

Myths and Ironies of AI-Assisted Design

Paul C. Parsons
Purdue University
West Lafayette, USA
parsonsp@purdue.edu

Abstract

Generative AI is rapidly becoming embedded in design tools and workflows, accompanied by promises of increased productivity, reduced cognitive burden, and improved decision-making. These promises shape how AI systems are adopted and evaluated in design practice. However, research on design, automation, and sociotechnical work suggests that labor-saving technologies often reorganize work in unintended ways rather than simply reducing effort. In this paper, I examine AI-assisted design by focusing on how AI reshapes the cognitive, coordinative, and organizational conditions under which design work unfolds. The analysis is organized around myths and ironies of AI-assisted design, treating myths as early abstractions of design work that become institutionalized as technologies are embedded in practice. It identifies six widely articulated myths—including higher-level thinking, explainability, reduced cognitive load, simplified workflows, democratization, and objectivity—and the patterned ironies that follow as omitted cognitive and coordinative work reasserts itself in practice. The paper concludes with implications for AI-assisted design tools, professional practice, and design education.

CCS Concepts

• **Human-centered computing** → HCI theory, concepts and models; • **Social and professional topics**;

Keywords

Design cognition, AI, Design judgment, Ironies

ACM Reference Format:

Paul C. Parsons. 2026. Myths and Ironies of AI-Assisted Design. In *Proceedings of Designing Interactive Systems (ACM DIS '26)*. ACM, New York, NY, USA, 13 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Generative AI is rapidly becoming embedded in the tools, workflows, and infrastructures of design. As AI-enabled features move from experimental add-ons to default features, expectations about productivity, accountability, and “best practice” are solidifying in ways that will shape design workflows and evaluation norms for

years to come. From copilot features embedded in design tools to research summarizers, inspiration generators, and auto-layout assistants, AI systems are increasingly being woven into everyday design practice. Across tool marketing, practitioner discourse, and the emerging HCI and design literature, a recurring promise is that AI will automate routine or low-level work so that designers can focus on higher-level thinking, become more productive, and “do more with less.” In this narrative, AI appears as an accelerator that cleans up messy early-stage work, reduces cognitive burden, and democratizes design expertise.

Design theory has long emphasized that design work cannot be reduced to a sequence of technical problem-solving steps that can simply be accelerated without changing the nature of the work itself. Design is fundamentally a practice of situated judgment, sensemaking, and framing carried out under shifting constraints, evolving understandings, and competing values [4, 32, 40]. Designers continually interpret situations, frame and reframe what is at stake, and co-evolve problems and possible solutions rather than simply optimizing within a fixed and fully specified space of possibilities [7, 34]. From this perspective, activities often described as “routine”—such as sketching, prototyping, or synthesizing research—are not merely mechanical labor to be automated away; they are integral to how designers develop understanding, surface tensions, and exercise judgment.

A parallel body of research on automation and labor-saving technologies reinforces this caution. Classic analyses show that attempts to remove humans from routine tasks frequently increase demands at precisely the moments when judgment is most critical [1]. Historical studies of domestic technologies similarly demonstrate how devices marketed as time-savers reorganized work and expectations so that effort increased rather than decreased at the level of the household system [3]. Across domains, studies of complex sociotechnical systems have repeatedly observed that new capabilities tend to redistribute work, shift burdens, and introduce new dependencies, rather than yield unencumbered efficiency gains [17, 49].

Taken together, these traditions suggest that the promises accompanying AI adoption should be understood not simply as optimistic claims, but as *early abstractions of design work*—that is, simplified accounts produced at the point of technological introduction, before design practice has resettled around new capabilities [46, 50]. As AI moves from experimentation to infrastructure, such abstractions risk becoming institutionalized, shaping workflows, evaluation criteria, and expectations before designers have had the opportunity to adapt practices, redistribute responsibilities, and make visible the coordinative and interpretive work required to sustain design judgment in practice. Importantly, these are not simply mistaken descriptions of design work, but necessarily partial representations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM DIS '26, Singapore

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXX.XXXXXXX>

that become authoritative as they are embedded in tools, workflows, and evaluation practices.

