

Designing AI to Think with Us, Not for Us

STEFANIE HUTKA, *Sendfull LLC*

This paper establishes a framework and toolkit for designing Generative Artificial Intelligence (genAI) that address foundational challenges of these technologies and reframe the problem-solution space. The stakes are high for human-centered solutions: genAI is rapidly disrupting existing markets with technologies that exhibit increasingly complex and emergent abilities, and accelerate scale, cognitive offloading, and distributed cognition. The Problem-Solution Symbiosis framework and toolkit extends, rather than displaces, human cognition, including tools for envisioning, problem (re)framing and selection, interdisciplinary collaboration, and the alignment of stakeholder needs with the strengths of a genAI system. Applying the toolkit helps us guide the development of useful, desirable genAI by building intuition about system capabilities, developing a systemic understanding of emerging problem spaces, and using a matrix to identify if and when to offload tasks to the system. The framework is informed by systems theory, frame analysis, human-computer interaction research on current AI design approaches, and analogous approaches from spatial computing design research.

Introduction

ChatGPT, Open AI's chatbot, reached 100 million monthly active users within two months of its November 2022 launch, making it the fastest-growing consumer application in history (Hu 2023). This milestone is widely seen as a breakthrough moment for generative AI (genAI), a class of machine learning (ML) that generates text, image or sound output, based on user input of the same modalities (Google 2024). Countless companies across sectors, industries and geographies have since integrated genAI into their offerings, optimizing existing *foundation models* (i.e., models trained on a wide variety of unlabeled data that can be used for a broad range of tasks, as Murphy (2022) explains) or training new models for their use cases. Furthermore, a majority of senior executives across industries agree that genAI will “substantially disrupt” their industry over the next five years (MIT Technology Review Insights 2024).

Business leaders seek to “ride the wave” of this disruptive technology (Christensen and Bower 1995), pushing for “AI-first strategy” (Acar 2024). This push results in the application of genAI across use cases, regardless of whether the technology is well-suited to solve for the target audience's needs. Coupled with unique aspects of genAI technology that we will subsequently discuss, this push

creates new challenges for ethnography practitioners seeking to guide organizations towards building technology that is useful and desirable for people.

This paper seeks to provide ethnography practitioners with guidance on how to “ride the wave” of genAI, especially given the rapidly-evolving, emergent genAI solutions driving organizational strategy. We will begin with an examination of the unique challenges genAI presents for developing human-centered solutions and discuss genAI’s potential to perform increasingly complex cognitive tasks, concluding that genAI urgently needs ethnography to build solutions that think with us, rather than for us. We then examine how teams currently design genAI systems, and discuss challenges and shortcomings of existing approaches to problem framing for genAI projects (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023), grounded in the practice of understanding and generating solutions for existing problems.

We posit that current approaches to guide the development of human-centered technology do not address a foundational challenge. Namely, genAI is a solution that unlocks new problem spaces, whereas current approaches focus on solving for existing problem spaces. We examine the reasoning mechanisms behind human-centered practitioners’ current approaches to problem framing, and the tension between this framing and genAI’s solution-first growth. We then offer a new framework to address the foundational challenge, helping reframe the interplay between problem and solution space to align stakeholders’ needs with strengths of a genAI system. We offer three tools that follow from the framework, which practitioners can use to help teams build useful, desirable genAI solutions that think with us, not for us.

Why GenAI Needs Ethnography

There are two unique characteristics of genAI systems that simultaneously make human-centered solution development challenging, while also creating the potential for powerful – and potentially harmful – capabilities. First, with each new generation of foundation model, it becomes increasingly difficult to anticipate the system’s behaviors *a priori* without ample hands-on experimentation. This challenge is rooted in three genAI system capabilities: *scale*, *homogenization* and *emergence* (Bommasani et al. 2021, De Paula et al. 2023; see Endnote 1 for further details).

Scale enables models to ingest massive amounts of data. Homogenization enables models to adapt to multiple tasks, modalities and disciplines, obviating the need to build separate models for each function. Emergence enables models to exhibit unprecedented and unexpected capabilities. For instance, when engineers

make a small change to one part of the model to better solve certain types of problems (e.g., math problems), new behaviors emerge as a result of the system's complexity (Klein 2024). These changes are not pre-programmed or trained into the model. The system may also “perceive” patterns that are not there (i.e., hallucinations), leading to unreliable output. Hands-on experimentation with genAI models has been proposed as an approach to develop intuition about the solutions' vast and often unpredictable capabilities (e.g., Mollick 2023, Walter 2024).

Second, as genAI systems' capabilities rapidly expand, so to do the range and complexity of tasks we can perform using these tools, and by definition, the potential problems that the technology could be applied to solve. We can examine ChatGPT's evolution as a case study (Open AI 2024a). The initial foundation model, GPT-3.5, was capable of natural language understanding and generation, with the ability to maintain context within shorter conversations. GPT-4, introduced in March 2023, offered more advanced problem-solving skills and more reliable responses, could retain context over extended conversations, and included multimodal abilities (i.e., processing and generating a range of data types beyond text, such as images and audio). GPT-4o, introduced in May 2024, honed its multimodal capabilities and added real-time functionality, with latencies comparable to human conversational turn-taking (Open AI 2024b, Stivers et al. 2009). This tool is not only capable of advanced cognitive tasks (e.g., thinking through a math problem), but also the ability to carry on a conversation approaching human capabilities and adopt a personality of the user's choosing.

