A Diffractive Analysis of GenAI through Perspectives on Productivity

Ian P. Swift
University of Illinois at Chicago
Chicago, USA

Debaleena Chattopadhyay
University of Illinois at Chicago
Chicago, USA

Abstract

Current conversations on productivity with Generative AI (GenAI) typically reduce productivity to inputs and outputs, overlooking the nuance of productivity. Simultaneously, existing research explores a variety of interactions with GenAI. To better understand these interactions, we apply a diffractive analysis of three perspectives on productivity beyond efficiency: Herbert Simon's procedural rationality, Hannah Arendt's work and worldliness, and Ivan Illich's conviviality, and analyze the different interactions with GenAI. In the discussion of these analyses, we explore how the observed differences can motivate choices between interactions, such as exploration for problem-solving versus inspiration for competence.

CCS Concepts

• **Human-centered computing → HCI theory, concepts and models**; *HCI design and evaluation methods*; • **Computing methodologies →** *Machine learning*.

Keywords

critical theory, productivity, generative ai, diffractive analysis

ACM Reference Format:
Ian P. Swift and Debaleena Chattopadhyay. 2026. A Diffractive Analysis of GenAI through Perspectives on Productivity. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '26), June 22–25, 2026, Linz, Austria.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3808045.3808074

1 Introduction

This paper addresses how to use GenAI to improve productivity when interactions are diverse, and productivity is defined broadly. We provide three analyses that consider different types of interactions with GenAI through three theoretical perspectives: Herbert Simon's concept of "procedural rationality", that focuses on the faculties of problem-solving as a necessary precursor to production, Hannah Arendt's concept of "work" and "worldliness", where humans are tasked with producing an artifice of durable things, and Ivan Illich's concept of "conviviality" that challenges industrial society to promote an individual's creative autonomy.

The motivation for this study is both a cynicism of the dominating definition of productivity, in which GenAI would be expected to increase the production of outputs (such as useful generated materials) with reduced use of inputs (such as labor), as well as a belief that different types of interactions with GenAI cannot be assumed to shape production in the same way. We believe that such an oversimplification of GenAI is a barrier to meaningful discussion of how to better motivate interactions. Further, the simplification results in conversations either focusing on whether GenAI can perform a given task or on debates about the value of automation in general. We believe GenAI warrants a broader discussion of how it can be used in different ways to meet a broader range of design decisions, enabling exploration of new ways of building with GenAI.

As such, this work presents an alternative to the conventional view of productivity in Taylorism and "scientific management" [53]. Taylor took advantage of both the Enlightenment's elevation of scientific thought and the growing industrial economy to define an efficient production process as one in which tasks are precomputed, standardized, and workers are assigned to reusable roles. However, not all instances of GenAI need to serve this purpose.

Analyzing alternative understandings of GenAI requires some analytical nuance that we aim to provide in this paper. For example, consider the interaction in which GenAI is given a series of descriptions and synthesizes a new output from them. In this work, we focus on the more nuanced and philosophical questions of "What mode of problem solving does such an interaction represent?" "Does the user create something meaningful through this interaction?" and "Does this interaction represent a user's autonomy?" We argue that these broader questions of production are necessary for a meaningful understanding of productivity with GenAI and serve to enable a more deliberate design of technology.

The methods used in this paper consist of constructing an analytical framework of 12 GenAI interactions derived from a close reading of 28 user studies. Then, we perform three analyses that map the interactions to aspects of the three theoretical perspectives. In a diffractive analysis, we investigate how our analyses show the different interactions and theories engaging with one another. Finally, we discuss these analyses to demonstrate their utility for critiquing choices in interaction design.

The paper is laid out as follows. Section 2 performs a review of related works in GenAI and explores the theoretical perspectives. Section 3 describes our methodology for constructing the analytical framework to be used throughout the subsequent analyses. Section 5 presents our three analyses: procedural rationality (Section 5.1), work and worldliness (Section 5.2), and conviviality (Section 5.3). After providing the results, we use diffractive analysis to discuss the design choices in Section 6, and conclude in Section 7.

CHIWORK '26, Linz, Austria

ACM ISBN 979-8-4007-2429-9/2026/06
https://doi.org/10.1145/3808045.3808074

2 Related Works

2.1 Generative AI

GenAI is a broad machine-learning term for models capable of generating content. The capabilities of these models include general-purpose "foundation models" [5], such as GPT [42] and DALL-E [43], as well as models trained on specific tasks [18]. General purpose foundation models often undergo prompt engineering [6, 59], fine-tuning [22, 39], or both to be set to a particular task [5]. The task-specific diversity prompted us to consider different interaction types and definitions of productivity.

Two trends are notable: the market assessment of GenAI and the related research focus. For the market assessment, in 2023, GPT-4 and GPT-4-powered technologies were reported to affect 47% to 56% of all tasks performed by US workers, and could reduce the time required to complete a task by at least half [16]. By 2025, 28% of Americans who were asked had previously used ChatGPT for work [47]. According to industry leaders, GenAI meets more of their productivity and efficiency goals, but falls short of expectations for decision-making, client relationships, and innovation [14]. Simultaneously, AI high performers are more likely to set additional growth and innovation goals, while other respondents typically focus solely on efficiency [52]. Accordingly, the popular trend in market discourse follows the efficiency narrative, in which GenAI exists to reduce inputs in the production process.

While market assessment focuses on efficiency, researchers have critically examined GenAI's efficiency capabilities. GenAI automation introduces new challenges for evaluating outputs and adapting them for use [48]. Users with greater confidence in GenAI to perform a task reported a decreased perceived enactment of critical thinking [32]. Researchers found that users creating an LLM-based ChatBot struggled to construct meaningful instructions [62]. Stories generated from ideas provided by GenAI exhibited greater similarity across different writers [15]. Use of GenAI to generate ideas led to increased fixation on the original design's features [55]. These studies offer a nuanced understanding of concerns around efficiency with GenAI.

2.2 Theoretical Background

2.2.1 Procedural Rationality. Procedural rationality is defined as a rationality of "appropriate deliberation" [50] of an individual with "bounded rationality" (by the limits of human cognitive processes [49]). Procedural rationality also describes the mechanisms an individual uses to engage in problem-solving, such as searching and the use of heuristics [51]. Drawing on procedural rationality, researchers also proposed a behavioral theory of firms, including principles such as "organizations avoid uncertainty" and "search that is stimulated by a problem" [13]. The architecture within which procedural rationality operates was described through a system called Soar, an implementation of a unified theory of cognition [37]. More recently, researchers have begun developing systems that compensate for the limits of human procedural rationality [33].