Seen through this lens, the dominant story about AI in design—more speed, less drudgery, more room for higher-level thinking—appears incomplete. Research has begun to document how AI reshapes collaboration, workflows, and creative processes in design, including paradoxes in co-creative systems, shifts in designers’ roles, and mixed evidence about productivity and reflection [28, 30, 35, 36, 39, 51]. Much of this work, however, takes tools, tasks, or workflow configurations as its primary units of analysis, making it harder to see how AI reorganizes design work beyond localized interaction. As a result, its most consequential effects on design practice are easily misattributed or rendered invisible: pressures on framing, sensemaking, and responsibility appear as secondary or incidental, rather than as central features of how work is reorganized. This gap is increasingly consequential because adoption decisions are being made by teams, organizations, and educators under productivity- and automation-oriented abstractions that fix what counts as “the work” of design too narrowly.

I examine AI-assisted design through a lens that attends to the cognitive, coordinative, and organizational conditions of design work. Within this lens, recurring *myths* are treated as early abstractions of design work, and corresponding *ironies* as the patterned consequences that emerge as practice resettles around them. Myths are widely circulating abstractions of what AI is expected to do for design work—for example, enabling higher-level thinking, reducing cognitive load, simplifying workflows, democratizing expertise, or increasing objectivity. They are not strawmen, but plausible and widely repeated claims that are often partially realized in practice. Ironies are the recurring patterns through which AI’s effects, under common organizational and temporal conditions, complicate or undercut those abstractions as design work unfolds.

This paper makes three contributions. First, it articulates a set of recurring myths and corresponding ironies of AI-assisted design as an analytic lens for understanding how AI reconfigures design work, judgment, and messiness in patterned ways. Second, it develops an account of AI’s systemic effects on design work that brings design theory into conversation with automation and work-practice research to explain why these dynamics are not simply design flaws of individual tools. Third, it outlines implications for design practice, AI tool design, and design education, focusing on how designers and organizations might better recognize and respond to these ironies in ways that sustain framing, sensemaking, judgment, and coordination.

2 Background

2.1 Framing, Sensemaking, and Judgment in Design

Design research has long characterized design work as a form of cognition that differs in important ways from problem solving under fixed goals and constraints [4]. Rather than selecting optimal solutions within a predefined space, designers engage in ongoing processes of framing, sensemaking, and judgment in situations that resist full specification [33]. These processes involve interpreting

ambiguous information, negotiating values and priorities, and iteratively constructing both problems and solutions as understanding develops [7]. Accounts of reflective practice emphasize that designers do not merely apply techniques to well-defined problems, but engage in cycles of action and reflection through which the situation “talks back,” prompting reframing and reinterpretation [32]. Subsequent research describes design as a process of co-evolution between problem and solution, with frames emerging, stabilizing, and being revised over time as designers explore alternatives and encounter constraints [6].

Central to these accounts is the role of judgment [26]. Designers must continually weigh incomplete, conflicting, and context-dependent considerations—technical feasibility, user experience, organizational goals, ethical concerns—without access to definitive criteria or exhaustive information. Design situations are characterized by an irreducible “messiness,” in which no amount of method, analysis, or formalization can eliminate the need for situated judgment [8, 40]. From this perspective, the quality of design work depends not only on the production of artifacts, but on the cognitive conditions that support framing, sensemaking, and reflective judgment over time.

This view of design cognition has an important implication for AI in design. If framing, sensemaking, and judgment are central to design quality, then technologies introduced to “optimize,” “accelerate,” or “automate” aspects of design work cannot be evaluated solely in terms of task efficiency or output quantity. They must also be understood in terms of how they reshape the cognitive work through which designers interpret situations, explore possibilities, and make commitments under uncertainty.