Looking ahead, the technology industry widely accepts that artificial general intelligence (AGI) – roughly defined as “AI systems that are generally smarter than humans” (Open AI 2023) – is inevitable (Dilmegani 2024). Industry leader Open AI states, “as our systems get closer to AGI, we are becoming increasingly cautious with the creation and deployment of our models”, indicating movement towards AGI as an end goal (Open AI 2023). Teams building genAI systems require ethnographers' guidance, not only for applying genAI where it is useful, but also extending, rather than replacing, human cognition in ways that are *desirable* to people using the technology. This requirement is especially urgent, given genAI's exponential, worldwide growth and evolution (McKinsey and Company 2023).

One could counter-argue that new technology has always enabled some amount of *cognitive offloading* (Risko and Gilbert 2016), enabling *distributed cognition* from us to our technology tools (e.g., Hutchins and Klausen 1996). For instance, outsourcing navigation in a familiar city to Google Maps, or using your phone's calculator to calculate a tip at a restaurant. We posit that genAI is already shifting us into a new realm of cognitive offloading given its increasingly advanced capabilities.

Ethnographers have the potential to help teams navigate questions about how much of our executive function we wish to distribute to a genAI system, and under what circumstances.

It is also important to highlight that the data on which genAI models are built can reflect cultural, social, ableist, gender, racial, ethnic and/or economic trauma (Sampson 2023), harming those who use genAI tools. For instance, Open AI's DALL-E 2 image generator showed people of color when prompted for images of prisoners, or exclusively white people when prompted for images of CEOs (Johnson 2022). Humans are also susceptible to *automation bias*, in which we tend to favor information from automated systems – even to the point of ignoring correct information from non-automated sources (Sampson 2023). This is particularly problematic given genAI's emergent behavior and potential for hallucination.

Best practices have begun to emerge for mitigating genAI's harm and bias risks, such as organizations establishing, *a priori*, what content the model will not generate, devising harms modeling scenarios, implementing adversarial testing, and heightened model performance monitoring. These approaches are well-documented in both the EPIC community (e.g., Sampson 2023; De Paula et al. 2023) and the broader ML community (e.g., Gallegos et al. 2023), and are therefore not a focus of the current paper.

Designing GenAI Systems: Current Approaches and Challenges

Ethnography practitioners have already been engaging with the development of genAI systems. Research on this engagement provides evidence that designing for complex, emerging ML systems like genAI presents three unique challenges (Yang et al. 2020). First, familiar user-centered design (UCD) approaches (e.g., sketching, and subsequently gathering feedback on, a low fidelity prototype of a user interaction, or conducting a “Wizard of Oz” evaluation; Klemmer 2002) have limited applicability, given the emergent nature of the genAI models' behavior. Relatedly, teams typically begin the design process with a solution in mind or in hand. If a team has the resources to train a foundation model from scratch, they are still likely seeking to “ride the wave” of genAI and know they want to explore using this specific technology for an application. More likely, the team is seeking to customize (e.g., fine-tune) an existing foundation model, meaning that there is a fully-fledged solution (i.e., foundation model) already available that the team is building from.

Early promising approaches to address these challenges include the development of prototyping tools that require minimal coding expertise, allowing practitioners

across backgrounds to experiment with AI and data (e.g., Carney et al. 2020). Teams have also reported success with regular, rapid experimentation with the technology solution, including a tight collaboration loop between product teams and engineers (e.g., ML engineers, data scientists; Yildirim and Oh et al. 2023, Walter 2024).

This experimentation approach overlaps with the second challenge – namely, genAI requires adaptations to how cross-functional teams collaborate – specifically, how product teams (e.g., ethnography practitioners, designers, product managers) and engineers building and/or customizing the foundation models (e.g., ML engineers, research scientists) work together. Members across teams have different ways of knowing (Hoy et al. 2023), including specialized language and workflows. Disconnects can quickly become amplified given the speed at which genAI technology develops.

Promising approaches to remedy these disconnects include ethnography practitioners and other product team members sensitizing ML engineers to user needs (Zdanowska and Taylor 2022) – for instance, using visuals and other *boundary objects* to help align cross-functional partners (Lee 2007, Star and Friesemer 1989), and close cross-functional collaboration between ML engineers and the product team to identify system capabilities (Yang et al. 2020). Some examples include a list of synthesized system capabilities (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023) and regular, cross-functional experimentation with an in-development system (Walter 2024).

The third unique challenge is a central tenet of this paper: genAI solutions unlock new problem spaces, standing in contrast to the human-centered playbook of solving for existing problem spaces. We have seen symptoms of this foundational challenge documented in human-computer interaction literature on how teams design AI-based solutions, manifesting in struggles practitioners face early in the design process.

Specifically, practitioners report needing support in “*getting the right design*” (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023). This phrase references Bill Buxton’s (2007) distinction between identifying that your design solution is solving a valuable problem for people, versus refining its usability after the solution and intended use case are identified (i.e., “*getting the design right*”). Existing AI design guidebooks (e.g., Google 2019, IBM 2022, Apple 2023) tend to focus on solution refinement, leaving practitioners seeking guidance on early-phase problem framing and reframing, and subsequent ideating on solutions that are both useful and technically feasible (Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023). The following section examines the reasoning mechanisms behind this foundational challenge.

Foundational Challenge: Problem-first Versus Solution-first Framing

GenAI is an example of a *disruptive* technology entering the market, displacing established markets and generating new ones (Christensen and Bower 1995). A famous past example is when Apple introduced the iPhone in 2007. This technology solution displaced mobile phone incumbents like Blackberry and Nokia, but was also the necessary precursor for mobile-first use cases and related business models to emerge – for instance, sharing photos via social media (e.g., Instagram), ridesharing (e.g., Uber) and mobile payments (e.g., Venmo). In the same way, genAI is starting to displace technologies we considered mainstream even mere months ago, and create new ways of accomplishing tasks – for instance, moving from a search bar to a prompt bar.

