Developing Design Tools to Support Envisioning with Data and AI

Nur Yildirim
HCI Institute
Carnegie Mellon University
Pittsburgh, PA, USA
yildirim@cmu.edu

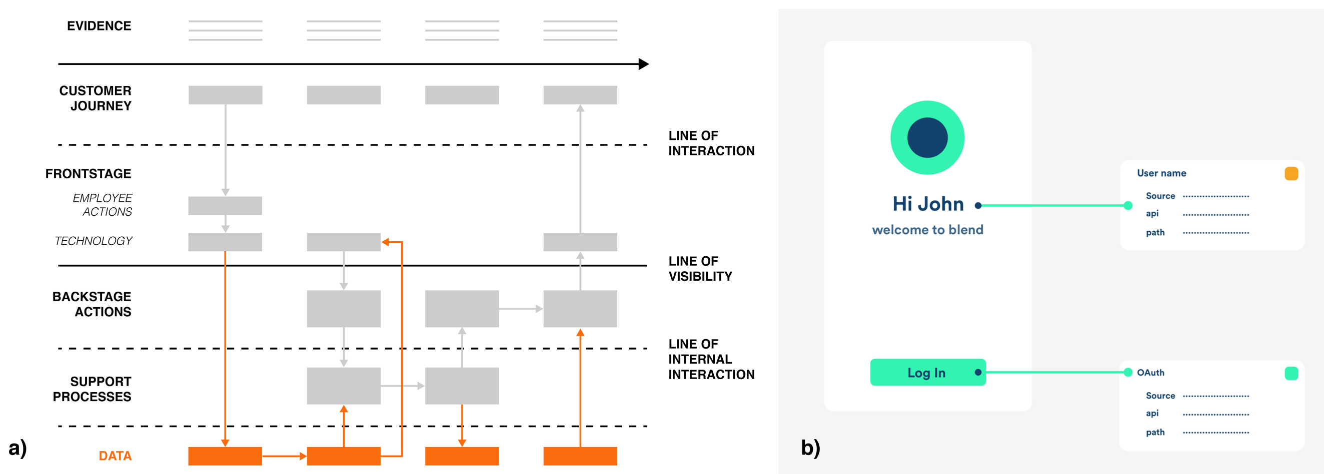

Figure 1: Augmented tools, such as (a) service blueprints with a data swim lane, and (b) annotated wireframes supported designers in understanding and communicating the role of data within their design.

ABSTRACT

This position paper makes the case for developing design tools, methods, and boundary objects for empowering designers to envision with data and AI. I describe a study that explored how experienced designers work with AI, and I introduce a new design tool –a taxonomy of AI capabilities– to support practitioners in envisioning with AI. I discuss the details of the knowledge these tools encode, and I provide discussion points to prompt future research directions.

KEYWORDS

user experience, artificial intelligence, design tools, futuring

ACM Reference Format:
Nur Yildirim. 2022. Developing Design Tools to Support Envisioning with Data and AI. In *CHI '22 Workshop Paper, June 03–05, 2022, New Orleans, LA.* ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

CHI '22 Workshop Paper, April 30–May 5, 2022, New Orleans, LA

ACM ISBN 978-1-4503-XXXX-X/18/06...$15.00
https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

My research focuses on designing human-AI interactions, and supporting design practitioners in designing for AI. In this workshop, I share two studies that reveal insights into developing future design tools that can empower design practitioners to envision with data and AI.

The first study describes how design practitioners working in innovation teams collaborated with AI experts to envision and prototype AI systems and data-driven products and services. Practitioners' reflections on their process and practices reveal that tools augmented with data, such as wireframes and service blueprints with data annotations, support designers in designing with data and AI. These tools also serve as boundary objects, facilitating collaboration between designers and AI experts. Second study describes a taxonomy of AI capabilities I developed to help designers envision buildable AI concepts. My collaborators and I created a corpus of 40 AI features that are commonly used across many products and services. We extracted high level AI capabilities, and we created communicative forms that can sensitize designers to AI capabilities.