2.2.2 Work and Worldliness. Arendt defines work in relation to labor in her *Vita Activa*, a hierarchy of human activities where, historically, labor (the "metabolic" aspect of production) is positioned below work (the creator of the "human artifice"), and both are positioned below action (the condition of "political life") [2]. Arendt's distinction between work and labor originates in Locke's quote, "the labour of his body, and the work of his hands" [34]; Arendt similarly uses craftsmanship to describe work, and relation to the body to describe labor [2]. Similarly, the concept of worldliness is rooted in Heidegger's *Weltlichkeit* ("worldhood"), in which humans are fundamentally in the world rather than outside observers [21]; related to how Arendt claims that work provides humans with a world in which they can live [2]. A few recent works have built upon this framework, including using the *vita activa* to analyze the future of work [56], or how it can be used to understand human activities [27].

2.2.3 Conviviality. The concept of "conviviality" represents "a modern society of responsibly limited tools", in response to Illich's observation of a pattern in which modern industries, once they reach a certain scale, begin to harm the people they were meant to serve [25]. Convivial tools can be controlled rather than controlling their user (by creating a system that requires their use). Illich developed these ideas in part through conversations at the Centro Intercultural de Documentación (CIDOC), building on earlier critiques of institutional systems such as schooling [24]. Related critiques of industrialism appeared in subsequent decades, such as Gorz's politics of liberation from work [19], and conviviality has more recently been cited as a foundational contribution to the degrowth movement [28]. In HCI, Fischer and Lemke proposed convivial computing as a system that provides users with multi-level control over their tools [17]. Similarly, Cibin et al. have used the concept of conviviality to bridge participatory design and critical media studies [12].

3 Methodology

Our methodology comprises two parts: an analytical framework and a diffractive reading informed by it. The analytical framework is derived from a review of designs and user studies. The diffractive analysis draws on three theoretical perspectives on productivity to generate new patterns within the analytical framework. The goal of analyzing interactions in this way is to develop new understandings of the different types of experiences with GenAI as a useful critical framework for design.

3.1 Analytical Framework

We developed the analytical framework by reviewing users' experiences and interactions with GenAI. We used peer-reviewed research to derive the interactions. As criteria for the types of peer-reviewed research, we considered studies that targeted some form of GenAI technology and included a user study in which users were tasked with a production task.

3.1.1 Paper Search. To aid our search, we used IEEE Xplore and the ACM Digital Library to query for papers published from 2015 to 2024. We used three concepts to perform the search: GenAI, productivity, and user studies. GenAI was operationalized with terms such as **"GenAI", "generative AI", "LLM", "foundation model", "diffusion", "GPT"**, etc., productivity was operationalized using the terms **"productivity", "efficiency", "performance", "effectiveness", "capacity"**, etc., and the presence of user studies was used in the search using the terms, **"user study", "user evaluation", "user testing", "interview", "focus group", "human subject", "user feedback"**, etc. After this initial search, we had n=16,554 results.

We first used the Zotero[1] citation manager to filter the search results. We used built-in queries to filter out citations that were from non-primary sources (exclude "survey" and "review" from the title), were listed as conference or journal, excluded "workshop" from the title, and checked our user study operationalization for its presence somewhere within the abstract. After this filtering, we had n=4,346 remaining studies.

To further identify papers related to producing with GenAI, we subjected the papers to several rounds of screening. The first screening was for titles that were off-topic (e.g., health diagnostics, or robotics), resulting in n=481 abstracts for screening. Abstracts were reviewed against three criteria: does the paper focus on GenAI, is GenAI integrated into a user workflow, and is the goal to produce an output, yielding n=111 papers. After a skim of these papers, we identified 32 papers that we believed could be relevant to our study, including two papers added from outside the search that we were aware of but that the search had not produced. After reading and analyzing the papers, 4 papers were excluded because they did not aid in understanding productive interactions with GenAI, resulting in **n=28 final papers**.

Among the topics of production are writing [10, 11, 45, 61], scene creation [20, 38], creating design artifacts [7, 9], music composing [35], professional content creation [29, 30, 40], working with technology [1, 3, 8, 23, 36, 41], and others.

3.1.2 Framework Construction. Our analytical framework, produced from reviewing the user studies, is structured from four thematic experiences: intent, synthesis, relief, and understanding. Figure 1 demonstrates the relationship between these experiences. In particular, we selected these experiences because of their prevalent role in shaping the AI production process. The first two categories, intent and understanding, focus on ideas: intent involves the user projecting an idea onto the system, whereas understanding involves ideas provided by the system. The categories of relief and synthesis focus on outcomes: relief on processes that the system replaces, and synthesis on outputs produced by the system.

After this initial analysis, we reexamined the papers to see where the reported experiences of the users agreed and disagreed. While these four experiences address distinct ideas, we found overlap and disagreement during data recoding. Accordingly, we further differentiated findings from the user studies by considering each experience in light of the others; we considered how specific interactions could be defined as one experience occurring through another. For example, when considering intent through understanding, we found that users can develop their intent by exploring the machine's ideas. There were also differences in the position of an experience: relieving a process through a synthesis meant a complete substitution of the workflow, whereas producing a synthesis by relieving a process involved an active construction of the user, but with key elements of production changed. This analysis yielded 12 interactions, identified by repeatedly examining structural relationships and findings from the user studies. The different definitions provide the necessary language to describe user interactions with GenAI. The full analytical framework is presented in Table 1.

Notably, the interactions *define the relationship with GenAI*, rather than strictly differentiating it from other forms of production. For example, direction interactions specify how a system should perform a task, a mechanism commonly used in computer programming. With GenAI, natural-language interfaces reduce the need to memorize syntax and write structured code. As such, direction interactions with GenAI are significantly different from writing code by hand. Similarly, exploration interactions are graph traversals, a long-standing representation in computer science. However, the generative aspect enables users to explore novel content beyond the system's provided inputs. As such, the interactions represent new developments with GenAI, even if they are, at times, familiar.

[1] https://www.zotero.org/

3.2 Diffractive Analysis

Using the analytical framework as a descriptive and analytical tool, we then performed three theoretical analyses and a subsequent diffractive analysis of productivity with GenAI. We base our method of diffractive reading on prior work in HCI [31, 54], which defines it as reading the same text through multiple theories to identify new observations that emerge when the results are stacked. Diffractive readings both enable theories, in combination, to produce insights beyond those of any single theory, and also represent an onto-epistemological commitment to showing how the differences matter [4]. The three perspectives used were Herbert Simon's procedural rationality, Hannah Arendt's work and worldliness, and Ivan Illich's conviviality. Procedural rationality is how an individual with limited cognition engages in problem-solving. Work and worldliness are the creation of a human artifice that stands outside of human subjectivity. Conviviality is an opposition to industrialism that reprioritizes individual autonomy and creativity. We chose these three perspectives for their distinct relationships to productivity and their distinction from efficiency measures, to provide a meaningful and critical diffractive analysis.