Disruptive technology – be it the iPhone or genAI systems – are, by definition, *solutions* that are then applied to use cases, rather than solutions developed around people’s needs. Per Paul Graham, co-founder of startup accelerator and venture capital firm Y Combinator, genAI is, “the exact opposite of a solution in search of a problem. It’s the solution to far more problems than its developers even knew existed” (Graham 2023). In other words, the solution (genAI) unlocks a number of *problem spaces* – all of the aspects related to understanding and defining a given problem people face (Simon 1969; Newell and Simon 1972).

The dynamic of a solution (genAI) unlocking new problem spaces creates a tension between human-centered design practitioners and genAI. We argue that this tension is rooted in how we approach problem framing in genAI solution development. Dorst (2011) examined the reasoning patterns that experienced human-centered design practitioners use to frame design problems. The following equation typifies a successful design solution:

$$\text{WHAT} + \text{HOW leads to VALUE}$$

In this equation, “WHAT” is an object, service, or system. “HOW” is a known working principle that will help achieve the value one aspires to offer customers (“VALUE”). The “WHAT”, in combination with the “HOW”, should therefore yield the aspired “VALUE”.

When presented with a complex design problem, experienced practitioners tend to encounter the following equation, in which the only “known” variable is the assumption about the value that they seek to achieve. The solution (“WHAT”) and the working principle (“HOW”) that achieves the value are unknown.

??? + ??? leads to VALUE

This requires abductive reasoning, working backwards from value. Here, experienced practitioners adopt a frame, shown in square brackets in the equation below. The frame implies that applying a certain working principle (“HOW”) will lead to the aspired value. The frame leaves the solution (“WHAT”) to be determined.

WHAT + [HOW leads to VALUE]

Before designing a solution, the practitioner tends to search for what Dorst calls the *central paradox* – the crux of what makes the problem difficult to solve. This is done through searching problem space (e.g., via ethnographic research, sense-making). Once the practitioner has a clearer view of the central paradox, they progress to the solution (“WHAT”).

In contrast, disruptive technology like genAI presents by *beginning* with the solution (“WHAT”), per the equation below:

WHAT + ??? leads to ???

Taking a conventional framing approach in this context will leave the practitioner with two unknown variables (“HOW”, “VALUE”), untethered from the solution (“WHAT”). This may explain what is observed when human-centered practitioners working on an AI-based project report user needs that the AI system cannot help address (Yildirim and Oh et al. 2023). The current approach is, “*What solution could unlock this value?*”, whereas we need to reframe to, “*What value could this solution unlock?*”

Adding further complexity to this reframing is that “WHAT” is difficult to define with specificity, due to genAI’s emergent properties. Teams must acknowledge this uncertainty, treating the potential problems that a genAI solution is uniquely positioned to solve as a hypothesis. That said, there is evidence that grounding a cross-disciplinary team in a general inventory of AI system capabilities at the start of an ideation session can help the group generate ideas for applications that are both useful and technically feasible (Yildirim and Oh et al. 2023). In this same study, a traditional UCD approach to brainstorming, which starts strictly from problem space, led to generating ideas that were constrained both in usefulness and technical feasibility. Even ideas that were viewed as high value to end users required

extremely high system accuracy (e.g., predicting sedation dose for ventilated patients in an intensive care unit, Yildirim and Oh et al. 2023).

This examination of problem framing mechanisms highlights the tension between the current human-centered playbook and the solution-first framing that disruptive technology dictates. The subsequent section proposes a new framework to help solve for this tension, and offer a path forward to align an understanding of a genAI system’s capabilities (i.e., solution space) and the needs of potential stakeholders of this system (i.e., problem space).

New Foundations: The Problem-Solution Symbiosis Framework

We offer the Problem-Solution Symbiosis (PSS) framework to help model the interplay between problem and solution space when building genAI solutions (Figure 1). We have previously demonstrated the dominance of solution space when working with disruptive technology, and how this can disrupt application of traditional problem framing if we assume problem and solution space are expanding in sequence, rather than in parallel. The PSS framework models this co-existence of problem and solution space. The arrows represent touchpoints at which each space can symbiotically inform and evolve the other. For instance, taking inventory of system capabilities can influence how we frame problem space (i.e., “what value could this solution unlock?” is predicated by an understanding of what the solution can do), and how subsequent work to understand problem space (e.g., ethnography) can inform development of the genAI system.




Figure 1. Problem-Solution Symbiosis (PSS) Framework. The solid purple line represents problem space. The dotted orange line represents solution space. The interplay of the arrows represents the symbiosis of problem and solution space as a team works towards launching a genAI solution. The arrows represent touchpoints at which each space can symbiotically inform and evolve the other. © Sendfull, LLC

As a function of time (i.e., going from left to right in Figure 1), the team developing the genAI solution can converge on a useful, desirable and feasible solution, mirroring traditional conceptualization of “controlled convergence” on a solution, such as those espoused by Pugh (1991) and Cross (1994). This convergence is typically marked by a product launch or other significant milestone. This paper focuses on early-stage framing and subsequent ideation, and therefore does not detail processes in later-stage or post-launch development. We encourage teams building genAI solutions to continue learning and iterating following product launch.

We posit that ethnography practitioners are well-situated to orchestrate the symbiosis of problem and solution space for genAI development. This orchestration requires the ethnographer’s expertise to build coalitions between different perspectives between cross-functional team members (Hasbrouck, Scull and DiCarlo 2016), with the goal of developing a useful, desirable genAI solution that can think with us (rather than for us). We will outline three tools the ethnography practitioner can use to facilitate this orchestration, drawing from the author’s lived experience in an analogous emerging technology space – zero-to-one spatial computing product development (Hutka 2021), systems theory, and human-computer interaction literature.