By sharing these insights and resources, I hope to deepen the discussion around futuring design tools that effectively scaffold design ideation for data and AI-enabled products and services, and that scaffold cross-disciplinary communication and collaboration.

2 TOOLS CREATED BY PRACTITIONERS

My collaboration with cross functional AI innovation teams explored the practices of design practitioners (e.g., user experience

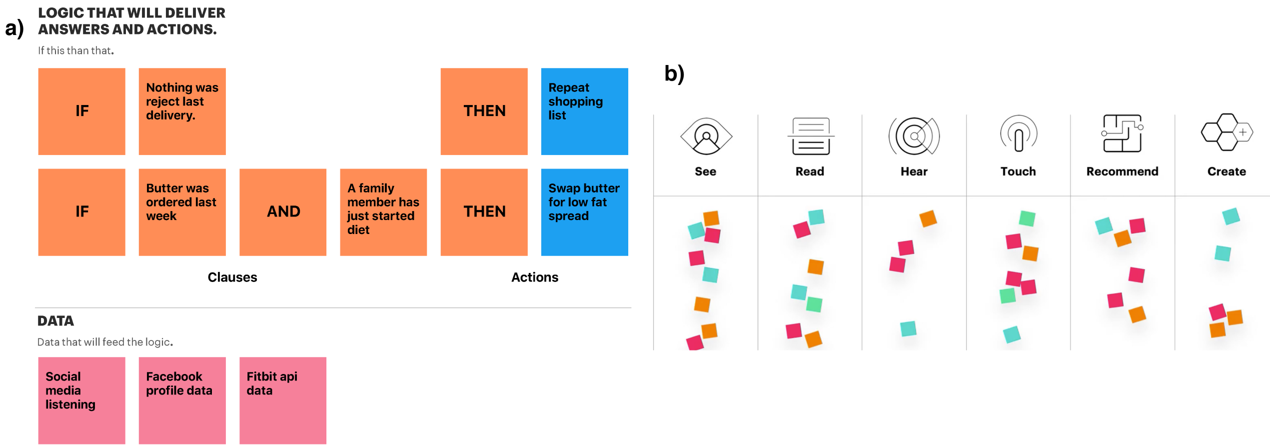

Figure 2: (a) Data-driven service design and systems design methods scaffolded designers' thinking around the AI system and its dependency on labelled data, (b) the AI Capability Matrix was used to learn about and ideate with AI capabilities.

designers, service designers) who regularly work with AI in the enterprise [4, 5]. We formed an interdisciplinary team consisting of HCI and AI researchers, design practitioners, and AI experts (e.g., data scientists, AI engineers). Our goal was to gain an understanding of how design practitioners approach designing for AI, and how they collaborate with other roles and stakeholders in AI development processes. We conducted design workshops and held discussions around current work practices. Practitioners' reflections revealed several tools and boundary objects they had developed to address the challenges of working with data and AI. Below are a few examples.

Service Blueprints with Data Swim Lane. Designers augmented service blueprints by adding a data layer as a distinct swim lane (Figure 1a). This helped them to understand and visualize the data pipeline, as it was important to *"understand which systems the data sits on, and whether data can be transferred across systems to be used together"*. Augmented service blueprints also served as boundary objects to facilitate design and data science collaboration: *"[This] made a huge difference in terms of the data scientists being able to talk through the process with the designers. Because it's really important for us where the data feeds in, so we know when we can use it for our analytics and AI."*

Wireframes with Data Annotations. Designers augmented their wireframes with data annotations to better understand and communicate data requirements (Figure 1b). Similar to the wireframe exemplar reported in [3], these served as boundary objects: *"[annotated wireframes] was designed to make our conversations with the development team a lot easier, because you're drawing this box, but where does this box pull its data from?"*

Data-driven Service Design Canvas. Designers tweaked the service model canvas to, creating what they referred to as a "data-driven service design canvas" (Figure 2a). They frequently used this tool to support ideation and team alignment around data needs. Based on the canvas, they created a set of logic statements using the structure, *"if this, then that."* These statements aided ideation, exploration, and scenario construction: *"We give people post-its where they put [if, and, then] clauses together with actions, so 'if nothing was rejected on the last delivery, then repeat shopping list'."* The canvas explicitly prompted designers to think about the AI's value proposition and required data through questions such as *"how will this service help to make people's lives better?"*, *"when is the service triggered?"* and *"what data is needed at each point?"*. This exercise helped them build elaborate and sophisticated data-driven services.