The analyses involve a close reading of the three authors to identify key aspects of production in their theoretical perspectives. Using these aspects and the twelve interactions, we determined if and how each interaction contributes to each aspect. Finally, we considered interactions among all aspects and theoretical frameworks simultaneously to identify new patterns of difference.

4 Note on Perspective

It has been noted in other works that the cost of GenAI is high [26], and is an aspect of industrial production. Illich, in particular, would be critical of such an industry; we welcome this critique alongside ours. However, rather than addressing the issue of GenAI production, we read Illich as addressing the emerging industrial use of GenAI. This is a limitation in perspective, but it affords us the opportunity to analyze in more detail how GenAI is used as an instrument. We hope this analysis of aspects of GenAI production would promote his ideas in limiting further industrial expansion, even if it does not reverse existing expansion.

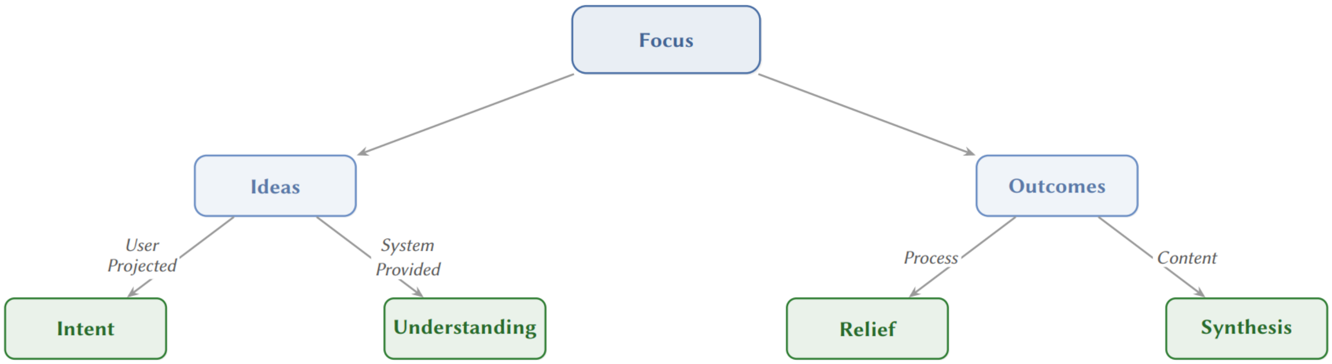

Figure 1: A decision tree to distinguish between different experiences with GenAI. On the left are experiences that focus on ideas, either projected by the user (Intent) or provided by the system (Understanding). On the right are particular outcomes, either replacing a process (Relief) or producing a particular output (Synthesis).

Another key limitation is our focus on positive interactions. We examined how GenAI can improve problem-solving, increase worldliness, and be more convivial. However, user interactions could degrade, for example, if a user repeatedly accepts auto-completions from a text generator without reading them, thereby taking a substantively rational role and ignoring the value of external memory provided by the GenAI. While this is always a risk, we primarily focused on best-case scenarios; interactions that would encourage these values more strongly, especially in cases where other interactions would be problematic. Accordingly, this paper should be understood as interactions that advocate these values, not as a guarantee of them.

5 Analyses

5.1 Procedural Rationality

While rationalism is inherent to the production philosophy of scientific management, Herbert Simon's rationalism focuses on a reinterpretation of traditional economics, claiming that the textbook definition of economic rationality "aims at maximizing profit, and in such simple circumstances that the computational ability to find the maximum is not in question," leading him to propose the required resolution where "the theory of the firm becomes a theory of estimation under uncertainty and a theory of computation" [51]. Accordingly, Herbert Simon emphasizes a rationality that is focused on how human beings (and artificial systems with similar capabilities) solve problems.

Inherent to a rationality focused on human problem solving are the characteristics of humans' both internal and external adaptations to the task. Simon says, "Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves" [51]; the processing systems that we use to compute logical problems rely on a limited short-term memory, and the creation of a variety of strategies to be invoked and adjusted to the task.

The question of how GenAI is to be understood in this procedurally rationalist perspective of problem-solving, we borrow three perspectives from Simon: 1) how GenAI can introduce new elements to long-term memory, as well as how it can utilize it and replace it, 2) how GenAI can enrich the search through a problem space through ideation and guidance, or create new avenues the user can search through, and 3) how GenAI usage relates to the changing the current representation of a problem, or providing entirely new ways of representing a problem.

5.1.1 Long-Term Memory. Given the computational limits of procedural rationality (i.e., limited observable information and short-term memory), problem-solving relies on the retrieval of previously retained data and procedures. Simon said it was a dependence upon "the cocoon of information, stored in books and in long-term memory" [51], which would certainly be updated to presently consider personal computers, search engines, and knowledge embedded in machine learning models such as GenAI. This can be demonstrated both in how GenAI contributes to and builds upon an internalized long-term memory and in how it serves as an external mechanism to replace reliance on one's own memory.

The information presented by GenAI can be communicated in a particular way to encourage long-term memory by enabling the user's internalization. One example is demonstration interactions, in which GenAI is able to associate a specific context (e.g., a bug, a musical phrase) with a related content (e.g., steps for debugging, corresponding musical notes), which it provides to the user that users of the system claimed enabled learning (e.g., how to debug, music theory) [3]. A similar interaction is differentiation, in which the GenAI uses its element's position in the foundational semantic space to demonstrate how elements relate, giving the users insight both into the differences between components, and the GenAI's part in interpreting those differences [11]. In both cases, GenAI communicates new information to the user (i.e., learning processes, theory, or differences) that the user can internalize and use later in problem-solving.

Conversely, GenAI systems can serve as an external memory store, replacing aspects of human recollection by leveraging their own access to information. For example, through indication interactions with GenAI, or by using GenAI to produce information based on a given context (e.g., identifying relevant machine learning