2.2 Automation, Labor-Saving Narratives, and Unintended Consequences

A substantial body of work has examined how technologies introduced under labor-saving or efficiency-oriented narratives tend to reshape work in unintended ways. In a classic analysis of industrial automation, Bainbridge describes the “ironies of automation,” showing how systems designed to reduce human involvement often increase demands for human intervention at precisely the moments when situations are most complex and time-critical [1]. Automation frequently removes routine activities while leaving people responsible for monitoring, diagnosis, and recovery—tasks that are cognitively demanding, context-sensitive, and difficult to standardize.

In addition to this insight, technological change should be understood not as the substitution of functions, but as an intervention into an ongoing field of practice that reorganizes roles, coordination, and performance expectations over time [46]. From this perspective, early representations of work reflect an *envisioned* or *imagined* form of practice—what has been termed work-as-imagined—while the realities of work-as-done continue to evolve as people adapt to new capabilities, constraints, and pressures [5, 18]. The resulting misalignment between how work is imagined and how it is actually carried out is a persistent source of unintended consequences, particularly as technologies become embedded and expectations harden.

Historical studies of everyday technologies reveal similar dynamics outside of industrial and safety-critical domains. Cowan’s account of domestic appliances shows how technologies marketed as time-saving reorganized household labor and social expectations, frequently increasing rather than decreasing the amount of work required [3]. In these cases, efficiency gains at the level of individual tasks translated into expanded standards, new responsibilities, or redistributed labor at the level of the system.

Subsequent research on complex sociotechnical systems has generalized these observations across domains such as aviation, healthcare, and energy. New capabilities tend to be absorbed into existing organizational pressures, reshaping what counts as acceptable performance and redistributing work rather than eliminating it [17, 49]. Importantly, these effects are not simply the result of poor implementation or immature technology; they arise from the interaction between finite cognitive and organizational resources, ongoing change, and interdependent roles.

Taken together, this work suggests that claims about technologies “freeing up” human effort should be treated with caution. Rather than assuming net reductions in work or cognitive demand, it is often more productive to ask how new capabilities reconfigure the nature, timing, and visibility of human judgment and responsibility. Studies of work practice show that early functional or task-based representations of work routinely omit the coordinative and interpretive effort required to make systems function in practice, pushing this work onto practitioners as technologies are adopted [50]. This perspective does not deny that technologies can deliver real benefits, but it highlights why efficiency gains are frequently accompanied by new forms of workload, coordination, and vulnerability.

A key insight for understanding these dynamics is that early abstractions of work systematically omit the effort required to make work hold together in practice. Formal representations of work—whether framed as tasks, functions, or roles—necessarily simplify complex activity in order to make technological change tractable. In doing so, they underrepresent the ongoing effort required to integrate, coordinate, and adapt activities in situated settings. Strauss and colleagues describe this effort as *articulation work*: the work required to align people, artefacts, and expectations so that nominal tasks can actually be carried out in practice [31, 42]. Studies of function allocation demonstrate that when technologies are introduced on the basis of such abstractions, the omitted coordinative and interpretive work does not disappear; it is displaced onto practitioners as systems are adopted and institutionalized [50]. From this perspective, the “extra work” that accompanies many labor-saving interventions is not accidental, but a predictable consequence of abstractions that simplify work in order to justify and guide change.

For design, this history provides a crucial backdrop. If design quality depends on framing, sensemaking, and judgment, then technologies introduced to accelerate or simplify design work are likely to interact with those cognitive or coordinative processes in complex ways. Understanding AI in design therefore requires attention not only to what tasks are automated, but to how the remaining work of interpretation and judgment is reshaped.

2.3 AI in Design: Tools, Workflows, and Emerging Tensions

A growing body of work explores AI-supported ideation and co-creation, treating generative models as collaborators or “design materials” that designers can probe, steer, and adapt through interaction [2, 24, 44, 52]. These studies document how AI can accelerate the production of plausible alternatives, support early-stage exploration, and lower barriers to generating design artifacts. At the same time, literature consistently reports tensions that complicate these benefits. Designers often face challenges related to control, understanding, and integration: AI outputs may be difficult to align with evolving design intentions; prompting and refinement can require substantial effort; and designers frequently oscillate between treating AI as a source of inspiration and as a system that must be carefully constrained [13, 20, 43]. Rather than simply reducing effort, AI can relocate work into new forms of diagnostic, communicative, and learning-oriented activity, as designers attempt to understand what a system is doing, why it produced a particular result, and how to shape its behavior.