AI Capability Matrix. To specifically help improve their understanding of AI capabilities, designers created an AI Capability Matrix (Figure 2b). They translated well known AI mechanisms (e.g., NLP, computer vision) into non-technical, relatable terms, such as see, read, and hear. Along with these capability abstractions, they used exemplars to sensitize designers to what AI can do. For example, designers described a system that could *"see"* text on packaging, to then *"read"* the text it found, extracting the ingredients in order to monitor for a conflict with a known set of food allergies. They thought of these capabilities as functions that could be combined, such as *"seeing"* text and then *"reading"* any found text. These capability abstractions and exemplars enabled designers to facilitate AI ideation sessions with various stakeholders. Using action verbs instead of technical AI terms and mechanisms made ideation workshops more accessible for designers, product managers, clients, and other stakeholders that did not have AI or data science training.

3 A TAXONOMY OF AI CAPABILITIES

Prior literature revealed challenges around envisioning and prototyping with data and AI. Designers find it difficult to grasp what AI can and cannot do, and they frequently envision ideas that exceed AI's capabilities and cannot be built. This difficulty stems from designers' lack of understanding of AI's capabilities and limitations [1, 2]. Building on these challenges and the insights detailing the types of tools designers use to effectively work with AI, my collaborators and I created a taxonomy of AI capabilities.

To help ground designers in capabilities they could reasonably ask of AI and avoid ideation of things that cannot be built, our