Table 1: The Analytical Framework

Interaction		Experience		Definition	Example
Concretization	=	*Synthesis*	*through Intent*	*GenAI produces an output from a description of its qualities*	*Users describe what operator they are looking for, and LowCoderNL configures it [44]*
Direction	=	*Relief*	*through Intent*	*GenAI follows a specific structure provided by the user*	*Users provide hierarchical prompts, and CoLadder follows the instructions [60]*
Demonstration	=	*Understanding*	*through Intent*	*GenAI expands upon user expressions with domain knowledge*	*Users compose a musical voice, and Cococo shows how corresponding voices would appear [35]*
Indication	=	*Intent*	*through Synthesis*	*GenAI informs the user through explanations of the task context*	*Users request information on code they are modifying, and GILT provides them with contextualized explanations [36]*
Substitution	=	*Relief*	*through Synthesis*	*GenAI completes the steps of a task, producing a near-finished output*	*Users upload a video, and PodReels transforms it into a video teaser [57]*
Extraction	=	*Understanding*	*through Synthesis*	*GenAI highlights specific concepts within a provided input*	*Users upload a press release, and AngleKindling provides the main points [40]*
Encapsulation	=	*Intent*	*through Relief*	*GenAI interactions are captured as reusable components*	*Users enter a prompt, and ABScribe makes it available for later re-use [45]*
Aggregation	=	*Synthesis*	*through Relief*	*GenAI generates content that the user otherwise would have to look up*	*Users write code, and a GenAI auto-complete fills in the details [58]*
Differentiation	=	*Understanding*	*through Relief*	*GenAI maps the relationships between multiple inputs or outputs*	*Users generate new characters, and Patchview places them on a semantic axis [11]*
Exploration	=	*Intent*	*through Understanding*	*GenAI iteratively suggests related elements that the user investigates for ideas*	*Users select from Kansei words, and AutoSpark repeatedly suggests more related words [7]*
Combination	=	*Synthesis*	*through Understanding*	*GenAI combines multiple inputs into a single output*	*Users select keywords generated earlier, and CreativeConnect combines them into a sketch [9]*
Inspiration	=	*Relief*	*through Understanding*	*GenAI generates content that the user reinterprets for their own work*	*TaleBrush generates stories, from which users draw characters, settings, etc. [10]*

models and parameters the user can use [8]), GenAI can generate meaningful information for the user's current task. Similarly, in aggregation interactions, GenAI generates the necessary material for the user to work with (e.g., providing information about the code they are currently writing). Similarly, in substitution interactions, the generative system produces content independently of the user, relying only on its own recall. In all of these instances, GenAI's ability to store information in long-term memory enhances overall problem-solving performance.

5.1.2 Search. Procedural rationality is largely defined by the rational processes by which an individual accomplishes a task. For example, Simon refers to the methods we use to enact judgment as "mainly a non-numerical heuristic search that draws upon information stored in large expert memories" [51]. While Simon goes on to say that AI at the time of writing was employed as "expert systems in a growing range of domains previously reserved for human expertise and judgment" [51], GenAI expands the support of rational search through both the ability to navigate a decision space, as well as the ability to generate new information sources that the user can search through.

Procedurally rational searches can take the form of a production process that aims to find a satisfactory configuration of components. One example of a productive search is determining which code to modify to eliminate a program bug [3]. GenAI can support code modification through indication interactions, such as using debugging context to generate a notification to the user indicating where else relevant code is located [3]. Alternatively, an individual may attempt to construct a structured output in which aggregation interactions yield the result of the next component, thereby eliminating sub-searches within the task. An example of GenAI aiding the solution-search is in using the system to generate code without checking documentation or internet forums for syntactic or procedural details [58].

The search process can be done in collaboration with GenAI. In this form, the GenAI produces ideational elements describing data (as in an extraction interaction), produces data that the user reinterprets and uses freely (as in an inspiration interaction), searches through suggestions to find new unexplored ideas (as in an exploration interaction), or combines ideational elements to create a new synthesized idea (as in a combination interaction). While these do not directly produce the search output, they do succeed in leveraging both the individual and GenAI to identify elements to search over, enabling the search to be conducted within an expanded domain.

5.1.3 Representation. In the context of procedural rationality, solving a problem inherently involves choosing an appropriate representation in which to explore. Simon says of representations that, "must begin with creating a representation for the problem... in which the search for the solution can take place", to the extent that he concludes that problem solving, "simply means representing it so as to make the solution transparent" [51]. In terms of representation, GenAI takes two forms. First, it transforms inputs from one form to another, and second, it provides new representations for approaching problems.

The transformation from one form to another is inherently useful in quickly alternating to a new representation. For example, an existing text could be transformed through input and generation into a set of summaries [40] or a set of design features could be generated from an image [7], both through an extraction interaction. Additionally, generating a product from a description, as in a concretization interaction, or a synthesis of multiple elements, as in a combination interaction, can create new representations of existing ideas. Generated outputs can also be displayed in their inherent positions in dimensional space, with differentiation interactions represented by their relative positions.

Additionally, the problem representation can be structured in new ways using the GenAI interactions as novel representations. For example, the natural language instruction of a direction interaction, such as the hierarchical prompt used to generate code in CoLadder [60], can be used to write code in a new way. Similarly, exploration interactions provide a new way to explore novel content, enabling the machine to suggest novel content alongside the user. Finally, encapsulation interactions use prompts as building blocks to transform outputs, such as the reusable, stackable atomic prompts in ABScribe [45].

5.2 Work and Worldliness

Contrary to the minimization of inputs and the maximization of outputs typically associated with scientific management, Arendt's ontological categorization of work focuses on whether the goal is to produce durable outputs that collectively form the world we live in. Arendt says that the creation of such things "is meant to outlast and transcend them all. The human condition of work is worldliness" [2]. Arendt contrasts this with a society of consumption and labor which "dazzled by the abundance of its growing fertility and caught in the smooth functioning of a never-ending process, would no longer be able to recognize its own futility-the futility of a life which 'does not fix or realize itself in any permanent subject which endures after [its] labour is past'" [2].

Arendt distinguishes between work and labor in the way the process of completion is achieved. Work is defined by a model that, in relation to the work, "not only precedes it, but does not disappear with the finished product, which it survives intact"; the process of creation in this sense is one of reification, destroying the natural world in order to bring into existence an idea [2]. These reified things, which are added to the world, differ from an act of labor such as tilling the soil in that, "soil, if it is to remain cultivated, needs to be labored upon time and again... it needs to be reproduced again and again in order to remain within the human world at all" [2]. Distinct from labor, which can be understood as a process of consumption, work is defined by its worldliness.

Certain key aspects of the work are evident in interactions with GenAI. Borrowing from Arendt, we note that work with GenAI can: 1) transform an idea into an output, as well as produce outputs from GenAI's own ideas, 2) involve the creation of outputs that are designed for or otherwise capable of persistence beyond use, 3) result in finalized "ends" as well as provide new "means".

5.2.1 Idea. A worldly output can be identified, in part, by its underlying idea, which guided the process of creation and persists beyond the end of production. Arendt calls such an idea a "model" that "can be an image beheld by the eye of the mind or a blueprint in which the image has already found a tentative materialization through work" [2]. Arendt speaks of this idea using a bed as an example, "we neither conceive of making a bed without first having some image, some 'idea' of a bed before our inner eye, nor can imagine a bed without having recourse to some visual experience of a real thing" [2]. GenAI can be used both as an implement for conveying a human idea and as a source of ideas in its own right.

Certain interaction modes with GenAI allow the user to transform an idea into an intentional output. For example, in concretization interactions, humans provide details about the intended output, and GenAI will generate an output following "what" they are aiming to produce (as one study notes in particular [44]). Conversely, through the encapsulation interaction, particular intentions with GenAI ("make more professional" [45]) are transformed into a reusable tool. In both interactions, the idea is transformed into an output.