Tool 1: Build Intuition About GenAI System Capabilities

In the earlier examination of how teams currently approach genAI system design, we observed that regular, rapid experimentation with technology solutions, close collaboration between product and engineering teams, and practitioners experimenting with code-free prototyping tools were useful for understanding these emerging systems’ capabilities. This process is analogous to *intuition building* that the author has applied when leading design research for zero-to-one spatial computing launches with large cross-functional teams (e.g., Hutka 2021).

Spatial computing describes a category of technologies in which computers understand people’s contexts, such as reading hand gestures, body position, and voice, and can add digital elements into the physical environment that can be manipulated similar to real-world objects (Bar-Zeev 2023). Augmented reality (AR), which overlays digital content on the physical environment, often falls under this category. One example of a spatial computing technology for which the author led design research is Adobe Aero, an AR authoring application that enables creative professionals to build digital spatial experiences (Hutka 2021). We outline different activities (solo and with a cross-functional team) that the author has led in spatial computing design research, and offer genAI analogs that ethnography practitioners can lead.

Table 1. Intuition Building Activities Used in Spatial Computing Design Research, and Corresponding GenAI Analogs

Activity Participant(s)	Spatial Computing Activity Description	Benefit	GenAI Analog
Ethnography practitioner	Spend time using existing technology tools (e.g., computing headsets like Magic Leap and HoloLens; AR mobile applications), for tasks related to the team’s product area.	Develop intuition about the AR medium’s strengths and weaknesses in different form factors.	Spend time using foundation models, such as GPT-4o and LLaMA, and/or existing products built on these models. Enter prompts related to tasks in your team’s product area.
	Try early versions of the product (i.e., “builds”) shared by engineers.	Develop intuition about the in-development product’s strengths and weaknesses; build shared language and rapport with engineers.	Partner with engineers to understand what they are building; interact with in-development systems as early as possible.
	Learn the basics of solution-adjacent tools. For example, learn the basics of Blender, a 3D authoring tool, which people frequently use to create content that will later be used in an AR experience built in a separate application.	Develop intuition about the product ecosystem into which AR fits.	If there is existing research on potential customers’ workflows, spend time in primary tools that are likely to be used alongside, or obviated by, a potential genAI solution.
Ethnography practitioner and their cross-functional team	Human-centered playtesting sessions, in which cross-functional teams were invited to use the latest build, with specific goals and/or tasks, akin to a usability test. Outcomes and next steps are documented.	Foster cross-disciplinary dialogue and awareness around solution strengths and weaknesses. Document observations and next steps.	Conduct a playtest session where team members all try using their own prompts in a given generative AI system, as relevant to your team’s product area.

Of these activities, cross-disciplinary sessions are uniquely valuable, due to their ability to organically build bridges across teams. For example, in such a session, one

team member outside of the engineering function says, “I noticed X was happening.” An engineer in the same session responds, sharing why X may occur from a technical perspective, and takes an action item to investigate offline. This exchange not only moves development forward, but fosters exposure to different ways of knowing. Note, this activity should not be used as a substitute for adversarial testing, in which a team systematically and deliberately introduces inputs designed to test how the model behaves if exploited by bad actors.

Tool 2: Stakeholder Ecosystem Mapping to Understand Problem Space

Designing for a complex, emergent system (e.g., any genAI-based solution) involves numerous stakeholders beyond the end user, all of whom hold different values and incentives (De Paula et al. 2023). For example, if building a genAI tool for internal enterprise employees, the ecosystem may include the end user, their collaborators, managers, information technology administrators, customers of the enterprise company, and the team developing the genAI tools (e.g., the ethnography practitioner, ML engineer, legal practitioner, business leaders). When exploring what problems a given genAI solution can solve for people, taking this holistic view can help teams reflect on assumptions about intended audience and motivations for building the system, in addition to generating new opportunity areas.

In an analogous approach, the author has used such mapping to understand the dynamic nature of how designers in agencies collaborate on spatial computing projects. Primary research was conducted in the form of in-depth interviews, starting with end users (e.g., designers), to learn about existing goals, behaviors, workflows and collaborators.

Through snowball sampling, participants offered introductions to these collaborators when possible (e.g., creative directors, information technology managers, business leaders). The subsequent output of synthesis was a collaboration model, demonstrating how these stakeholders worked together, as well as the goals, values and pain points of each stakeholder. This stakeholder-sensitive approach provided the team with knowledge not only about the end user, but how upstream product adoption decisions were made. The outcome informed the go-to-market strategy, addressing upstream barriers to adoption faced by senior decision makers, as well as product requirements to meet the end users’ needs. This approach adapts the service design practice of service blueprinting (Shostack 1982), in which a team maps the interactions between all stakeholders of a system. This differs from a UCD approach, which would more narrowly focus on the end user of the system.

The ethnography practitioner can take a similar approach on genAI projects. Following primary research to learn about the target audience, the ethnography

practitioner can map stakeholders, as well as their values and incentives. We recommend extending the aforementioned example to include the team building the genAI tool to promote reflexivity. For example, if a team is looking to fine-tune an existing foundation model to more effectively address customer service interactions for their company, they would consider: the customer (i.e., the end user; values: accuracy, efficiency; incentives: quickly resolve a problem); current customer service representatives (values: being able to focus on complex cases; incentives: skill development); business owners (values: cost reduction, scalability; incentives: increased profitability); internal product team (values: building useful tools; incentives: customer engagement and return use). This mapping exercise may reveal both opportunities for the product, but also expose risks (e.g., roles that genAI technology risks replacing) and new testable hypotheses for subsequent research. This activity also serves to bring interdisciplinary team members together, with the ethnography practitioner again serving as the “bridge” between disciplines.

