful when they expose dimensions designers had not explicitly considered; for example, T6 expanded from a primarily reactive cooking assistant toward a proactive assistant after exploring initiative alternatives, and T8 generated filtering strategies at different hierarchical levels rather than relying on one filter for the whole interface. Second, conceptual completeness creates workload and interpretation costs. The teams' questions about automation, mixed initiative, filtering, and decision-making show that a theoretically complete framework still needs examples, lighter-weight entry points, and language that helps designers map abstract concepts to their current artifact.

5.7. Summary of Framework Translation and Workbook Revisions

Taken together, the deployment study shows how A3I's theoretically derived components were translated into workbook prompts, where project teams encountered friction, and how those findings informed revisions. Table 3 summarizes this connection across the major framework components.

Table 3. How A3I connects theoretical sources, framework components, workbook prompts, observed difficulties in the first deployment, and revisions accordingly.

Theory	Framework Component	Workbook Prompt	Observed Difficulty	Revision
Sense-making and information visualization	Hierarchic information view	Create a view that organizes information chunks across levels of detail.	Project teams needed examples for applying hierarchy across diverse interfaces.	Added examples and clarified hierarchy as an information organization problem rather than only a visual layout.
Information foraging and filtering	Filter	Specify criteria that reduce the information users need to inspect.	Some teams needed clarification that filtering can apply with or without hierarchy.	Clarified filtering as rules that reduce information to task-relevant subsets.
Bounded rationality and predictive/generative AI	Decision support	Explore predictive and generative ways AI can help users process a selected chunk.	Some teams needed to define what "decision" meant in their task.	Added examples emphasizing decisions as parts of larger workflows.
Overreliance and cognitive forcing	Upfront vs. in-due-course automation	Decide when AI findings should be revealed.	Teams asked how automation differs from mixed initiative.	Reworded around whether AI automates most of the task.
Mixed-initiative interaction	Reactive-proactive continuum	Decide whether AI should act on request, suggest action, or act on its own.	Teams wanted more concrete examples for exploring AI initiative.	Added examples and identified pattern-library linkage as a future direction.

6. Expert Interview

To help our studies transition from a classroom to an industrial setting, we conducted an expert interview (IRB-approved) with three experts (E1-E3) with industrial research experience, seeking their opinions to evaluate the usefulness of the A3I WORKBOOK and to strategize for its deployment to support professional UX designers' work.

6.1. Participants

Two of the experts currently worked in a software company (one as a researcher and the other in UX research) and the third expert was in academia but used to be a member of an industrial research lab. The experts had between five and nine years of experience working on research topics ranging from creating responsible, user-centered AI, and human-AI interaction and collaboration, resulting in over 12 related publications in CHI, CSCW, and TOCHI. All experts had experiences collaborating with and/or studying professional designers working on AI-infused systems.

6.2. Process

Each interview took place via a video conference between a research team member as the interviewer and one expert as the interviewee. We started with an introduction of the overall A3I framework, followed by some question-answering for clarification. Then we asked each expert to walkthrough and review the workbook via a shared Google Doc. Finally, we probed for the overall usefulness of the workbook in guiding the design of AI-infused information-intensive applications as well as specific aspects that are challenging and that would need further improvement. Each session of the interview was recorded and took between 45 minutes to an hour. Each expert received $50 Amazon gift card for their participation.

6.3. Analysis & Findings

We took a thematic approach [54] to analyze qualitative data transcribed from the interviews, building an affinity diagram that organizes excerpts of the transcripts to hierarchically build up higher-level themes shared amongst experts, which we report below.

6.3.1. Helping Designers Decide Whether to Use the Workbook Based on Its Relevance to Their Projects

Foremost, experts pointed out the importance of making it clear upfront who should use the workbook, why, and when, so that designers will be motivated to invest their time and efforts to follow the steps. In addition, E3 suggested using an example to show "what is the exact outcome" of using the workbook. All experts also mentioned the importance of scope: it is important to help designers decide whether their project fits within the scope of the workbook. While the workbook does cover a broad range of information processing tasks, designers might hesitate to decide whether their project fits in this scope.

6.3.2. Using Designer-Friendly Language and Examples

E1 pointed out that the workbook's language should be "designer-friendly", i.e.,, removing or replacing academic terms with lay descriptions, such as "information chunk", "sense-making", and "mixed-initiative". Further, E1 suggested the use of "call-to-action" language. For example, instead of saying "Can you design a hierarchic information view ...", use the "do this" phrasing—"Come up with a design of a hierarchic information view ...". In particular, the headings that describe the major steps should also use the same style of language. Instead of "Sensing-making / hierarchy", which contains no call-to-action, rephrase the heading as "Step 2: Create a Hierarchic Information View".

E1 also suggested using online shopping—a scenario quite familiar to designers—as a cross-cutting example throughout the workbook. E2 mentioned the need to differentiate what the workbook envisions (i.e., involving AI throughout a task) vs.a conventional online shopping UI design, so that designers can think creatively about how AI can enable new user experiences beyond existing online shopping examples.

6.3.3. Integration with Designers' Existing Workflow and Tools

Experts questioned the relationship between the workbook and other design activities to better understand its role contextually. E2 pointed out that there should be a "Question 0"—how people currently process information without using AI, which needs to be answered as the prerequisite of the two main questions in the framework. E1 compared this workbook with the human-AI experience (HAX) toolkit (https://www.microsoft.com/en-us/haxtoolkit/ (accessed on 19 August 2026)) and pointed out that they are complementary with each other with the following differences: (i) Our workbook serves an even earlier design stage than HAX; (ii) Our workbook is better for "initial

design ideation" while HAX is for "planning the interaction details"; and (iii) Our workbook has a more specific focus on information systems while HAX applies to more general cases of involving AI.

All experts suggested that we provide a version of the workbook with tighter integration with tools that designers frequently use. E2 pointed out that, without such integration, going through the workbook would seem more like extra work. For example, one possibility is integration with Figma, e.g., showing the questions in a FigJam board or providing a project template where designers who already use Figma can directly do prototyping in response to each question.

6.3.4. Benefits and Cost of Using the Workbook

Overall, experts saw clear benefits of the workbook for the ideation stage when designing an A3I application: E1 liked the workbook's idea of "holding hands", i.e., guiding designers step-by-step through the series of questions that prompt their design ideation. E3 considered the workbook "overall useful" during the "first diamond" [56]. Further, E3 appreciated how the workbook's mapping from AI capabilities to information processing serves to sensitize [57] designers' abilities to comprehend and define AI's role in their project. Further, both E2 and E3 mentioned the benefit of the workbook to involve other team members, e.g., data scientists who can help designers identify AI models that work with the type of information the system deals with (E3). E2 felt this workbook provides a vocabulary for designer to think about, present, and advocate for their ideas.

In terms of cost, E1 and E2 pointed out the effort of completing the workbook is "demanding" (E2) and might discourage designers from going through the process. E2 suggested ways to reduce the workload, such as combining certain steps and, rather than responding to each question with sketches, allowing designers to create one single sketch and then use the workbook's questions to reflect on and iterate their design.