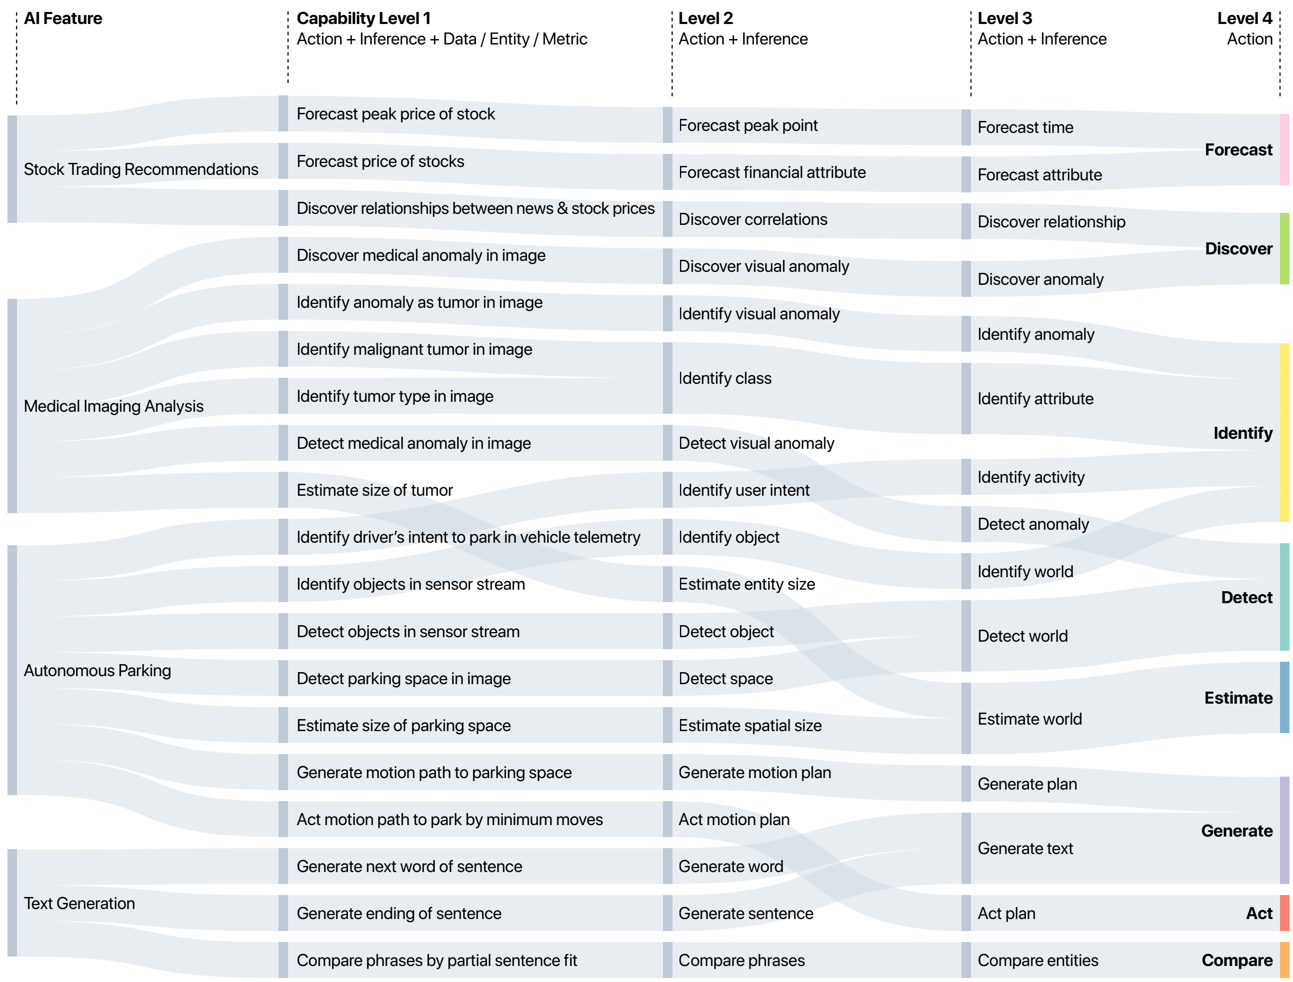

Figure 3: A small excerpt of AI capability taxonomy showing four of the AI features that breakdown into the eight Actions we identified.

taxonomy draws on a corpus of AI features frequently found in services used across different industries. We surveyed commercial AI products and services across a broad set of domains (e.g. healthcare, transportation, education) and identified 40 AI features that are commonly used. We created a taxonomy of capabilities by breaking down AI features into distinct capabilities. We then worked from small sets of capabilities to iteratively define the right levels of abstraction for the taxonomy. For example, *Medical Imaging Analysis* (Figure 3) detects a medical anomaly in an image, then identifies the anomaly as tumor, and identifies whether it is malignant or benign.

Through repeated rounds, we developed specific language for describing each capability including the action (verb), various entities, types of data, and the focus of the inference. Our taxonomy captured each AI capability on Capability Level 1, detailing the Action; Inference; and Data, Entity or Metric (Figure 3). Next, we abstracted the initial capability level (Capability Level 1) to a higher level by reducing Data, Entity, or Metric (Level 2), which would then abstract to a higher level of inference (Level 3). For example, identifying an object is ultimately about identifying a world entity. Generating text might include generating words, sentences, or phrases. All these low level inferences would abstract to "text". The final level of the taxonomy held the top-level Action for the taxonomy (Level 4).

Our efforts to encode the 40 AI features as a taxonomy produced 202 specific AI capabilities. These abstract to eight representative Actions that capture AI capabilities on a high level: *Detect*, *Identify*, *Estimate*, *Forecast*, *Compare*, *Discover*, *Generate*, and *Act*. Below, I briefly describe each capability along with some example applications.

Detect determines an entity's presence or absence within a data set or data stream. Examples: Detect human voice in audio (smart speaker); Detect face in image (camera); Detect step in motion sensor stream (smartwatch).