Tool 3: The Cognitive Offloading Matrix to Identify Which – If Any – Tasks to Offload to GenAI

We previously posed the questions: how much of our executive function do we wish to delegate to a genAI system, and under what circumstances? To this end, we offer the Cognitive Offloading Matrix (Figure 2), which builds on human-computer interaction literature. The goal of this matrix is to help ethnography practitioners consider what cognitive tasks people may want to consider – or avoid – delegating to AI. While it can be applied within the context of current genAI solution development, it aims to serve as a durable tool as genAI capabilities increase and approach AGI.

The matrix has two axes: Unique strengths of humans and genAI systems (x-axis), and desirability of a given task (y-axis). The goal is to only leverage genAI systems for tasks that genAI is better equipped to solve for relative to people, *and* is something people want to offload.

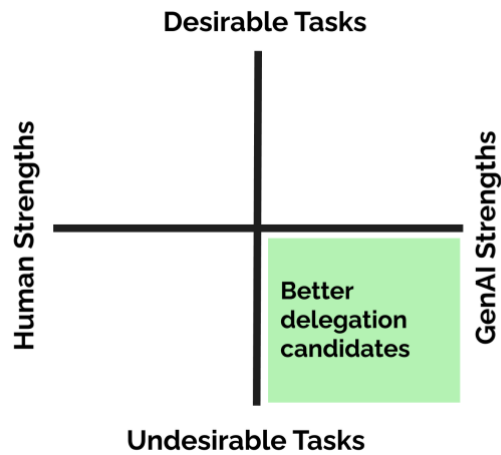


Figure 2. Cognitive Offloading Matrix. The matrix includes Strengths on the x-axis (human strengths at left, genAI strengths at right), and Tasks on the y-axis (desirable tasks at top, undesirable tasks at bottom). Tasks that fall in the bottom right quadrant (i.e., are undesirable and leverage genAI strengths) are better candidates to delegate to genAI (marked in green). © Sendfull, LLC.

To populate the x-axis, we will begin by examining a framework called the *Fitts List* (Fitts 1951), developed by psychologist and early human factors engineer, Paul Fitts. The Fitts List detailed the strengths of human versus machine capabilities. One example of human strengths is reacting to low-probability events, such as accidents (note how this activity is something at which machines, such as autonomous vehicles, remain poor). In contrast, machines excel at detecting stimuli beyond the capabilities of the human perceptual system, such as infrared detection (de Winter and Dodou 2011). Practitioners can refer to the *People + AI Guidebook* (Google 2019) for a list of “when AI is probably better” and “when AI is probably not better” as a starting point, though we caution that the lists are not specific to genAI technology solutions, with their properties of scale, homogenization and emergence.

Ethnography practitioners can adapt the Fitts List when working on genAI projects to identify potential areas where the technology can extend – rather than replace – human cognition. We can consider an example scenario in which a team is considering applying genAI to a music streaming tool. This activity can be done solo or in a workshop setting with a cross-functional team.

In this example, let us assume that a strength of a given genAI system was generating large-scale, real-time personalized recommendation. Let us also assume that a team has established (e.g., as informed by ethnographic research) that a human strength is making recommendations based on cultural nuances or non-obvious connections. These strengths can be mapped in a table, as shown in Table 2.

Table 2. Example of Fitts List for a genAI Music Streaming Tool

Human Strength(s)	GenAI Strength(s)
Making recommendations based on cultural nuances or non-obvious connections.	Generating large-scale, real-time personalized recommendations.

Populating the y-axis requires an understanding of a potential audience’s workflows. If there is no evidence on the audience’s workflows, this is cause for either engaging in primary research or treating tasks as hypotheses until otherwise investigated. This axis differentiates between desirable tasks and undesirable tasks to the potential audience. This concept is inspired by a survey investigating what activities on which people would like a robot’s help (Li et al. 2022), as part of a larger effort to develop a human-centered benchmark for robots with AI capabilities (i.e., “embodied AI”). Respondents were asked, “*How much would you benefit if a robot did this for you?*” (ibid., p.3). The highest-ranked responses were: “wash floor”, “clean bathroom” and “clean after a wild party.” The lowest-ranked responses were: “buy a ring”, “play squash” and “opening presents.” These results supported that people sought to remain engaged in pleasurable and/or meaningful tasks, and only sought to outsource laborious tasks to robots, if given the opportunity.

We posit that these results offer relevant takeaways, when considering what problems to solve with a genAI system capable of increasingly complex “cognitive” capabilities. Considering the music streaming example. Let’s recall that in this scenario, ethnographic research revealed that amongst people who used current streaming services, creating and sharing custom playlists for loved ones was an enjoyable and valuable activity. While we may have determined from our previous Fitts List that genAI is highly capable of recommendations for such a playlist, we can safely infer that this is not something people wish to outsource (Figure 3).

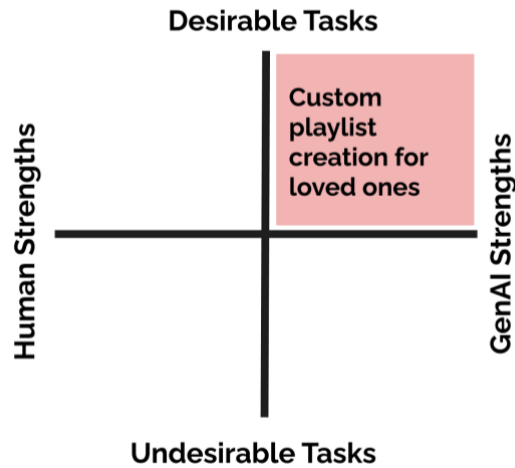


Figure 3. While ‘Custom playlist creation for loved ones’ is something genAI is highly capable of, it is also a desirable task for people. It therefore falls into the upper right quadrant of the Cognitive Offloading Matrix, and is not a good candidate for offloading to a genAI system, as indicated by the red square. © Sendfull, LLC.