6.4. Summary

The expert interviews suggest three broader lessons for translating a conceptual framework into a professional design scaffold. First, designers need an explicit relevance check before investing effort: the workbook should help teams decide whether their project is information-intensive, whether they are early enough in the process for divergent exploration, and what output they should expect. Second, theoretical constructs must be translated into practitioner vocabulary and familiar artifacts, because terms such as "sense-making," "information chunk," and "mixed initiative" can obscure otherwise useful distinctions. Third, design scaffolds must balance conceptual completeness against perceived workload by supporting both fuller walkthroughs and lighter reflective use around an existing sketch or prototype.

7. Design Exercise with UX Professionals

Based on findings of the above two studies, we iterated on the A3I WORKBOOK and then conducted a week-long design exercise with professional UX designers to understand how they would use the A3I WORKBOOK at an early stage of designing an A3I application.

7.1. Changes Made to Improve the A3I WORKBOOK

The expert interviews surfaced a set of issues different from the first deployment study, such as the needs to help designers decide whether the workbook is relevant to their projects, to use non-academic and designer-friendly language, and to reduce workload. Based on these feedbacks, we made the following changes to improve the workbook:

- We added an additional "Before You Start" page, briefly answering three questions to help a designer determine the workbook's relevance to their projects: What is this workbook for? Who should use this workbook? When should this workbook be used?
- We conducted a pass of "jargon-proofing", replacing academic terms with lay descriptions and further used a single online shopping use case (replaced by a photo application in this study, because the study's design task itself involved online shopping) throughout the workbook to help designers understand key concepts behind each questions.
- We created A3I WORKBOOK-lite—A lighter version of the workbook that condensed the original long document into a few slides (Figure 3) to reduce the barrier and workload of adoption. The slides also naturally compartmentalize the steps, allowing a designer to focus on one slide (step) at a time. We also used different colors for instructional slides and workspace slides, the latter of which serve as a canvas for designers to create or attach their prototypes in response to each step.
- We called out a "branching point" prior to "When/how to involve AI ...", making it more clear that the designer

should decide either AI will automate most of the task or work mixed-initiatively with humans.

- Throughout, we added visual in addition to textual examples to illustrate each step, e.g., diagrams of file systems, a treemap visualization, and a screenshot of a photo organizing app to illustrate the concept of hierarchic information view.



Figure 3. Our design exercise with UX professionals used the A3I WORKBOOK-lite—a lighter version of the workbook that condensed the original long document into a few slides to reduce the barrier and workload of adoption. Shown are 8 of its 16 slides.

7.2. Participants

We recruited 7 participants, 2 male and 5 female, aged 25 to 55. All participants worked as UX designers in a North America software company, with professional experiences ranging from 3 to 29 years. In a pre-screening survey, all participants mentioned having experiences designing AI features in various applications, ranging from virtual reality to cloud computing platforms. We recruited these participants through the company's internal communication channel. This study was approved by the IRB of the author's institute and followed the company's policy on engaging in research with human subjects. Upon completion of the study, participants received an internal incentive—a choice of company merchandise—for their time.

7.3. Tasks & Procedure

We met with each participant over a video conference to start off the design exercise. First, we described to each participant the task as follows:

> *Users of online shopping platforms (e.g., Amazon) often browse other customers' reviews before buying a product; however, the shear number of reviews makes it challenging for a user to process all the information. You are part of an online shopping team that wants to create AI features on the platform to help a user better decide whether to purchase a product based on other customers' reviews. The deliverable is a low- or medium- fidelity prototype.*

Having established the design task, we then presented each participant a step-by-step walkthrough of the A3I WORKBOOK-lite. We asked each participant to use the workbook as an ideation and exploration tool as much as they feel it is useful and appropriate for the design exercise. They could also use any other tool they preferred to assist with the task, e.g., Figma. Specifically, we asked each participant to fill in a copy of the workbook (in the form of Google Slides) as they performed the design exercise, including using the commenting features to communicate any questions and feedback. We asked participants to spend a total of two hours of time over the course of one week before we scheduled them for an exit interview, during which we walked through each participant's workbook to follow up on the questions or feedback they raised as comments. Finally, we asked participants to summarize their overall experience of using the workbook for the design exercise, answering the following questions that mirrored the ones we asked in the first study:

- Did you have any difficulty understanding/following the workbook?
- Is this workbook useful for the design task? Why or why not?
- Do you see yourself using this workbook in your work? Why or why not?

7.4. Analysis & Findings

The qualitative data included participants' completed slide decks, comments within the slide deck, and their responses during the interview. The research team conducted a thematic analysis [54]. First, we broke participants' comments and responses into atomic segments, each of which was assigned a code. Next, we developed categories from segments that shared the same code. Finally, we merged related categories into higher-level recurring themes.

Overall, participants felt the workbook was "clear" (P3) and "straightforward" (P5). Participants appreciated the examples being "closely related to the designer's task" (P4), that helped them to "learn and remember" (P1), to understand "what the questions mean" (P3), and to direct their thinking (P2, P3).

7.4.1. Participants Perceived That A3I WORKBOOK Supported a Structured thought Process

One of the two major perceived benefits of using the workbook is how it supports a step-by-step, structured thought process, which "prompted me to think about things in a certain order" (P7), "reminding designers to think about what information users need and how to help them process the information" (P1) and "trying to let me think about design in a more structured way, building into each step" (P5).

> *The workbook forces me not to get ahead of yourself, e.g., start with thinking of the info architecture. Usually when designing in Figma, I do not get a chance to stop and think about details. It's good for the workbook to call the things out.* (P3)

Participants pointed out how the steps about when and how to involve AI were "really helpful ... giving a very good framework for helping the brainstorming activities" (P6). Participants also mentioned how such a structured thought process could facilitate communication (P1), following the same steps to present one's design ideas (P3). In the meantime, P5 also pointed out a potential downside of limiting oneself to the workbook's structure, "having a tendency to not explore beyond what the book says" (P5).

7.4.2. Participants Perceived That A3I WORKBOOK Surfaced Implicit Design Dimensions

The other major perceived-benefit of the workbook is how it helps discovering unknown or implicit design dimensions "to think about AI" (P5) and to "brainstorm from different aspects that I might not think about" (P2). Specifically, participants called out the steps about when and how to involve AI—

> *I didn't think of upfront vs. in due course design patterns. It's very useful yet usually happen in the design process at a subconscious level. I appreciate putting those out as individual and separate steps, calling out the timing dimension which is very important in a user journey.* (P5)

Relatedly, P7 commented on the part of AI being reactive, proactive or somewhere in between—"I mostly had the ideas before that, but I did not look at it through this lens" (P7).

7.4.3. Difficulties of Understanding and Responding to "Hierarchy"

Perhaps the most surprising finding was participants' struggle in understanding and responding to the step of creating a hierarchic information view. We had expected participants to come up with novel ways of organizing and visualizing reviews hierarchically (e.g., similar to what an affinity diagram would look like). However, all but two participants seemed to have interpreted "hierarchic information view" in unexpected ways, such as showing a list of review attributes, e.g., location, date, score (P4, P6, P7), a sequence of steps to retrieve a review (P2), or dividing all reviews by individual reviewer (P5), which was just a long list of reviews. Despite the examples (file system, tree-map, and digital photo organization), participants were unsure "what the ask is" (P7) in that step. Participants felt it did not quite fit the design process they were used to—"usually not starting with thinking in terms of info hierarchy" (P6). Some felt this step of "focusing on the information hierarchy may not have been very useful in this task" (P7). Overall, the findings suggest that "creating a hierarchic information view" is not a type of language designers have common understanding of, nor is it clear to them why this step is necessary in the flow of the workbook. One possible explanation is that the professional designers were not designing from scratch; they tended to imagine AI features within existing product ecosystems where information architecture may already be established. Another possibility is a terminology mismatch: what the framework calls a hierarchic information view may overlap with designers' familiar practices of information architecture, navigation, or dashboard organization, but the workbook did not make that connection explicit. The confusion may also reflect an ambiguity between an information hierarchy (how task-relevant information is conceptually organized) and an interface hierarchy (how screens, panels, and controls are visually arranged). This disconnect became a substantive finding rather than merely a usability problem: translating an academic concept into a design scaffold requires aligning the concept with practitioners' existing vocabulary and work constraints.