Just as a blueprint serves as a source of an idea, an idea for what to produce can be sourced within the GenAI. For example, indication interactions direct the user in producing a particular output (e.g., guiding a user toward a code modification [3]), and provide the underlying ideas that appear in the final product (e.g., a particular resolution to the coding problem). Similarly, aggregation interactions produce material using information that the generative system has learned from or otherwise been provided with internally, such as solutions to a particular problem that are reused in the generation of outputs.

5.2.2 *Durability.*

The human artifice comprises a world of durable things. For a product to be durable, it must be that "their proper use does not cause them to disappear," with Arendt noting the contrast between durable and consumer goods in that, "what distinguishes the most flimsy pair of shoes from mere consumer goods is that they do not spoil if I do not wear them, that they have an independence of their own, however modest" [2]. GenAI both creates products intended to survive the interaction of their creation and outputs that, while solid, are neither intentionally preserved nor consumed.

Interactions with GenAI are durable when they are intended to be built upon after the initial interaction. For example, encapsulation interactions involve creating reusable prompts that the user maintains and applies repeatedly to their work over time. Similarly, differentiation interactions reverse the generation process to create a semantic position, thereby enabling continued contrast and comparison as more elements are added alongside one another. Notably, in line with Arendt's observation that durable goods exist outside people and enable them to relate [2], one study found that differentiation could be used to discuss elements among co-writers [11].

Additionally, GenAI can produce outputs that satisfy the condition of producing things in the world, where the particular content produced and its use determine its durability. For example, concretization interactions could generate an output intended for sharing, such as a stylized graph [46], or a combination interaction could be the intentional final product of a designed image, such as in the merging of branches in Spellburst [1]. However, unlike encapsulation and differentiation, it is not inherent in these interactions that the output be reused over time.

5.2.3 *Means and Ends.*

Worldly products can be separated from the products of labor in that it "is an end product in the twofold sense that the production process comes to an end in it ('the process disappears in the product,' as Marx said) and that it is only a means to produce this end" [2]. Labor does not have any such end, and Arendt says that "the end of the process is not determined by the end product but rather by the exhaustion of labor power" [2]. GenAI can contribute to this relationship between means and ends in two ways: either as a process that terminates to achieve a particular goal, or as a means by which the user can achieve a particular end.

Interactions with GenAI produce an end when the interaction's purpose yields a finalized product. One example is concretization, in which GenAI generates a final concrete output from the user's input, or, similarly, combination interactions, which produce a combined output from elements of the work process. Additionally, aggregation interactions function similarly: the user begins with a working context, and the GenAI generates a final product to fill the gap. In these cases, the interactions culminate in a final state.

GenAI can also provide interactions that enable the user to produce in new ways (a means), but do not prescribe an immediate particular output (an end). One example is a direction interaction, in which the user explains a process to the GenAI, with the goal of having the GenAI generate according to the user-set limits. Similarly, encapsulation interactions involve using reusable prompts as individual actions, providing the user with a set of tools to generate material.

5.3 Conviviality

While the measure of efficiency is a hallmark of the industrial mode of production, Illich criticizes this measure to the extent that "the effects of compulsive efficiency do more damage than good to most people in our generation" [25]. In response to the threats posed by industrialization, conviviality represents an inversion of society that prioritizes individuals and their creative autonomy, thereby allowing them to be responsible for themselves and their own production. With respect to production, Illich focuses on tools that enable conviviality, which he claims "enhances the ability of people to pursue their own goals in their unique way." [25].

According to Illich, "Tools foster conviviality to the extent to which they can be easily used, by anybody, as often or as seldom as desired, for the accomplishment of a purpose chosen by the user," as they are "enhancing each person's range of freedom... [and] make the most of the energy and imagination each has" [25]. As such, conviviality is the condition in which a person is free to produce without dependence or manipulation, and convivial tools are those that support an individual in this initiative.

Specific interactions with GenAI can support conviviality. Using Illich's work as a guide, we note three ways in which GenAI supports conviviality: 1) in relationship to the individual's ability to control the tool, or how it changes their relationship with control, 2) through the relationship with the individual's competence, both in supporting their competence through use and requiring it to produce an output, and 3) as it relates to an individual's ability to express themselves, both in providing the materials for expression, and generating following a user's expressions.

5.3.1 *Control.*

Conviviality requires that the producer maintain control over the tool used, rather than lose control of the productive process. Illich says of tools that, "An individual relates himself in action to his society through the use of tools that he actively masters" [25]. In terms of convivial AI tools, in certain interactions with GenAI, the user can experience control over the outcome, whereas in other interactions, the user collaborates with GenAI to produce an output.

Conviviality, which reinforces an individual's independent action, benefits from GenAI interactions that allow the user to intervene directly in the production process. In direction interactions, the user specifies the structure of what they are intending to produce, as in CoLadder's hierarchical prompts that generate a particular piece of code [60], thereby maintaining convivial control through fine-tailored prompts. Similarly, in encapsulation interactions, the user specifies the desired prompts as reusable actions, thereby maintaining control over the creation of the reusable components. Through interactions that provide careful guidance in this way, the user can express convivial control.

Additionally, GenAI tools enable production in which the user actively collaborates with the AI, allowing the user to work through ideas and thoughts with feedback from the GenAI. For example, rather than either the GenAI or the individual being fully responsible for an idea, an inspiration interaction involves the GenAI producing an initial output, from which the producer takes aspects of inspiration for their use. Similarly, in exploration interactions, the user and the GenAI can collaborate on AI-generated successors and on the user's careful exploration. Additionally, a user can specify the details of what they are trying to produce, either by providing a desired output in a concretization interaction or by selecting elements for a combination interaction. In these interactions, the GenAI is responsible for part of the control in producing the output.

5.3.2 *Competence.* For the autonomy associated with conviviality, depending on and working with an individual's competence is a key factor in the tools' conviviality. Looking at the nature of competence, Illich notes that "Progress should mean growing competence in self-care rather than growing dependence." [25]. Accordingly, Illich views the development and cultivation of general competence as an inherent feature of conviviality. GenAI interactions that involve a convivial competence either build upon a user's existing competence or position it as central to the productive process.

GenAI can increase a user's competence through two types of interactions: demonstration and indication. In demonstration interactions, the GenAI generates the remainder of a task that the user can learn from, such as how users co-composing music with GenAI would make an edit to the musical score, and the GenAI would demonstrate aspects of music theory as it re-composed other voices [35]. Similarly, indication interactions generate task information, which, when open-ended, enables "learning by doing" through user experimentation [8]. These interactions describe two distinct ways in which the user can acquire convivial competence with GenAI.