Let us assume this same research study demonstrated that people want music to match their context, and that people are frustrated when the streaming service recommends the same low-energy music they want to hear when winding down in the evening, while embarking on an early-morning drive. Here, real-time generation plus contextual awareness could be something at which genAI excels. People are already open to algorithmically-generated playlists, and seek the added convenience of having music to match their context. This is a strong candidate for a “thinking” task to delegate to genAI (Figure 4).

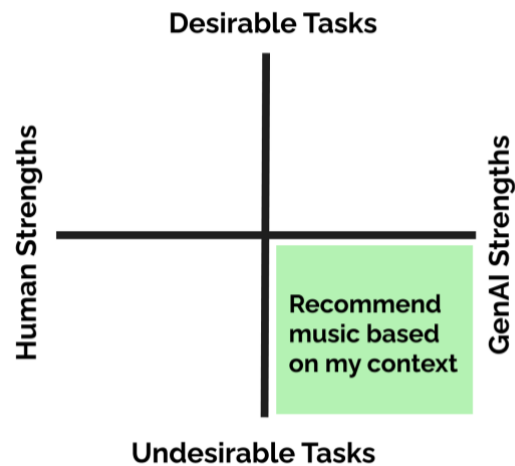


Figure 4. ‘Recommend music based on my context’ is something genAI is highly capable of, and a task people see as undesirable to perform themselves. It therefore falls into the lower right quadrant of the Cognitive Offloading Matrix, and is a good candidate for offloading to a genAI system, as indicated by the green square. © Sendfull, LLC.

General Discussion

The PSS framework proposes that ethnography practitioners move beyond a UCD approach by working *with* solution space early in the design process, and simultaneously, exploring problem space via a systemic approach to map stakeholders' values and incentives. This framework responds to observations that familiar UCD approaches often are of limited applicability in the development of genAI systems (e.g., Yildirim and Pushkarna et al. 2023; Yildirim and Oh et al. 2023), focused on how to identify and solve existing problems.

We posit that there are two sub-themes to tease apart regarding the limitations of a UCD approach. The first is that genAI is distinguished by unique characteristics (e.g., speed of advancement, complex, emergent behavior, coupled with being a disruptive technology) and therefore requires an updated approach (e.g., the PSS framework, and corresponding toolkit). The second is a more general observation that human-centered practitioners need to move beyond UCD – an observation that extends beyond the development of genAI systems. For example, Chesluk and Youngblood (2023) proposed a shift towards *user ecosystem thinking*, applying a systems-sensitive approach to industry ethnography to understand a broader range of human subjects and settings than UCD traditionally considers.

Forlizzi (2018) also advocated for moving beyond UCD to a more systemic approach. Forlizzi argues that we require stakeholder-centered design, accounting for “different entities interacting with and through products, services, and systems to achieve a desired outcome”; this is in contrast to designing “one thing for one person” (ibid., p. 1). The need for this shift beyond UCD approaches is a result of how computing technology has developed. In the earliest days of human-computer interaction, experts developed computer “programs” for themselves. Next came designing computers for others, as these machines became used in workplaces and homes. Consumer devices such as smartphones led to designing for entertainment and engagement, accelerating the development of user experience design. In the current age, we are building complex systems (genAI included) that require a systemic approach. The second tool proposed in this paper adopts this systemic view of problem space, grounded in this literature espousing a shift towards systematically building genAI solutions.

We recognize that systemic approaches can be met with friction when applied in industry settings. Product teams are often familiar with UCD language and methods, and therefore hesitant to adopt a broader view (e.g., Chesluk and Youngblood 2023). However, we are cautiously optimistic that the challenges of designing useful, desirable human-centered AI can accelerate the adoptions of systems-level thinking.

Relatedly, growing awareness of genAI “hype” (e.g., Chowdhury 2024, invoking the Gartner hype cycle [Gartner n.d.]), and the pitfalls of indiscriminately applying an “AI-first strategy” (Acar 2024) may increase business leaders’ appetite for greater specificity around which problems are best suited for a genAI solution. The ethnography practitioner remains in a unique position to move human-centered AI forward given both their toolkits and mindsets, including the ability to build bridges between disciplines and perspectives.

Conclusion

Before *generating* new frameworks to guide human-centered development of genAI solutions, it is important to examine our *foundations* – for example, how we frame problems as human-centered practitioners, and identify foundational challenges associated with genAI development. Through this examination, we identified a foundational challenge, namely that current human-centered approaches focus on solving for existing problem spaces, and do not fully address the foundational challenge genAI presents – namely, a solution unlocking new problem spaces.

To address this foundational challenge, this paper proposed the PSS framework to reframe the interplay between problem and solution space when building genAI solutions. Relatedly, we shared three corresponding tools that the ethnography practitioner is well-equipped to apply. These tools included approaches to building intuition about solution space (i.e., genAI system capabilities) and stakeholder ecosystem mapping to understand problem space, as well as the Cognitive Offloading Matrix, to help identify if and when to offload tasks to genAI. We posit this Matrix can be used as genAI systems become increasingly capable of advanced cognitive tasks. We invite ethnography practitioners to apply the framework and toolkit offered in this paper, and adapt them to their needs. There are also opportunities to explore the application of this framework and tools to other emerging technology applications beyond genAI, such as solutions built on robotics (e.g., embodied AI) and computer vision technology.