7.4.4. Trade-Offs of Using A3I WORKBOOK-Lite

As participants used the "lite" version of the workbook, they noticed the trade-offs of leaving out more detailed instructions. Some requested additional examples specific to the task (P7) or "success stories from actual products" (P6). P2 suggested breaking down "what AI can do" into two smaller steps focused on predictive and generative help, respectively; P3 suggested filling in the reactive-proactive spectrum with more concrete designs; P6 suggested the workbook to call out new and diverse capabilities of AI to further help designers imagine what AI can do. Since different designers might expect to see more details in different parts of the workbook, one possible revision could be linking each step in the lite version to the ones in the full version.

7.4.5. Various Ways of Adopting A3I WORKBOOK in Day-to-Day Work

Overall, despite coming from different product backgrounds, participants felt the A3I WORKBOOK is relevant to their work in the company, as "more and more designers' work will be AI-assisted features" (P6). The focus on information-intensive applications also resonated with a few designers' work products where users often needed to interact with a large amount of text, logs, charts, or documents (P2, P3, P6, P7). However, participants also mentioned practical constraints of using the workbook as-is. Foremost, all participants spent the majority of their work on established products rather than ideating new products or features at an early stage; even when creating new products or features, the design space was often constrained by the existing eco-system of products, e.g., a dashboard-like environment (P5). Considering such constraints, participants suggested various ways to adopt specific facets of the workbook. For example, P1 would gather comprehensive information about users to port into the workbook, P5 suggested using the workbook when a product already has an established UI and wants to add some AI features, and P7 considered using the workbook as "a very basic checklist just to make sure I've thought about different AI angles" (P7). Quite a few participants emphasized the use of the reactive-proactive spectrum, using the workbook to "break down the interaction flow to explore at what steps to surface AI actions reactively or proactively" (P3), or to enable "AI proactiveness when detecting capacity to save customer unnecessary cost" (P7). P6 was designing automatically-generated widgets in a console for investigating issues, which currently worked reactively; the workbook could help him *think through both scenarios*P6 and possibly expand the design to provide proactive support.

7.5. Summary and Comparisons with Findings in the Earlier Study

Our findings echo those (Section 5.5.2) from the previous study (deployment in course projects): (i) the A3I WORKBOOK supports a structured thought process (helping project teams think more structurally) and helps discovering unknown or implicit design dimensions (think broader). However, the unexpected finding is participants'

struggle in understanding or responding to the step of creating a hierarchic view—in contrast, only one project team in the earlier study reported difficulty with this concept. Participants in the design exercise found hierarchy to have little relevance to their design process or the study's task. Overall, while participants might not use the workbook in their work exactly as they did in this exercise, they identified relevance and various opportunities to employ specific facets of the workbook.

This comparison surfaces two additional lessons. First, designers reinterpret theoretical constructs through their existing professional vocabularies and product constraints. The same hierarchy prompt that students used while designing from scratch became difficult for professionals who imagined adding AI features to existing product ecosystems with established information architecture, navigation patterns, and dashboard conventions. Second, AI-design ideation benefits from separating AI role, AI timing, and AI initiative. Participants repeatedly found the upfront/in-due-course and reactive/proactive distinctions useful because these prompts made implicit journey-level choices explicit, even when they did not adopt the workbook as a complete process.

8. Reflections & Discussions

In this section, we discuss limitations of our existing work, opportunities for future research, and expansions based on our study findings.

8.1. The Value and Limitation of the A3I WORKBOOK *as a Structured "Thinking Aid" to Support Designers' Early Ideation Stage*

Reflecting on the findings from the studies, we believe that the value of our workbook is providing a "thinking aid" that supports designers in surfacing and comparing their own ideas for creating an A3I application. The exercise of reflecting on one's design vis-à-vis the workbook's steps was reported by participants as helping them articulate design choices and explore alternatives. Focusing on guiding designers through a high-level structured process, the workbook leaves the execution of each step to the designers, thus trading off low-level, detailed instructions, e.g., providing a "look-up" manual finding design solutions. For example, the workbook guides designers to consider that there are more than one way to invoke AI automation (upfront vs.in due course) without listing all possible interaction designs to realize each invocation approach. Overall, we consider our workbook analogous to other aids for high-level, structural thinking, such as SWOT [58] or TRIZ [59]. This positioning also clarifies the relationship between A3I and design patterns. A3I is a thinking aid that helps designers identify a need or possibility, while a pattern library can suggest candidate instantiated solutions. For example, a team may use A3I to identify that AI should summarize evidence only after a user inspects a subset of reviews; a pattern library could then suggest concrete disclosure, comparison, or annotation patterns for realizing that idea. One possible future direction is to use the current workbook as an index that links to a library of "design patterns" [12] for each aspect of an A3I application.

8.2. What the Design Artifacts Reveal across the Workbook-Use Studies

The design artifacts provide concrete traces of how participants interpreted and used the workbook prompts. In the student deployment, T6's workbook and prototype documented a three-level recipe view as visual backup for an audio assistant and a proactive assistant that inferred the next needed step from speech. T4's workbook walkthrough recorded reconsideration of a plain-text generative explanation, while T7 considered proactive support before retaining a reactive design to preserve user control. In the professional exercise, the completed decks made different interpretations of the hierarchy prompt visible: P2 represented review access as a retrieval sequence, P4, P6, and P7 listed review attributes, and P5 grouped reviews by reviewer.

Taken together, these traces show that the workbook's prompts made alternatives and interpretations external and available for reflection. They also reinforce the contrast between initiative, which participants could use to compare interaction alternatives, and hierarchy, which the professionals often translated into sequence, attributes, or grouping rather than a nested structure. However, the studies did not preserve a consistent pre-workbook artifact baseline, and the artifacts were not collected or coded for a comparative outcome analysis. They therefore corroborate what participants considered but cannot establish that the workbook caused additional or better ideas.

8.3. The A3I WORKBOOK*'s Scope and Relations to Other Design Activities*

The scope of the A3I WORKBOOKis focused on A3I applications to support users' information processing tasks. We believe this type of applications is fairly common amongst knowledge workers—from physicians using AI to assist diagnoses based on medical images to artists who need to develop their work amongst a large number of AI-generated examples. Our studies suggest that the workbook's utility lies in informing someone who is at the early stage of ideating designs of A3I applications. In the meantime, there are various other important design

issues that our workbook does not support, such as ways to design AI that can achieve human trust (addressed in [60]), explainability (reviewed in [61]), agency (discussed in [62]), and mental models (studied in [63]). We believe that our workbook is extensible to connect to other design activities and to encompass the above-mentioned unaddressed issues. Foremost, as pointed out by one of our expert interviewees, the use of the workbook can serve as a natural extension of some initial user research and understanding of how users' current approach to process certain information. As another example, we can also extend the design choice of reactive vs.proactive AI to take into consideration Horvitz's multiple principles of mixed-initiative interaction [7]. Likewise, the "in due course" design choice can be further expanded to consider different nuanced techniques discussed in AI overreliance research [8,9].