Several studies have begun to examine how different ways of integrating AI into workflows affect designers’ cognition and sense of agency. Xu et al. show that positioning AI as a leading versus supporting actor in a workflow produces different patterns of cognitive activity, confidence, and perceived ownership, with implications for how reflective and exploratory design work unfolds over time [51]. Similarly, research on timing effects suggests that introducing AI too early in a creative process can undermine autonomy and originality, whereas delayed or conditional use can better support human judgment [29]. These findings underscore that AI’s impact on design depends not only on what the technology can do, but on how it is embedded in the temporal and organizational rhythms of design work.

Other work focuses explicitly on designers’ need to understand the limitations, trade-offs, and assumptions embedded in AI systems. Studies of transparency and model understanding emphasize that designers seek forms of “designerly understanding” that support contextual reasoning and exploration, rather than exhaustive technical explanations [23]. Tools such as fAllureNotes demonstrate that supporting designers in understanding model failures and edge cases can improve reasoning and decision-making, but also require additional effort and structured reflection [25]. Likewise, work on communicating algorithmic trade-offs shows that values such as fairness, accuracy, and performance cannot be optimized simultaneously, and that designers must actively negotiate these tensions in context [53].

Across these studies, a common pattern emerges: while AI can offer meaningful support for certain activities, its integration into design practice often introduces new demands on designers’ framing, sensemaking, and judgment. Much of the existing literature, however, approaches these dynamics through the lenses of tools, tasks, or discrete workflow configurations. As a result, it can be difficult to see how localized findings—about ideation speed, prompt crafting, transparency needs, or agency—accumulate into broader patterns that reshape how design work is understood, organized, and evaluated at the cognitive and organizational level. From a

work-practice perspective, this fragmentation is not incidental; rather, empirical findings about individual tools or workflows can accumulate without converging on a shared account of how work is reconfigured at the system level [50].

3 Analytic Lens: Myths and Ironies of AI-Assisted Design

To examine how AI reshapes design work, I adopt an analytic lens organized around myths and ironies of AI-assisted design. This lens builds on prior work on the unintended consequences of automation and technological change, while shifting the focus from operator performance or system safety to the cognitive and coordinative work of design. In this paper, myths are treated as *early abstractions of design work*: simplified accounts of what matters, what can be delegated, and what constitutes “good practice” when new capabilities are introduced. These abstractions circulate widely in tool marketing, practitioner discourse, and research framing, where they foreground aspects of design that are immediately visible and measurable, such as speed, output, and automation.

This lens was developed through an interpretive synthesis. It draws on recurring promises articulated across AI-assisted design literature and practitioner/tool discourse, and interprets them through scholarship on design theory, work practice, and automation. The six myths were selected for their recurrence, salience, distinctiveness, and explanatory usefulness; they are not intended as an exhaustive taxonomy or claim about prevalence, but as a bounded analytic lens for organizing recurring patterns in current discourse and practice. In this paper, myths refer to the recurring abstractions that shape expectations and adoption, while narratives refer to the discursive forms through which those abstractions circulate.

Ironies, in turn, are the patterned consequences that emerge as design practice resettles around these early abstractions. As AI becomes embedded in everyday work, aspects of design that were omitted or backgrounded in these early abstractions—such as framing, sensemaking, articulation, and judgment—reassert themselves in practice. Ironies are therefore not isolated failures or implementation errors, but recurrent dynamics that arise when simplified abstractions of work encounter the situated, evolving realities of practice under finite resources, time pressure, and organizational constraint. From this perspective, myths and ironies are analytically inseparable: myths provide the abstractions that make adoption possible and legible, while ironies reveal the limits of those abstractions as practice unfolds.