Some GenAI interactions use a human's inherent competence to produce an outcome. For example, direction interactions require continuous prompting and prompt refinement to produce the intended generation; however, once the effect is achieved, it has been reported to be rewarding [30]. Similarly, inspiration interactions generate ideas that the user must parse and interpret, and it is the user's responsibility to use them independently. In both interactions, individual competence is a key component of the process.

5.3.3 *Expression.* Illich warns that we need protection from unregulated tools and institutions that "curtail or negate any person's right to the creative use of his or her energy" [25]. For individuals in a convivial society, expression is tantamount to freedom. Illich says that convivial tools "allow the user to express his meaning in action" [25]. Certain interactions with GenAI tools can both provide the convivial user with materials for expression and generate content that contains their intentions.

Convivial expression is enabled when GenAI enables the user to create new material. For example, extraction interactions generate information that the user can utilize in their production, such as keywords describing thematic elements in an image; further, exploration interactions can be used to expand upon an idea and consider alternatives, ultimately resulting in a more diverse set of keywords than they would have come up with on their own [9]. Alternatively, inspiration interactions can provide source material that the user can extract ideas from and reuse elsewhere, such as elements of a generated story (characters, settings, etc.) that are incorporated into their own writing [10].

GenAI interactions can enable the user to amplify their meaning into a desired output, combining the machine's power with the user's inherent expression. For example, the interactions of concretization and combination enable the user to create outputs from their intention, such as in BISCUIT's generated interface, the user can specify various parameters they want their output to have, which the GenAI will then turn into code [8]; and similarly, in CreativeConnect, the keywords that the user selects over the course of an interaction are combined into a generated sketch [9]. In these cases, the user-provided pieces constitute the final synthesis.

6 Discussion

Following Barad's commitment "to understanding which differences matter, how they matter, and for whom" [4], this section examines how our three analyses produce findings that none of them could produce alone. We look for four kinds of patterns: constructive interference, where the interaction engages multiple theories strongly; destructive interference, where the interaction engages multiple theories weakly; partial engagement, where an interaction engages one or two theories but not the others; and fine-grained differences between similar interactions where details of our analysis produce divergent readings. In Figure 2, we display the diffraction pattern. Details of our analysis are shown by shading: low-engagement areas are shaded darker, and high-engagement areas are shaded lighter. One notable pattern of constructive interference appears in combination and indication, which engages strongly across all three theories. Combination produces value across theoretical commitments: transforming representations in Simon's terms, amplifying user expression in Illich's, and producing durable outputs in Arendt's. Indication does similar work along different dimensions: aiding search, providing access to ideas, and building user competence. Alternatively, we argue that concretization and encapsulation, as interactions, tend to engage with work and worldliness more than with conviviality, and notably more than with procedural rationality. Similarly, our analysis argues that inspiration interactions are those that engage most strongly with conviviality, but have limited engagement with procedural rationality, and even less with work and worldliness. Finally, we claim that substitution interactions mostly do not engage with the theories we consider here, suggesting that their benefit would be lower in this context.

In the following sections, we examine in more detail the differences among our analyses of four sets of interactions: aggregation vs. indication, direction vs. encapsulation, concretization vs. combination, and exploration vs. inspiration.

6.1 Aggregation and Indication

Aggregation is the interaction where the information necessary to produce an outcome is contained within the generated material. For example, Weber et al. found that programmers writing code with GenAI would not need to go to StackOverflow to look up how to implement a feature, but could instead use the generator to produce the feature for them [58]. Alternatively, indication interactions provide the user with the necessary information to complete a task, but they are expected to implement the feature themselves. For example, the GILT system provides information on the functions and parameters under consideration [36].

In alignment with procedural rationality, both interactions use externalized information to guide the productive process: automatically generated from the external information in the case of

Table 2: Cross-reference of the twelve interactions (top) against the three theoretical lenses (sections) and the eighteen aspects identified across them (left). Marks indicate where an interaction engages with an aspect of a theory. Each theory block is color-coded (Procedural Rationality in green, Work & Worldliness in blue, Conviviality in red), and within each theory block, columns are shaded by the strength of each interaction's engagement: lighter shading indicates stronger engagement, darker shading indicates weaker engagement.

	Combination	Indication	Concretization	Encapsulation	Direction	Aggregation	Differentiation	Demonstration	Inspiration	Exploration	Extraction	Substitution
Procedural Rationality												
Adds to Long-Term Memory							■	■				
Externalizes Long-Term Memory		■				■						■
Aids in Search		■				■						
Collaborates in Search	■								■	■	■	
Transforming Representation	■		■				■				■	
GenAI as new Representation				■	■					■		
Work & Worldliness												
Producing from Ideas			■	■								
Containing Production Ideas		■				■						
Intended Durability				■			■					
Allows Durability	■		■									
Finalization of Ends (From Means)	■		■			■						
Provides Means (For Ends)				■	■							
Conviviality												
User Can Control				■	■							
Share Control with GenAI	■		■						■	■		
Enables Competence		■						■				
Uses Competence					■				■			
Enables Expression									■	■	■	
Amplifies Expression	■		■									

aggregation, and generating relevant task information in the case of indication. Similarly, the worldly idea behind the output guides production in both interactions. However, while aggregation meets the worldly criteria of producing an end to the process in the form of the generated output, indication does not produce any final output. Additionally, the indication meets the value of conviviality by enabling the producer's competence by providing them with information, whereas the automatic generation of aggregation has no such impact. Accordingly, indication and aggregation are similar in their ability to use information to guide the task, but whereas aggregation provides an immediate output, indication first endows the user with competence and then requires them to produce the output. If the goal of the interaction is simply to produce material, then it would be implied to prefer aggregation over indication interactions. However, if the goal is for the human to have knowledge and understanding of the generated output before it is produced, then indication interactions would be preferred.

6.2 Direction and Encapsulation

Direction interactions involve providing detailed instructions to the GenAI regarding an intended output. For example, CoLadder allows the user to create hierarchical prompts to instruct GenAI on how to produce a particular intended function, where the users would spend significantly more time editing the prompt than they would editing the code produced by the GenAI [60]. Conversely, encapsulation interactions involve breaking the instruction to GenAI into modular components (e.g., "make it shorter", "make it more professional" [45]) that can be reapplied across multiple instances or combined to create stacked effects.

Regarding procedural rationality, both direction and encapsulation interactions yield new representations of problem-solving: direction interactions frame the problem-solving process as constructing an elaborate prompt, whereas encapsulation frames it as the creation of modular instructions. Additionally, both direction and encapsulation are viewed as worldly means, as they can be used to achieve a particular end: the final result of using prompts, whether detailed or modularized. Furthermore, as a convivial tool, both direction and encapsulation are controllable, either via a single prompt or through modular components. However, while direction interactions require a competent formulation of the prompt, the modular encapsulation components allow experimentation without this requirement. Additionally, encapsulation supports worldliness by reifying ideas into modular components and by giving those components durability through their intent to be reused. They are similar in their presentation of the task as a configurable means to produce an output, but whereas direction requires the user's competence, encapsulation avoids this requirement by providing a sustained output that the user can continue to work with over time. As such, directional interactions pose a challenge due to their reliance on a learned skill set, whereas encapsulation interactions are grounded in their more worldly creation of components.