8.4. From Ideation to Implementation

A3I's output is intended to inform downstream design and engineering artifacts, not replace them. For example, a workbook response that AI should proactively recommend a subset of information can become a user story about receiving timely recommendations, an interaction flow showing when recommendations appear and how users accept or dismiss them, an AI behavior specification describing triggers and confidence thresholds, and a data requirement describing which signals are needed to make the recommendation. Similarly, a response that AI should summarize evidence in due course can become a prototype annotation, a content-quality rubric, and a technical brief for summarization model evaluation. Table 4 lists further examples. This bridge is especially important for cross-functional teams: designers can use A3I to articulate interaction intent, while data scientists and engineers can examine whether the required data, model capability, latency, privacy constraints, and failure-handling assumptions are plausible.

Table 4. Examples of how workbook outputs can feed downstream design and engineering artifacts.

Workbook Output	Downstream Artifact
Hierarchic information view.	Information architecture diagram; navigation model; content inventory.
Filter criteria.	User flow; prototype state diagram; data requirement for filterable attributes.
Predictive AI support.	Model requirement specification; evaluation metric; confidence and error-handling plan.
Generative AI support.	Interaction specification; prompt/content-quality rubric; generated-output review flow.
Initiative choice.	User story; trigger specification; notification and override behavior.
Data and feasibility assumptions.	Data inventory; privacy and fairness review; implementation risk register.

8.5. Data and Feasibility Assumptions

The current workbook primarily encourages divergent ideation and intentionally avoids suppressing ideas too early based on feasibility concerns. However, our expert interviews and professional-design exercise also suggest that teams need a way to surface data assumptions before moving downstream. A future version of the workbook should add a preliminary prompt, perhaps after the system overview, asking what data is required, where it will come from, whether it is available, and what privacy, bias, quality, or governance concerns might shape the design. We view this as surfacing assumptions and constraints without prematurely closing the design space. Because this prompt was not part of the evaluated workbook, it should be treated as a proposed revision rather than a finding supported by the present studies.

8.6. Hierarchy as a Translation Problem between Theory and Practice

The professional design exercise suggests that the difficulty with "hierarchy" was both terminological and conceptual. The term could be confused with interface hierarchy or established information-architecture work, and participants reported that they were not accustomed to beginning AI-feature design by treating task-relevant information as a hierarchy. This contrast was most visible between students designing applications from scratch, who could invent nested structures for recipes or course information, and professionals working within product ecosystems whose navigation and dashboard conventions already existed.

We therefore retain hierarchy as one recurring organization strategy that can prompt exploration, but not as a universal prerequisite or the only way to structure information. Grouping, clustering, sequence, and spatial organization may be equally appropriate for a task. A post-study workbook revision will ask designers to "organize or represent the information," with hierarchy offered as one example and with links to familiar artifacts such as information-architecture diagrams and task flows. The evaluated workbook used the original hierarchy wording; this broader prompt is a revision motivated by the findings and will be incorporated into a future version of the workbook.

8.7. Generalizable Lessons for Design-Support Tools

Across the three studies, we draw four lessons for designing design-support tools beyond A3I. First, scaffolds are most useful when they expose dimensions that designers might otherwise leave implicit; participants repeatedly valued prompts that separated AI role, timing, and initiative. Second, designers reinterpret theoretical constructs through existing professional vocabularies, so a scaffold must bridge academic terms to familiar design artifacts and language. Third, scaffolds must balance conceptual completeness with perceived workload by supporting both full walkthroughs and lightweight reflective use. Fourth, AI-design ideation benefits from separating AI role, AI timing, and AI initiative because these distinctions help teams compare alternatives before committing to a single interaction direction.

8.8. Evaluation Limitations

The studies provide formative evidence about how participants interpreted and used A3I, but they do not provide comparative evidence that A3I produces broader, more diverse, higher-quality, or more feasible designs than other approaches. We did not compare A3I against unstructured brainstorming, other frameworks, generic workbooks, or no-workbook conditions. We also did not use an objective design-quality metric, and the present studies relied on interviews, workbook artifacts, walkthroughs, and short-term design exercises rather than longitudinal deployment. The classroom deployment provided rich access to design process and artifacts, but the participants were students rather than professional designers. The expert interviews helped identify industry-facing concerns, but experts reviewed and discussed the workbook rather than using it in their own design projects. The UX professional design exercise brought the workbook closer to practice, but it was short, task-based, and conducted with one company. The classroom study also involved the first author as instructor, grader, recruiter, and interviewer. Although participation and responses did not affect grades and projects were peer-graded by other teams, this relationship may have introduced demand effects or discouraged criticism. The samples were also culturally and disciplinarily narrow: all course participants studied Electrical Engineering and Computer Science, the professional designers came from one North American company, and the three-expert sample was not designed to test cultural transfer. Whether the workbook's vocabulary, organization strategies, and workflow assumptions transfer across cultures, languages, disciplines, and different information-processing practices remains untested.

8.9. Scope Limitations

A3I is focused on information-intensive applications where the user's primary challenge is collecting, organizing, interpreting, comparing, transforming, or acting on substantial information. It is not intended to cover all AI-infused applications. Applications centered mainly on physical manipulation, delegated command execution, generic chat, or human-robot collaboration may require different scaffolds unless information processing remains the central user task. The current evidence also leaves open how well A3I transfers across organizations, domains, team structures, and product maturity levels.

8.10. Framework Limitations

The current framework and workbook have limited treatment of data feasibility constraints, which is why we propose adding the Data and Feasibility Check as a preliminary prompt. Terminology also remains under refinement, especially around hierarchy, information chunks, and mixed initiative. Finally, the workbook may require domain-specific adaptation: teams may need examples, templates, and prompts calibrated to particular information types such as text, images, logs, medical data, or design examples. Thus, the current evidence supports claims about preliminary utility, interpretive friction, and revision opportunities, while leaving comparative effectiveness and long-term adoption as future work.

8.11. Implication for Broader Types of AI-Infused Systems

As we focus on the information-intensive applications [2], there are other types of AI-infused systems that A3I has not considered, such as people delegating a task to a voice assistant (e.g., Siri and Alexa), resolving an issue by messaging with a customer service chatbot, and working with robots to perform physical tasks. One primary distinction is that these other types of scenarios normally do not directly expose the user to a large amount of information (even though such information might still exist, e.g., the many paths a field robot can maneuver autonomously). As such, it becomes questionable whether the sense-making part of the A3I WORKBOOK still applies. Further, since some other human-AI collaboration workflow is governed by turn-taking (e.g., chatting with a bot), it also becomes questionable whether the when/how-to-invoke-AI part of our workbook still applies. To

better explore these limits, future work can evaluate our workbook in these above-mentioned scenarios and perhaps develop variations of the workbook that cater to these other types of AI-infused systems.