About the Author

Stefanie Hutka (stef@sendfull.com) is the founder and Head of Design Research at Sendfull. She previously led research for 0-to-1 spatial computing product launches at Adobe and Meta, and established the design research function at AR start-up, DAQRI. She holds a PhD in Psychology from the University of Toronto.

Notes

Thank you to the anonymous reviewers of this paper submission for their valuable feedback, and Juliana Saldarriaga for their support and guidance as paper curator.

1. Scale, homogenization and emergence can be traced back to two techniques proposed in the seminal ML paper, Attention is All You Need (2017). These techniques were self-supervised learning, which enabled the ingestion of billions of data sources (e.g., documents) and multi-head self-attention, in which a ML model selectively chooses which input to pay attention to, rather than attending to each input equally (De Paula et al. 2023). These techniques scaffolded future genAI solutions like ChatGPT, and enabled scale, homogenization and emergence.

References Cited

Acar, Oguz A. 2024. “Is Your AI-First Strategy Causing More Problems Than It’s Solving?” *Harvard Business Review*. Accessed May 7, 2024. <https://hbr.org/2024/03/is-your-ai-first-strategy-causing-more-problems-than-its-solving>

Apple. 2023. “Human Interface Guidelines: Machine Learning.” *Apple Developer Documentation*. Accessed May 10, 2024. <https://developer.apple.com/design/human-interface-guidelines/machine-learning>

Bar-Zeev, Avi. 2023. “What is Spatial Computing?” *Medium*. Accessed May 21, 2024. <https://avibarzeev.medium.com/what-is-spatial-computing-48d5da22b09d>

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, et al. 2021. “On the Opportunities and Risks of Foundation Models.” *arXiv preprint*, arXiv:2108.07258.

Bower, Joseph L., and Clayton M. Christensen. 1995. “Disruptive Technologies: Catching the Wave.” *Harvard Business Review*, 43-53.

Buxton, Bill. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design*. Amsterdam: Elsevier.

Carney, Michelle, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. “Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification.” *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-8.

Chesluk, Benjamin, and Mike Youngblood. 2023. “Toward an Ethnography of Friction and Ease in Complex Systems.” *2023 EPIC Proceedings*, 517-537. ISSN 1559-8918.

- Chowdhury, H. 2024. "It Looks Like It Could be the End of the AI Hype Cycle." *Business Insider*. Accessed May 20, 2024. <https://www.businessinsider.com/ai-leaders-worry-that-their-industry-is-based-on-hype-2024-4>
- Cross, N. 1994. *Engineering Design Methods – Strategies for Product Design*. Chichester, UK: Wiley.
- De Paula, Rogério, Britta Fiore-Gartland, Jofish Kaye, and Tamara Kneese. 2023. "Foundational AI: Navigating a Shifting Terrain." *EPIC Learning and Networking Week 2023*. September 18, 2023. <https://www.epicpeople.org/foundational-ai/>
- De Winter, Joost C. F., and Dimitra Dodou. 2014. "Why the Fitts List has Persisted Throughout the History of Function Allocation." *Cognition, Technology and Work*, 16: 1-11.
- Dilmegani, Cem. 2024. "When Will Singularity Happen? 1700 Expert Opinions of AGI 2024." *AI Multiple Research*. Accessed May 10, 2024. <https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/>
- Dorst, Kees. 2011. "The Core of 'Design Thinking' and Its Application." *Design Studies* 32, no. 6: 521-532.
- Fitts, Paul M. 1951. *Human Engineering for an Effective Air Navigation and Traffic Control System*. Washington, DC: National Research Council.
- Forlizzi, Jodi. 2018. "Moving Beyond User-centered Design." *Interactions* 25, no. 5: 22–23.
- Gallegos, Isabel O., Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2023. "Bias and Fairness in Large Language Models: A Survey." *Computational Linguistics* 50, no. 3. *arXiv:2309.00770v3*
- Gartner. (n.d.) "Gartner Hype Cycle." Accessed May 21, 2024. <https://www.gartner.com/en/research/methodologies/gartner-hype-cycle>
- Google. 2019. "People + AI Guidebook." Accessed May 7, 2024. pair.withgoogle.com/guidebook
- Google. 2024. "What is Machine Learning?" Accessed May 12, 2024. <https://developers.google.com/machine-learning/intro-to-ml/what-is-ml>
- Graham, Paul (@paulg). 2023. "AI is the exact opposite of a solution in search of a problem. It's the solution to far more problems than its developers even knew existed." *X*, Accessed May 7, 2024. <https://x.com/paulg/status/1689874390442561536>
- Hasbrouck, Jay, Charley Scull, and Lisa DiCarlo. 2016. "Ethnographic Thinking for Wicked Problems: Framing Systemic Challenges and Catalyzing Change." *2016 EPIC Proceedings*, 560-561. <https://www.epicpeople.org/ethnographic-thinking-wicked-problems/>