Identify searches for a specific item or class of items among a set or stream of similar things. Examples: Identify if message is

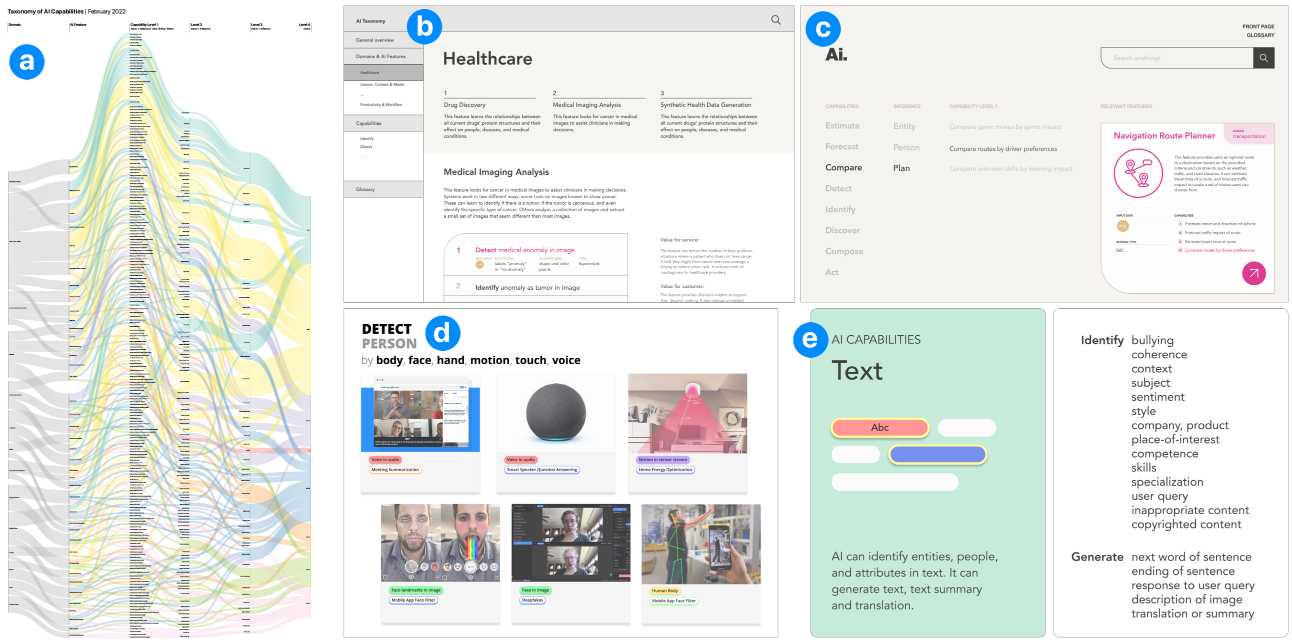

Figure 4: Sankey diagram of the taxonomy (left) allowed us to sketch various communicative forms for different use cases (right). For example, web based forms (b, c) can allow designers to browse the taxonomy corpus and each AI exemplar, and print forms (d, e) can support more directed ideation with specific capabilities or data types.

spam (email); Identify if Steve's face (security); Identify the type of cancer (medical imaging).

Estimate infers a value (e.g., position, size, duration, cost) related to the current situation. This is about making an inference about now. Examples: Estimate driving time (navigation maps); Estimate chances this is spam (email); Estimate direction sound came from (smart speaker).

Forecast infers a value (e.g., position, size, duration, cost) related to a future situation. Examples: Forecast best time to buy stock (financial planner); Forecast tomorrow's weather (weather app); Forecast max price for my house (real estate app).

Compare evaluates and organizes a collection of like items based on a metric. Examples: Compare items by likelihood of purchase (online store); Compare posts by likely engagement (social media); Compare movies by likelihood of watching (media).

Discover reveals patterns and relationships within a dataset. Examples: Discover how people use this site (usage mining); Discover unusual bank transactions (fraud detection); Discover relationships between drugs and disease (drug discovery).

Generate creates new content based on knowledge of similar content. Examples; Generate chat response (chat agent); Generate detail in image (photo retouching); Generate synthetic medical records (medical data synthesis).

Act executes a strategy to achieve a specific goal or outcome. Examples: Act: Park the car (autonomous parking); Act: Play poker (gambling agent); Act: Fly drone to location (drone pilot).