8.12. Towards a Library of Reusable Design Examples and Templates

Currently, the A3I WORKBOOKaims to provide high-level guidance that directs designers' ideation to encompass different aspects of AI-infused information processing. In the future, we plan to expand the workbook with organized sets of examples extracted from literature and existing software products. In so doing, we hope to bridge the existing gaps of understanding concepts in the workbook. For example, a designer can browse examples of a hierarchic information view while ideating their own design. Further, we can gather more examples by considering each aspect in the framework separately in the context of the general HCI techniques, such as early work on Treemaps [64] and Zoomable UI [16], which would be excellent examples to help a designer understand what a hierarchic view looks like. Moreover, future work can be devoted to provide a library of reusable mark-up templates, sample code, or enabling APIs pertinent to these examples. One approach is to consider different types of information (e.g., text, image, and audio) and organize resources accordingly, e.g., abstractive summarization and question generation models for using AI generatively to help humans process text.

8.13. Lessons Learned of Deploying the A3I WORKBOOK *from Academic to Industrial Settings*

We reflect on two lessons learned in the process of developing the A3I WORKBOOKand deploying it from academic to individual settings. First, the trade-offs between lightweight and depth. Unlike course project teams who used the workbook for weeks, designers in industry are often busy and would be reluctant and unable to devote as much their time and effort. Our solution to encourage adoption was making the workbook more lightweight, which, however, traded-off details that designers might find necessary once they adopt the workbook. Second, the differences in language and concepts. In contrast to the open-endedness in an academic setting, designers in industry might have developed their own idiosyncratic vocabulary of characterizing the design process. We had to translate academic terminologies and concepts to the ones appropriate for the industrial context. Still, we found gaps where certain concepts that make sense in the academic deployment (e.g., creating a hierarchic information view) struggled to resonate with professional designers. These two lessons combined suggest the need of gradually introducing and bridging an academically-developed process to the industry. Our design exercise can be thought of as the "first look" of the workbook to surface potential usages, confusions and difficulties. Subsequently, there need to be more in-depth sessions, each helping designers better understand and respond to one small facet of the workbook.

9. Conclusions

AI holds the promise of augmenting a wide range of knowledge workers' tasks, many of which involve processing domain-specific information. A3I addresses this information-intensive subset; physical manipulation and generic chat fall outside its scope when information processing is not the user's central challenge. Designing such applications is especially challenging in the early divergent stage, when teams must decide what role AI should play, where it should enter the information-processing workflow, and how initiative should be divided between people and AI. By unifying theories and design examples, we developed A3I as a conceptual framework for structuring this early design exploration [2]. The accompanying A3I WORKBOOK operationalizes the framework as a high-level thinking aid that helps designers explore what an AI system can do and when/how to invoke AI to help a human process information. Our formative studies document how students, industrial researchers, and UX professionals interpreted, used, and critiqued the workbook. Participants reported value in the workbook's structured thought process and its ability to surface design dimensions they had not explicitly considered, while also revealing practical frictions around terminology, workload, and especially the translation of hierarchy into professional design practice. These findings support A3I as complementary early-stage scaffolding, not as a replacement for existing design methods or as evidence of superiority over other ideation approaches. We hope the lessons learned (see Section 8.13) in developing and deploying the A3I WORKBOOKfrom academic to industrial settings can inform future efforts to translate HCI theory into actionable design scaffolds. In the short term, future work should link the workbook's prompts to reusable patterns and worked examples for text- and image-intensive applications. In the longer term, A3I should be adapted to foundation-model and agentic workflows and developed into domain-specific versions whose vocabulary, risks, and examples reflect users' professional practice.

Funding

This research was funded by the National Science Foundation, grant number 2047297.

Institutional Review Board Statement

The studies were approved by the Institutional Review Board of the University of California, Los Angeles (protocol code IRB-21-0139, renewed annually; most recent renewal on 29 May 2026).

Informed Consent Statement

Informed consent was obtained from all subjects involved in the study.

Data Availability Statement

The A3I workbook is available at https://hci.ucla.edu/a3i. The interview and study data are not publicly available to protect participant privacy, in line with the IRB protocol.

Conflicts of Interest

The author declares no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this work, the author used Claude (Anthropic) to assist with proofreading and copy-editing. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Appendix A. Literature Search Process for Finding Design Examples

It is not our main objective to conduct a thorough literature survey on human-AI interaction; rather, we reviewed literature most relevant to A3I applicationsto inform the formulation of the A3I framework. Below, we report our search process for future work to build on and continue to develop A3I.

Given the multitude of contribution types [65] in human-AI collaboration research, our particular focus is on prior work with an artifact contribution, i.e., that invents new designs of AI-infused systems with the primary task of processing some sorts of information.

Appendix A.1. Inclusion and Exclusion Criteria

Given the aforementioned focus, we only included research that resulted in a system artifact. Further, such systems need to contain intelligent features, i.e., be able to automate tasks for humans involving processing domain-specific information. Since both human and AI are critical in our analysis, we had to exclude fully or mostly automated systems (i.e., those that reduce human input to the minimum, such as a few button presses). We excluded systems whose primary task was generic dialog or delegated command execution rather than information-intensive work; conversational modality by itself was not an exclusion criterion. Finally, human-robot collaborative systems intended for physical tasks, rather than information processing, were also excluded. Note that we excluded a large amount of work on human-AI interaction and collaboration because their main contribution type is empirical studies and findings whereas our main focus is on artifacts and systems.

Appendix A.2. Information Source, Search Strategy and Selection Process

We conducted the search on Google Scholar, ACM Digital Library, and IEEE Xplore as the main sources. Our search query was a tuple (A, B) where $A \in$ {"AI-infused", "mixed-initiative", "human-centered AI", "collaborative AI"} and $B \in$ {"data", "information", "system", "tool"}. Upon gathering the initial search results, we performed a manual examination (by the first author) to determine whether a paper met the inclusion / exclusion criteria for further review. No automation tool was used. For each paper we included, we focused on the system artifact as the main subject of analysis (even though some papers position a theory or empirical study as part of their contributions).

Appendix A.3. Classification Rules

Table A1 classifies reported system features or workflow states rather than assigning each paper to one exclusive system type. Multiple marks are therefore allowed when a system exhibits different behaviors in different workflows, and a blank cell means that we did not identify that feature in the paper.

Overlap between reactive and proactive behavior exposes a control–anticipation tradeoff: a system can preserve explicit user control in some states while initiating support where needs can be inferred. Likewise, combining

predictive and generative support lets a workflow alternate between narrowing alternatives and expanding or transforming the information available for judgment.

Table A1. We identified prior work that provides system artifacts as examples to illustrate the A3I design strategies.