6.3 Concretization and Combination

A concretization interaction involves describing the desired output that the GenAI produces. An example is LowCoderNL, where users describe a machine learning component they are looking for so that GenAI can configure it; users were found to need to provide information about "what, not how" they were trying to produce [44]. A combination interaction involves providing multiple elements to the GenAI, which then uses them to create a combined output. This can be seen in AutoSpark, where the user selects several design features from the workspace, and GenAI uses them to generate a synthesized design [7].

Concerning procedural rationality, both concretization and combination transform an input, a description for concretization, a selection of items for combination, into a new format, a synthesized output. In terms of worldliness, both interactions produce material, finalized outputs as a result of a fabrication process; for concretization interactions, the output is fabricated from a description, and for combination interactions, the output is produced by selecting multiple features. Additionally, in terms of conviviality, both of these interactions represent a collaboration with GenAI and are used as a means to express an idea: concretization, collaborating by translating the intention into output, and combination, collaborating by synthesizing the elements into a whole. However, while concretization serves the worldly purpose of a production bound to an idea, there is no such clear intention for combination. Similarly, a combination interaction serves the procedurally rational purpose of searching a problem space by connecting ideas into a new synthesis, whereas concretization adds no value when searching for a solution. Both interactions collaborate with GenAI to transform an input into an expressive product: concretization transforms an idea into a final output, while combination enables further exploration. As such, if the goal of the synthesis is to create a specific intentional object, then concretization would be preferred over combination. However, if the goal of the synthesis is to obtain an unknown solution, then combination would be preferred over concretization.

6.4 Exploration and Inspiration

Exploration interactions involve iterations of the GenAI providing suggestions and the user selecting between them to explore deeper. For example, in AutoSpark, the system provides related Kansei words that the user can select to explore further [7]. Alternatively, inspiration interactions involve the user generating content using some initial mechanism, which they then take ideas from for their own use. For example, in TaleBrush, the user could sketch a line to specify the fortune of a generated story, from which the users would take inspiration in the form of characters, settings, expressions, etc. [10]

With respect to procedural rationality, both interactions enable the user to search through ideas: exploration interactions provide iterative exploration by following suggestions, while inspiration interactions present ideas for the user to contrast against their internal understanding. Similarly, for conviviality, both interactions involve collaboration with the GenAI to produce material with which the user can express themselves: exploration interactions, which produce expressive materials from the contents of the exploration, and inspiration interactions, which involve collaboration with the GenAI to produce source material that the user then takes from. However, with respect to conviviality, inspiration-based interactions require the competence to apply GenAI ideas to the user's own work, whereas exploration does not. Conversely, with respect to procedural rationality, exploration offers a novel representation of a problem space in the exploration of suggestions, while inspiration does not create a new representation. Both interactions involve a collaborative effort in the search process to generate creative material, but inspiration relies on the human's own faculties to produce the final content, whereas exploration creates a new problem representation to explore it. As such, if the goal is to create something more original in the individual, then inspiration interactions are preferred, whereas if the goal is to effectively search for a solution to a problem, then exploration interactions are preferred.

7 Conclusion

In this work, we studied three theoretical perspectives on procedural rationality, work and worldliness, and convivial tools, and how they intersect to create intricate definitions of productivity. To this end, we developed a 12-part analytical framework for interactions with GenAI to improve productivity. We analyzed the framework using each of the three theories. Using this analysis, we performed a diffractive analysis to inform design implications, yielding nuanced choices such as aggregation and indication interactions, offering reification and competence, respectively. We present this critical framework as a tool for designing GenAI applications that account for values beyond efficiency. We hope this analysis can encourage a more diverse conversation around productivity with GenAI.

8 Disclosure

The authors wish to acknowledge the use of Gemini in the writing of this paper. This tool was used to assist in generating table formats. The paper remains an accurate representation of the authors' underlying work and novel intellectual contributions.

References

[1] Tyler Angert, Miroslav Suzara, Jenny Han, Christopher Pondoc, and Hariharan Subramonyam. 2023. Spellburst: A node-based interface for exploratory creative coding with natural language prompts. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22.

[2] Hannah Arendt. 1998. *The Human Condition* (2nd ed.). University of Chicago Press, Chicago. Introduction by Margaret Canovan.

[3] Yasharth Bajpai, Bhavya Chopra, Param Biyani, Cagri Aslan, Dustin Coleman, Sumit Gulwani, Chris Parnin, Arjun Radhakrishna, and Gustavo Soares. 2024. Let's Fix this Together: Conversational Debugging with GitHub Copilot. In *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 1–12.

[4] Karen Barad. 2007. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. duke university Press.

[5] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).

[6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.

[7] Liuqing Chen, Qianzhi Jing, Yixin Tsang, Qianyi Wang, Ruocong Liu, Duowei Xia, Yunzhan Zhou, and Lingyun Sun. 2024. AutoSpark: Supporting Automobile Appearance Design Ideation with Kansei Engineering and Generative AI. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–19.

[8] Ruijia Cheng, Titus Barik, Alan Leung, Fred Hohman, and Jeffrey Nichols. 2024. BISCUIT: Scaffolding LLM-Generated Code with Ephemeral UIs in Computational Notebooks. In *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 13–23.

[9] DaEun Choi, Sumin Hong, Jeongeon Park, John Joon Young Chung, and Juho Kim. 2024. CreativeConnect: Supporting Reference Recombination for Graphic Design Ideation with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–25.

[10] John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching stories with generative pretrained language models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–19.

[11] John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-Powered Worldbuilding with Generative Dust and Magnet Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–19.

[12] Roberto Cibin, Tiziano Bonini, Maurizio Teli, et al. 2023. Participatory design and critical media studies: a convivial conversation. *INTERACTION DESIGN & ARCHITECTURE (S)* 56 (2023), 156–173.

[13] Richard M Cyert and James G March. 1963. *A behavioral theory of the firm*. Prentice-Hall.

[14] Deloitte Global AI Institute. 2026. *The State of AI in the Enterprise: The Untapped Edge*. Technical Report. https://www.deloitte.com/us/en/about/press-room/state-of-ai-report-2026.html

[15] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science advances* 10, 28 (2024), eadn5290.

[16] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384, 6702 (2024), 1306–1308.

[17] Gerhard Fischer and Andreas Lemke. 1987. Constrained Design Processes: Steps Towards Convivial Computing. (06 1987), 50.