- Hoy, Tom, Iman Munire Bilal, and Zoe Liou. 2023. "Grounded Models: The Future of Sensemaking in a World of Generative AI." *2023 EPIC Proceedings*, 177-200.
<https://www.epicpeople.org/grounded-models-future-of-sensemaking-and-generative-ai/>
- Hu, Krystal. "ChatGPT Sets Record for Fastest-growing User Base – Analyst Note." *Reuters*. Accessed May 15, 2023. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Hutka, Stefanie. 2021. "Building for the Future, Together: A Model for Bringing Emerging Products to Market, Using Anticipatory Ethnography and Mixed Methods Research." *2021 EPIC Proceedings*, 59-74. ISSN 1559-8918.
- Hutchins, Edwin, and Klausen, Tove. 1996. "Distributed Cognition in an Airline Cockpit." *Cognition and Communication at Work*, 15-34.
- IBM. 2022. "Design for AI." Accessed May 15, 2024. <https://www.ibm.com/design/ai/>
- Johnson, Khari. 2022. "DALL-E 2 Creates Incredible Images – and Biased Ones You Don't See." *WIRED*. Accessed May 20, 2024. <https://www.wired.com/story/dall-e-2-ai-text-image-bias-social-media/>
- Klein, Ezra. "How Should I Be Using A.I. Right Now?." *The Ezra Klein Show*. April 2, 2024.
<https://podcasts.apple.com/us/podcast/how-should-i-be-using-a-i-right-now/id1548604447?i=1000651164959>
- Klemmer, Scott R., Anoop K. Sinha, Jack Chen, James A. Landay, Nadeem Aboobaker, and Annie Wang. 2000. "Suede: A Wizard of Oz Prototyping Tool for Speech User Interfaces." *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology (UIST '00)*, 1–10.
<https://doi.org/10.1145/354401.354406>
- Lee, Charlotte P. 2007. "Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work." *Computer Supported Cooperative Work (CSCW)* 16, no. 3: 307–339.
- Li, Chengshu, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang et al. 2024. "Behavior-1k: A Human-centered, Embodied AI Benchmark with 1,000 Everyday Activities and Realistic Simulation." *arXiv preprint*, arXiv:2403.09227.
- McKinsey and Company. 2023. "The State of AI in 2023: Generative AI's Breakout Year." Accessed February 11, 2024.
https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai%20in%202023%20generative%20ais%20breakout%20year/the-state-of-ai-in-2023-generative-ais-breakout-year_vf.pdf
- MIT Technology Review Insights. 2024. "Generative AI: Differentiating Disruptors from the Disrupted." Accessed May 27, 2024. www.technologyreview.com/2024/02/29/1089152/generative-ai-differentiating-disruptors-from-the-disrupted

Mollick, Ethan (@emollick). 2023. “Or a cutting-edge anything-and-AI expert. Spend 10 hours just playing with a Frontier Model of your choice (ChatGPT Plus if you want to spend cash, Bing Creative Mode or Claude 2 otherwise) and you will learn to do something no-one else in the world knows.” X, Accessed July 31, 2024. <https://x.com/emollick/status/1717741987846037600>

Murphy, Mike. 2022. “Design for AI.” Accessed June 26, 2024. <https://research.ibm.com/blog/what-are-foundation-models>

Newell, Allen, and Herbert Alexander Simon. 1972. *Human Problem Solving*. Prentice-Hall: Englewood Cliffs, NJ.

Open AI. 2023. “Planning for AGI and Beyond.” Accessed May 14, 2024. <https://openai.com/index/planning-for-agi-and-beyond/>

Open AI. 2024a. “ChatGPT – Release Notes.” Accessed August 5, 2024. <https://help.openai.com/en/articles/6825453-chatgpt-release-notes>

Open AI. 2024b. “Introducing GPT-4o and More Tools to ChatGPT Free Users.” Accessed May 13, 2024. <https://openai.com/index/gpt-4o-and-more-tools-to-chatgpt-free/>

Pugh, Stuart. 1991. *Total Design: Integrated Methods for Successful Product Engineering*. Wokingham, UK: Addison-Wesley Publishing Company.

Risko, Evan F., and Sam J. Gilbert. 2016. “Cognitive Offloading.” *Trends in Cognitive Sciences* 20, no. 9: 676-688.

Sampson, Ovetta. 2023. “Pulp Friction: Creating Space for Cultural Context, Values, and the Quirkiness of Humanity in AI Engagement.” *2023 EPIC Keynote*. <https://www.epicpeople.org/pulp-friction-creating-space-for-cultural-context-values-humanity-in-ai/>

Shostack, G. Lynn. 1982. “How to Design a Service.” *European Journal of Marketing* 16, no. 1: 49-63.

Simon, Herbert Alexander. 1969. *The Sciences of the Artificial*. The MIT Press: Cambridge, MA.

Star, Susan Leigh, and James R Griesemer. 1989. “Institutional Ecology, Translations’ and Boundary Objects: Amateurs and Professionals in Berkeley’s Museum of Vertebrate Zoology.” *Social Studies of Science* 19, no. 3: 387–420.

Stivers, Tanya, Nicholas J. Enfield, Penelope Brown, Christina Englert, Makoto Hayashi, Trine Heinemann, Gertie Hoymann et al. 2009. “Universals and Cultural Variation in Turn-taking in Conversation.” *Proceedings of the National Academy of Sciences* 106, no. 26: 10587-10592.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. “Attention is All You Need.” *Advances in Neural Information Processing Systems*, 1-15.

Walter, Aaron. 2024. "Adobe Design Talks on AI and the Creative Process." Design Better Podcast (note: Podcast was a recording of an in-person panel, which the author attended and asked a question of the panelists during the Question and Answer period). Accessed June 26, 2024.
<https://designbetterpodcast.com/p/bonus-ai-and-the-creative-process>

Yang, Qian, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. "Re-examining Whether, Why, and How Human-AI Interaction is Uniquely Difficult to Design." In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-13.

Yildirim, Nur, Mahima Pushkarna, Nitesh Goyal, Martin Wattenberg, and Fernanda Viégas. 2023. "Investigating How Practitioners Use Human-AI Guidelines: A Case Study on the People + AI Guidebook." In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1-13.

Yildirim, Nur, Changhoon Oh, Deniz Sayar, Kayla Brand, Supritha Challa, Violet Turri, Nina Crosby Walton et al. 2023. "Creating Design Resources to Scaffold the Ideation of AI Concepts." In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, 2326-2346.