Human–AI Collaborative Systems Processing Information	What Can an AI System Do to Help a Human Process Information?				When and How to Invoke AI as a Human Processes Information?			
	Sense-Making		Decision-Making		AI Automation		Mixed Initiative	
	Hierarchy	Filter	Predictive	Generative	Upfront	In Due Course	Reactive	Proactive
Horvitz [7]			✓	✓			✓	✓
Igarashi & Hughes [66]				✓			✓	
Ai-Chang et al. [67]		✓		✓			✓	
Xiao et al. [21]	✓		✓					✓
Fogarty et al. [68]		✓					✓	
Chau et al. [20]	✓	✓		✓			✓	
Guo et al. [69]		✓		✓				✓
Umetani et al. [29]		✓	✓	✓			✓	
Gao et al. [50]		✓					✓	✓
Kazi et al. [30]		✓		✓			✓	
Chen et al. [28]		✓	✓	✓			✓	
Nguyen et al. [31]		✓	✓	✓			✓	
Wall et al. [27]		✓	✓				✓	
Lee et al. [70]				✓			✓	
Cai et al. [45]		✓		✓			✓	
Yang et al. [38]			✓		✓			
Xie et al. [24]	✓	✓	✓		✓			
Koch et al. [43]		✓		✓			✓	✓
Wang et al. [71]	✓	✓	✓		✓			
Jeon et al. [26]		✓		✓	✓			
Gu et al. [49]		✓	✓		✓	✓	✓	
Lee et al. [48]		✓	✓		✓			
Chan et al. [42]				✓			✓	
Liu et al. [39]	✓	✓	✓		✓			
Lee et al. [23]	✓		✓		✓			
Gu et al. [25]	✓	✓	✓		✓			
Zhang et al. [41]				✓			✓	✓
Yuan et al. [72]				✓			✓	
Koyama & Goto [32]		✓		✓			✓	
Ko et al. [73]		✓		✓			✓	
Liu et al. [74]		✓		✓			✓	
Gu et al. [52]	✓	✓	✓				✓	✓
Gebreegziabher et al. [40]		✓	✓	✓			✓	
Chen et al. [44]		✓	✓	✓		✓	✓	

Note. Multiple marks in a row indicate that the reported system exhibits multiple support or invocation behaviors in different features or workflow states; blank cells mean “not identified”.

Appendix A.4. Analysis Method

We gathered a total of 34 papers from the above-mentioned search process, most of which were published at CHI (15), UIST (9), and CSCW (3). For this revision, we assigned each paper one primary application domain

based on its target task and summarized the user-facing predictive/generative support coding shown in Table A1.

Creative, design, and media-authoring systems form the largest domain group (13/34, 38.2%), followed by healthcare and clinical systems (7/34, 20.6%). Because the corpus is purposive rather than exhaustive and is concentrated in these domains, the distribution characterizes the examples used to derive and illustrate A3I; it does not establish the framework's generality across all AI-infused applications. Using an iterative qualitative coding process, the first author itemized key aspects of each system's design, including its information operations, UI structure, workflow, type of AI support, and when/how AI was invoked. The first author then conducted multiple iterations to merge the initial codes into higher-level themes and organize them into the framework (Table A1).

Table A2. Descriptive composition of the 34-paper corpus. Percentages use the full corpus as denominator.

Coding Category	Count	Percent
Primary application domain		
Creative/design/media authoring	13	38.2%
Healthcare/clinical	7	20.6%
Knowledge work/information seeking/news	5	14.7%
Data analysis/programming	4	11.8%
Accessibility	3	8.8%
Planning/scheduling	2	5.9%
User-facing AI support coding		
Predictive only	11	32.4%
Generative only	15	44.1%
Both	6	17.6%
Neither identified	2	5.9%

References

1. Card, S.K.; Moran, T.P.; Newell, A. *The Psychology of Human-Computer Interaction*; CRC Press: Boca Raton, FL, USA, 2018. https://doi.org/10.1201/9780203736166.
2. Zarmer, C.L.; Chew, C. Frameworks for Interactive, Extensible, Information-Intensive Applications. In Proceedings of the 5th Annual ACM Symposium on User Interface Software and Technology, Monterey, CA, USA, 15–18 November 1992; pp. 33–41. https://doi.org/10.1145/142621.142626.
3. Riedl, M.O. Human-Centered Artificial Intelligence and Machine Learning. *Hum. Behav. Emerg. Technol.* **2019**, *1*, 33–36. https://doi.org/10.1002/hbe2.117.
4. Amershi, S.; Weld, D.; Vorvoreanu, M.; et al. Guidelines for Human-AI Interaction. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–13. https://doi.org/10.1145/3290605.3300233.
5. Shneiderman, B. *Human-Centered AI*; Oxford University Press: Oxford, UK, 2022. https://doi.org/10.1093/oso/9780192845290.001.0001.
6. Russell, D.M.; Stefik, M.J.; Pirolli, P.; et al. The Cost Structure of Sensemaking. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Amsterdam, The Netherlands, 24–29 April 1993; pp. 269–276. https://doi.org/10.1145/169059.169209.
7. Horvitz, E. Principles of Mixed-Initiative User Interfaces. In Proceedings of the SIGCHI conference on Human Factors in Computing Systems, Pittsburgh, PA, USA, 15–20 May 1999; pp. 159–166. https://doi.org/10.1145/302979.303030.
8. Buçinca, Z.; Malaya, M.B.; Gajos, K.Z. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making. *Proc. ACM Hum.-Comput. Interact.* **2021**, *5*, 1–21. https://doi.org/10.1145/3449287.
9. Passi, S.; Vorvoreanu, M. Overreliance on AI Literature Review. Available online: https://www.microsoft.com/en-us/research/wp-content/uploads/2022/06/Aether-Overreliance-on-AI-Review-Final-6.21.22.pdf (accessed on 19 October 2026).
10. Hands-On Tools for Building Effective Human-AI Experiences . Available online: https://www.microsoft.com/en-us/haxtoolkit/ (accessed on 19 August 2026).
11. People + AI Guidebook. Available online: https://pair.withgoogle.com/guidebook (accessed on 19 October 2026).
12. Tidwell, J. *Designing Interfaces: Patterns for Effective Interaction Design*; O'Reilly Media, Inc.: Sebastopol, CA, USA, 2010.
13. Pirolli, P.; Card, S. Information Foraging. *Psychol. Rev.* **1999**, *106*, 643–675. https://doi.org/10.1037/0033-295x.106.4.643.
14. Shneiderman, B. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In Proceedings of the 1996 IEEE Symposium on Visual Languages, Boulder, CO, USA, 3–6 September 1996; pp. 336–343. https://doi.org/10.1109/vl.1996.545307.