[18] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

[19] André Gorz. 1982. *Farewell to the working class: an essay on post-industrial socialism*. Pluto Press.

[20] Rui He, Huaxin Wei, and Ying Cao. 2024. An Interactive System for Supporting Creative Exploration of Cinematic Composition Designs. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–15.

[21] Martin Heidegger. 1962. *Being and Time*. Harper, New York. Translated by John Macquarrie and Edward Robinson.

[22] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *Iclr* 1, 2 (2022), 3.

[23] Yanwei Huang, Yunfan Zhou, Ran Chen, Changhao Pan, Xinhuan Shu, Di Weng, and Yingcai Wu. 2023. Interactive table synthesis with natural language. *IEEE Transactions on Visualization and Computer Graphics* 30, 9 (2023), 6130–6145.

[24] Ivan Illich. 1971. *Deschooling Society*. Harper & Row, New York.

[25] Ivan Illich and Anne Lang. 1973. Tools for conviviality. (1973).

[26] Nanna Inie, Jeanette Falk, and Raghavendra Selvan. 2025. How CO2STLY Is CHI? The Carbon Footprint of Generative AI in HCI Research and What We Should Do About It. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–29.

[27] Inger Jansson and Petra Wagman. 2017. Hannah Arendt's vita activa: A valuable contribution to occupational science. *Journal of Occupational Science* 24, 3 (2017), 290–301.

[28] Giorgos Kallis. 2018. *Degrowth*. Agenda Publishing, Redwood City, CA.

[29] Jeongyeon Kim, Sangho Suh, Lydia B Chilton, and Haijun Xia. 2023. Metaphorian: Leveraging large language models to support extended metaphor creation for science writing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 115–135.

[30] Charlotte Kobiella, Yarhy Said Flores López, Franz Waltenberger, Fiona Draxler, and Albrecht Schmidt. 2024. " If the machine is as good as me, then what use am I?"–How the use of ChatGPT changes young professionals' perception of productivity and accomplishment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.

[31] Amanda Lazar, Ben Jelen, Alisha Pradhan, and Katie A Siek. 2021. Adopting diffractive reading to advance hci research: A case study on technology for aging. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 5 (2021), 1–29.

[32] Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–22.

[33] Falk Lieder, Owen X Chen, Paul M Krueger, and Thomas L Griffiths. 2019. Cognitive prostheses for goal achievement. *Nature human behaviour* 3, 10 (2019), 1096–1106.

[34] John Locke. 2021. *Second Treatise of Government*. Project Gutenberg, Salt Lake City, Utah. https://gutenberg.org Release Date: April 22, 2003 [eBook 7370] [Most recently updated: December 25, 2021].

[35] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J Cai. 2020. Novice-AI music co-creation via AI-steering tools for deep generative models. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.

[36] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an llm to help with code understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.

[37] Allen Newell. 1990. *Unified theories of cognition*. Harvard University Press, Cambridge, MA.

[38] Jeongseok Oh, Seungju Kim, and Seungjun Kim. 2024. LumiMood: A Creativity Support Tool for Designing the Mood of a 3D Scene. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.

[39] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.

[40] Savvas Petridis, Nicholas Diakopoulos, Kevin Crowston, Mark Hansen, Keren Henderson, Stan Jastrzebski, Jeffrey V Nickerson, and Lydia B Chilton. 2023. Anglekindling: Supporting journalistic angle ideation with large language models. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–16.

[41] Crystal Qian and James Wexler. 2024. Take it, leave it, or fix it: measuring productivity and trust in human-AI collaboration. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 370–384.

[42] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).

[43] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*. Pmlr, 8821–8831.

[44] Nikitha Rao, Jason Tsay, Kiran Kate, Vincent Hellendoorn, and Martin Hirzel. 2024. AI for Low-Code for AI. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 837–852.

[45] Mohi Reza, Nathan M Laundry, Ilya Musabirov, Peter Dushniku, Zhi Yuan "Michael" Yu, Kashish Mittal, Tovi Grossman, Michael Liut, Anastasia Kuzminykh, and Joseph Jay Williams. 2024. ABScribe: Rapid Exploration & Organization of Multiple Writing Variations in Human-AI Co-Writing Tasks using Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.

[46] Chuhan Shi, Weiwei Cui, Chengzhong Liu, Chengbo Zheng, Haidong Zhang, Qiong Luo, and Xiaojuan Ma. 2023. Nl2color: Refining color palettes for charts with natural language. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2023), 814–824.

[47] Olivia Sidoti and Colleen McClain. 2025. 34% of U.S. Adults Have Used ChatGPT, About Double the Share in 2023. Pew Research Center. https://www.pewresearch.org/short-reads/2025/06/25/34-of-us-adults-have-used-chatgpt-about-double-the-share-in-2023/

[48] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2025. Ironies of generative AI: understanding and mitigating productivity loss in Human-AI interaction. *International Journal of Human–Computer Interaction* 41, 5 (2025), 2898–2919.

[49] H. A. Simon. 1957. *Models of man; social and rational*. Wiley.

[50] Herbert A Simon. 1976. From substantive to procedural rationality. In *25 years of economic theory: Retrospect and prospect*. Springer, 65–86.

[51] Herbert A Simon. 1996. *The Sciences of the Artificial*. MIT press.

[52] Alex Singla, Alexander Sukharevsky, Bryce Hall, Lareina Yee, Michael Chui, and Tara Balakrishnan. 2025. *The State of AI in 2025: Agents, Innovation, and Transformation*. Technical Report. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

[53] Frederick Winslow Taylor. 2004. *Scientific management*. Routledge.

[54] Jordan Taylor, Wesley Hanwen Deng, Kenneth Holstein, Sarah Fox, and Haiyi Zhu. 2024. Carefully Unmaking the "Marginalized User": A Diffractive Analysis of a Gay Online Community. *ACM Transactions on Computer-Human Interaction* 31, 6 (2024), 1–30.

[55] Samangi Wadinambiarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The effects of generative ai on design fixation and divergent thinking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.

[56] Rosalie A Waelen. 2025. Rethinking automation and the future of work with Hannah Arendt. *Journal of Business Ethics* 201, 1 (2025), 3–14.

[57] Sitong Wang, Zheng Ning, Anh Truong, Mira Dontcheva, Dingzeyu Li, and Lydia B Chilton. 2024. PodReels: Human-AI Co-Creation of Video Podcast Teasers. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 958–974.

[58] Thomas Weber, Maximilian Brandmaier, Albrecht Schmidt, and Sven Mayer. 2024. Significant productivity gains through programming with large language models. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (2024), 1–29.

[59] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.

[60] Ryan Yen, Jiawen Stefanie Zhu, Sangho Suh, Haijun Xia, and Jian Zhao. 2024. Coladder: Manipulating code generation via multi-level blocks. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–20.

[61] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. 841–852.

[62] J Diego Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21.