Our taxonomy captures and documents the AI capabilities within a hierarchical, extensible structure. While an effective classification system, the taxonomy is a fairly abstract artifact. It does not directly communicate these capabilities to designers. We created a sankey diagram (Figure 4a) as a way to develop, discuss, and critique the taxonomy. This provided a starting place for thinking about new communicative forms that could sensitize designers to AI's capabilities. Figure 4 presents four sketches of possible communicative forms meant to help communicate the breadth of forms the taxonomy might take to support different use cases. These include (b) a website, (c) a web-based interactive data visualization, (d) a mood board, and (e) a card deck. To illustrate how these forms might be useful, here I describe two use cases with specific design challenges.

Use case 1: Exploring specific inferences. A design team works on new concepts for an IoT-enabled smart home. They want to explore how the home might notice if someone is at home or in a specific room. They worry that homeowners might not want constant monitoring in the form of cameras. They use the AI Taxonomy moodboard posters (Figure 4d), selecting the "Detect person" moodboard. It details different ways a system might know if a person is present.

Use case 2: Exploring capabilities related to types of data. A design team works on a new service for building inspectors. It records them when they inspect a building, turning their words into text. The design team wants to explore what AI might be able to do to make the text more useful to both the inspector and to the building owner. They draw a data type card from the AI capability card deck (Figure 4e) to see what they might be able to do with text as an input.

4 PROMPTS FOR WORKSHOP DISCUSSION

The above studies showed two things. First, there is a great need for design tools, methods, and processes to aid designers in working with data and AI. Design practitioners in industry shared that they often repurpose existing tools or create bespoke tools to ease these challenges. Second, designers need boundary objects that can support them in collaborating with AI experts. To collectively envision AI innovations, they should be able to discuss the AI capability and the inference of the system, and whether the kind of data required is readily available to run a particular idea.

The following are some starting points for this discussion:

(1) **How much AI knowledge is needed for domain experts to participate in ideation?** Design practitioners worked to elicit input from domain experts to inform the design of AI systems. What kind of AI literacy is needed for domain experts to effectively participate in AI envisionment? Can a taxonomy of capabilities and value co-creation be helpful for other roles and stakeholders who do not have a technical background in AI/ML?

(2) **What do boundary objects need to ground to effectively facilitate discussion around the AI system and data requirements?** What are the qualities of an effective boundary object? What type of information is critical for collective ideation of AI systems?

(3) **Can HCI facilitation bridge the gap between AI's technical advances and human-centered AI products?** HCI routinely facilitates the process of technology innovation to reduce the risk of developing products and services nobody wants. What is uniquely difficult about facilitating AI ideation? How can HCI practices effectively span the gap across multiple roles and stakeholders in AI development processes?

REFERENCES

[1] Michael Chromik, Florian Lachner, and Andreas Butz. 2020. ML for UX?-An Inventory and Predictions on the Use of Machine Learning Techniques for UX Research. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–11.

[2] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 278–288.

[3] Qian Yang, John Zimmerman, Aaron Steinfeld, and Anthony Tomasic. 2016. Planning adaptive mobile experiences when wireframing. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 565–576.

[4] Nur Yildirim, Connor Upton, Costello Donnacha, Robert Giusti, Sinem Lacin, Sara Lovic, James O'Neill, Rudi O'Reilly Meehan, Eoin Ó Loideáin, Azzurra Pini, Medb Corcoran, Jeremiah Hayes, Diarmuid Cahalane, Gaurav Shivhare, Luigi Castoro, Giovanni Caruso, Alex Kass, Teresa Tung, Changhoon Oh, James McCann, Jodi Forlizzi, and John Zimmerman. 2022. How Experienced Designers of Enterprise Applications Engage AI as a Design Material. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.

[5] John Zimmerman, Changhoon Oh, Nur Yildirim, Alex Kass, Teresa Tung, and Jodi Forlizzi. 2020. UX designers pushing AI in the enterprise: a case for adaptive UIs. *Interactions* 28, 1 (2020), 72–77.