15. Shneiderman, B. Tree Visualization with Tree-Maps: 2-D Space-Filling Approach. *ACM Trans. Graph.* **1992**, *11*, 92–99. https://doi.org/10.1145/102377.115768.
16. Bederson, B.B.; Hollan, J.D. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology, Marina del Rey, CA, USA, 2–4 November 1994; pp. 17–26. https://doi.org/10.1145/192426.192435.
17. Simon, H.A. *The Sciences of the Artificial*; The MIT Press: Cambridge, MA, USA, 2019. https://doi.org/10.7551/mitpress/12107.001.0001.
18. Meadows, D.H.; Wright, D. *Thinking in Systems*; Routledge: London, UK, 2008.
19. Shneiderman, B.; Maes, P. Direct Manipulation vs. Interface Agents. *Interactions* **1997**, *4*, 42–61. https://doi.org/10.1145/267505.267514.
20. Chau, D.H.; Kittur, A.; Hong, J.I.; et al. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. In Proceedings of the CHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada, 7–12 May 2011; pp. 167–176. https://doi.org/10.1145/1978942.1978967.
21. Xiao, J.; Zhang, X.; Cheatle, P.; et al. Mixed-Initiative Photo Collage Authoring. In Proceedings of the 16th ACM international conference on Multimedia, Vancouver, BC, Canada, 26–31 October 2008; pp. 509–518. https://doi.org/10.1145/1459359.1459427.
22. Glueck, M.; Grossman, T.; Wigdor, D. A Model of Navigation for Very Large Data Views. In Proceedings of the Graphics Interface, Regina, SK, Canada, 29–31 May 2013; pp. 9–16.
23. Lee, J.; Herskovitz, J.; Peng, Y.H.; et al. ImageExplorer: Multi-Layered Touch Exploration to Encourage Skepticism towards Imperfect AI-Generated Image Captions. In Proceedings of the CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April–5 May 2022; pp. 1–15. https://doi.org/10.1145/3491102.3501966.
24. Xie, Y.; Chen, M.; Kao, D.; et al. CheXplain: Enabling Physicians to Explore and Understand Data-Driven, AI-Enabled Medical Imaging Analysis. In Proceedings of the CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 25–30 April 2020; pp. 1–13. https://doi.org/10.1145/3313831.3376807.
25. Gu, H.; Liang, Y.; Xu, Y.; et al. Improving Workflow Integration with xPath: Design and Evaluation of a Human-AI Diagnosis System in Pathology. *ACM Trans. Comput.-Hum. Interact.* **2023**, *30*, 28:1–28:37. https://doi.org/10.1145/3577011.
26. Jeon, Y.; Jin, S.; Shih, P.C.; et al. FashionQ: An AI-Driven Creativity Support Tool for Facilitating Ideation in Fashion Design. In Proceedings of the CHI Conference on Human Factors in Computing Systems, Yokohama, Japan, 8–13 May 2021; pp. 1–18. https://doi.org/10.1145/3411764.3445093.
27. Wall, E.; Das, S.; Chawla, R.; et al. Podium: Ranking Data Using Mixed-Initiative Visual Analytics. *IEEE Trans. Vis. Comput. Graph.* **2018**, *24*, 288–297. https://doi.org/10.1109/tvcg.2017.2745078.
28. Chen, X.A.; Tao, Y.; Wang, G.; et al. Forte: User-Driven Generative Design. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 21–26 April 2018; pp. 1–12. https://doi.org/10.1145/3173574.3174070.
29. Umetani, N.; Igarashi, T.; Mitra, N.J. Guided Exploration of Physically Valid Shapes for Furniture Design. *ACM Trans. Graph.* **2012**, *31*, 1–11. https://doi.org/10.1145/2185520.2185582.
30. Kazi, R.H.; Grossman, T.; Cheong, H.; et al. DreamSketch: Early Stage 3D Design Explorations with Sketching and Generative Design. In Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology, Québec City, QC, Canada, 22–25 October 2017; pp. 401–414. https://doi.org/10.1145/3126594.3126662.
31. Nguyen, A.T.; Kharosekar, A.; Krishnan, S.; et al. Believe It or Not. In Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology, Berlin, Germany, 14–17 October 2018; pp. 189–199. https://doi.org/10.1145/3242587.3242666.
32. Koyama, Y.; Goto, M. BO as Assistant: Using Bayesian Optimization for Asynchronously Generating Design Suggestions. In Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology, Bend, OR, USA, 29 October–2 November 2022; pp. 1–14. https://doi.org/10.1145/3526113.3545664.
33. Simon, H.A. *Models of Bounded Rationality: Empirically Grounded Economic Reason*; MIT Press: Cambridge, MA, USA, 1997.
34. Horvitz, E.J.; Breese, J.S.; Henrion, M. Decision Theory in Expert Systems and Artificial Intelligence. *Int. J. Approx. Reason.* **1988**, *2*, 247–302. https://doi.org/10.1016/0888-613x(88)90120-x.
35. Shortliffe, E.H. A Rule-Based Computer Program for Advising Physicians regarding Antimicrobial Therapy Selection. In Proceedings of the 1974 Annual Conference, San Diego, CA, USA, November 1974; Volume 2, pp. 739. https://doi.org/10.1145/1408800.1408906.
36. Blois, M.S. *Information and Medicine: The Nature of Medical Descriptions*; University of California Press: Berkeley, CA, USA, 1984.
37. Guilford, J.P. Three Faces of Intellect. In *Studies in Individual Differences: The Search for Intelligence*; Appleton-Century-Crofts: East Norwalk, CT, USA, 1961; pp. 756–774. https://doi.org/10.1037/11491-066.
38. Yang, Q.; Steinfeld, A.; Zimmerman, J. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. In Proceedings of the CHI'19: CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–11. https://doi.org/10.1145/3290605.3300468.

39. Liu, X.B.; Wang, R.; Li, D.; et al. CrossA11y: Identifying Video Accessibility Issues via Cross-Modal Grounding. In Proceedings of the UIST'22: The 35th Annual ACM Symposium on User Interface Software and Technology, Bend, OR, USA, 29 October–2 November 2022; pp. 1–14. https://doi.org/10.1145/3526113.3545703.
40. Gebreegziabher, S.A.; Zhang, Z.; Tang, X.; et al. PaTAT: Human-AI Collaborative Qualitative Coding with Explainable Interactive Rule Synthesis. In Proceedings of the CHI'23: CHI Conference on Human Factors in Computing Systems, Hamburg, Germany, 23–28 April 2023; pp. 1–19. https://doi.org/10.1145/3544548.3581352.
41. Zhang, Z.; Xu, Y.; Wang, Y.; et al. StoryBuddy: A Human-AI Collaborative Chatbot for Parent-Child Interactive Storytelling with Flexible Parental Involvement. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI'22), New Orleans, LA, USA, 29 April–5 May 2022; pp. 1–21. https://doi.org/10.1145/3491102.3517479.
42. Chan, J.; Chang, J.C.; Hope, T.; et al. SOLVENT: A Mixed Initiative System for Finding Analogies between Research Papers. *Proc. ACM Hum.-Comput. Interact.* **2018**, *2*, 1–21. https://doi.org/10.1145/3274300.
43. Koch, J.; Taffin, N.; Beaudouin-Lafon, M.; et al. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. *Proc. ACM Hum.-Comput. Interact.* **2020**, *4*, 1–27. https://doi.org/10.1145/3392850.
44. Chen, X.A.; Wu, C.S.; Murakhovs'ka, L.; et al. Marvista: Exploring the Design of a Human-AI Collaborative News Reading Tool. *ACM Trans. Comput.-Hum. Interact.* **2023**, *30*, 1–27. https://doi.org/10.1145/3609331.
45. Cai, C.J.; Reif, E.; Hegde, N.; et al. Human-Centered Tools for Coping with Imperfect Algorithms during Medical Decision-Making. In Proceedings of the CHI'19: CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–14. https://doi.org/10.1145/3290605.3300234.
46. Park, J.S.; O'Brien, J.; Cai, C.J.; et al. Generative Agents: Interactive Simulacra of Human Behavior. In Proceedings of the UIST'23: The 36th Annual ACM Symposium on User Interface Software and Technology, San Francisco, CA, USA, 29 October–1 November 2023; pp. 1–22. https://doi.org/10.1145/3586183.3606763.
47. Lee, J.D.; See, K.A. Trust in Automation: Designing for Appropriate Reliance. *Hum. Factors* **2004**, *46*, 50–80. https://doi.org/10.1518/hfes.46.1.50.30392.
48. Lee, M.H.; Siewiorek, D.P.; Smailagic, A.; et al. A Human-AI Collaborative Approach for Clinical Decision Making on Rehabilitation Assessment. In Proceedings of the CHI'21: CHI Conference on Human Factors in Computing Systems, Yokohama, Japan, 8–13 May 2021; pp. 1–14. https://doi.org/10.1145/3411764.3445472.
49. Gu, H.; Huang, J.; Hung, L.; et al. Lessons Learned from Designing an AI-Enabled Diagnosis Tool for Pathologists. *Proc. ACM Hum.-Comput. Interact.* **2021**, *5*, 1–25. https://doi.org/10.1145/3449084.
50. Gao, T.; Dontcheva, M.; Adar, E.; et al. DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In Proceedings of the UIST'15: The 28th Annual ACM Symposium on User Interface Software and Technology, Charlotte, NC, USA, 11–15 November 2015; pp. 489–500. https://doi.org/10.1145/2807442.2807478.
51. Weiser, M.; Brown, J.S. The Coming Age of Calm Technology. In *Beyond Calculation*; Springer New York: New York, NY, 1997; pp. 75–85. https://doi.org/10.1007/978-1-4612-0685-9_6.
52. Gu, H.; Yang, C.; Haeri, M.; et al. Augmenting Pathologists with NaviPath: Design and Evaluation of a Human-AI Collaborative Navigation System. In Proceedings of the CHI'23: CHI Conference on Human Factors in Computing Systems, Hamburg, Germany, 23–28 April 2023; pp. 1–19. https://doi.org/10.1145/3544548.3580694.
53. Liu, X.B.; Kirilyuk, V.; Yuan, X.; et al. Visual Captions: Augmenting Verbal Communication with on-the-Fly Visuals. In Proceedings of the CHI'23: CHI Conference on Human Factors in Computing Systems, Hamburg, Germany, 23–28 April 2023; pp. 1–20. https://doi.org/10.1145/3544548.3581566.
54. Braun, V.; Clarke, V. Thematic Analysis. In *APA Handbook of Research Methods in Psychology, Vol 2: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological*; Cooper, H., Camic, P.M., Long, D.L., Eds.; American Psychological Association: Washington, DC, USA, 2012; pp. 57–71. https://doi.org/10.1037/13620-004.
55. Grant, A.M. *Think Again: The Power of Knowing What You Don't Know*; Viking: New York, NY, USA, 2021.
56. Design Council. A Study of the Design Process. Available online: https://www.designcouncil.org.uk/fileadmin/uploads/dc/Documents/ElevenLessons_Design_Council%2520%25282%2529.pdf (accessed on 23 February 2024).
57. Yang, Q.; Steinfeld, A.; Rosé, C.; et al. Re-Examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In Proceedings of the CHI'20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 25–30 April 2020; pp. 1–13. https://doi.org/10.1145/3313831.3376301.
58. SWOT Analysis. Available online: https://en.wikipedia.org/w/index.php?title=SWOT_analysis&oldid=1231505823#History (accessed on 19 October 2026).
59. TRIZ. Available online: https://en.wikipedia.org/w/index.php?title=TRIZ&oldid=1233265087 (accessed on 19 October 2026).
60. Ueno, T.; Sawa, Y.; Kim, Y.; et al. Trust in Human-AI Interaction: Scoping out Models, Measures, and Methods. In Proceedings of the CHI'22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April–5 May 2022; pp. 1–7. https://doi.org/10.1145/3491101.3519772.
61. Mueller, S.T.; Hoffman, R.R.; Clancey, W.; et al. Explanation in Human-AI Systems: A Literature Meta-Review, Synopsis of Key Ideas and Publications, and Bibliography for Explainable AI. *arXiv* **2019**, arXiv:1902.01876.
62. Sundar, S.S. Rise of Machine Agency: A Framework for Studying the Psychology of Human–AI Interaction (HAII). *J. Comput.-Mediat. Commun.* **2020**, *25*, 74–88. https://doi.org/10.1093/jcmc/zmz026.

63. Khadpe, P.; Krishna, R.; Fei-Fei, L.; et al. Conceptual Metaphors Impact Perceptions of Human-AI Collaboration. *Proc. ACM Hum.-Comput. Interact.* **2020**, *4*, 163:1–163:26. https://doi.org/10.1145/3415234.
64. Shneiderman, B.; Plaisant, C. Treemaps for Space-Constrained Visualization of Hierarchies. Available online: https://www.cs.umd.edu/hcil/treemap-history/ (accessed on 19 October 2026).
65. Wobbrock, J.O.; Kientz, J.A. Research Contributions in Human-Computer Interaction. *Interactions* **2016**, *23*, 38–44. https://doi.org/10.1145/2907069.
66. Igarashi, T.; Hughes, J.F. A Suggestive Interface for 3D Drawing. In Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST'01), Orlando, FL, USA, 11–14 November 2001; pp. 173–181. https://doi.org/10.1145/502377.502379.
67. Ai-Chang, M.; Yglesias, J.; Chafin, B.G.; et al. MAPGEN: Mixed-Initiative Planning and Scheduling for the Mars Exploration Rover Mission. *IEEE Intell. Syst.* **2004**, *19*, 8–12. https://doi.org/10.1109/mis.2004.1265878.
68. Fogarty, J.; Tan, D.; Kapoor, A.; et al. CueFlik: Interactive Concept Learning in Image Search. In Proceedings of the CHI'08: CHI Conference on Human Factors in Computing Systems, Florence, Italy, 5–10 April 2008; pp. 29–38. https://doi.org/10.1145/1357054.1357061.
69. Guo, P.J.; Kandel, S.; Hellerstein, J.M.; et al. Proactive Wrangling: Mixed-Initiative End-User Programming of Data Transformation Scripts. In Proceedings of the UIST'11: The 24th Annual ACM Symposium on User Interface Software and Technology, Santa Barbara, CA, USA, 16–19 October 2011; pp. 65–74. https://doi.org/10.1145/2047196.2047205.
70. Lee, B.; Jin, T.; Lee, S.H.; et al. SmartManikin: Virtual Humans with Agency for Design Tools. In Proceedings of the CHI'19: CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–13. https://doi.org/10.1145/3290605.3300814.
71. Wang, R.; Chen, Z.; Zhang, M.R.; et al. Revamp: Enhancing Accessible Information Seeking Experience of Online Shopping for Blind or Low Vision Users. In Proceedings of the CHI'21: CHI Conference on Human Factors in Computing Systems, Yokohama, Japan, 8–13 May 2021; pp. 1–14. https://doi.org/10.1145/3411764.3445547.
72. Yuan, A.; Coenen, A.; Reif, E.; et al. Wordcraft: Story Writing with Large Language Models. In Proceedings of the IUI'22: 27th International Conference on Intelligent User Interfaces, Helsinki, Finland, 22–25 March 2022; pp. 841–852. https://doi.org/10.1145/3490099.3511105.
73. Ko, H.K.; An, S.; Park, G.; et al. We-Toon: A Communication Support System between Writers and Artists in Collaborative Webtoon Sketch Revision. In Proceedings of the UIST'22: The 35th Annual ACM Symposium on User Interface Software and Technology, Bend, OR, USA, 29 October–2 November 2022; pp. 1–14. https://doi.org/10.1145/3526113.3545612.
74. Liu, V.; Qiao, H.; Chilton, L. Opal: Multimodal Image Generation for News Illustration. In Proceedings of the UIST'22: The 35th Annual ACM Symposium on User Interface Software and Technology, Bend, OR, USA, 29 October–2 November 2022; pp. 1–17. https://doi.org/10.1145/3526113.3545621.