Organizing the analysis around myths and ironies serves two purposes. First, this lens connects localized empirical observations—about ideation speed, prompting effort, explainability, or workflow integration—to broader patterns that recur across tools, teams, and organizational settings. Second, it foregrounds the cognitive and organizational consequences of AI adoption without assuming that technologies are either inherently beneficial or inherently harmful. The goal is not to adjudicate whether AI is good or bad, but to make visible the trade-offs that accompany its institutionalization in design practice.

In the sections that follow, I examine a set of recurring myths of AI-assisted design and articulate the corresponding ironies that emerge as practice resettles around them. The myths are ordered to move from individual cognitive conditions to organizational and epistemic consequences. The first three focus on time, representation, and cognitive effort; the remaining myths address workflow institutionalization, redistributed expertise and responsibility, and the displacement of value judgments. Each myth–irony pair includes a description of the abstraction that motivates adoption, the patterned counter-effects that arise under common conditions, and the implications for design practice. Together, these analyses provide an account of how AI reshapes design work over time, and why its most consequential effects are often less visible than its immediate gains.

4 Six Myths and Their Ironies

4.1 Myth 1: AI Frees Designers to Focus on Higher-Level Thinking

AI is frequently framed—by tool vendors, practitioners, and recent research accounts—as taking over routine or low-level work so that designers can focus on higher-level thinking and be more productive [13, 20, 51]. In this narrative, AI-generated drafts, layouts, summaries, and alternatives reduce the effort required for execution, leaving designers free to concentrate on strategy, concept development, and problem framing—treating these activities as separable from the time and interaction through which they are actually accomplished. This promise recurs across tool marketing and practitioner discourse, and it is echoed in research that highlights gains in speed and output when AI is integrated into design workflows. The appeal of this myth is understandable. Many activities in design—such as producing initial artifacts, generating variations, or synthesizing large amounts of information—can indeed be accelerated by AI. Designers often report that these capabilities feel empowering, particularly in early stages of exploration or under tight deadlines. At a local level, AI can reduce the time required to produce plausible representations and can lower barriers to trying out ideas.

Irony: acceleration compresses time for framing and sensemaking. More broadly, acceleration reorganizes the temporal conditions under which design work unfolds, constraining not only framing and sensemaking, but also the articulation, coordination, and judgment required to sustain them. The abstraction underlying this myth foregrounds the speed of artifact production while backgrounding the temporal slack, iterative engagement, and coordination through which framing and sensemaking are sustained in practice. Rather than freeing designers to engage in higher-level thinking, AI-driven acceleration can compress the temporal conditions under which such thinking is sustained in design. As production becomes faster and cheaper, organizational expectations tend to adjust accordingly: more alternatives are requested, more iterations are produced, and timelines tighten. What appears as “freed capacity” at the level of individual tasks is frequently absorbed as increased throughput at the level of projects and teams. This pattern reflects a well-documented dynamic in automation and productivity contexts: efficiency gains are rarely preserved as slack,

but are instead recaptured as increased demand, tighter timelines, or expanded scope [1, 38, 46].

From a design cognition perspective, this matters because framing and sensemaking depend on time, iteration, and repeated engagement with evolving situations. Framing a design problem, interpreting feedback, and revising one's understanding of what matters require temporal slack and repeated engagement with ambiguity. When AI enables rapid production of artifacts, it can inadvertently shift attention toward selection and delivery—choosing among outputs and meeting deadlines—at the expense of revisiting assumptions, exploring alternative framings, or questioning whether the right problem is being addressed.

Empirical studies of AI-supported design workflows illustrate this tension. Research comparing different modes of AI integration shows that workflows positioned to maximize productivity and output can change how designers allocate cognitive effort, increasing activity around comprehension and synthesis of AI-generated material while reducing opportunities for reflective exploration [51]. Other studies report that designers experience pressure to keep pace with AI's generative capacity, leading them to accept or lightly modify outputs rather than engaging deeply with the underlying design space [13, 20]. In these settings, the promise of "higher-level thinking" is undermined by the reorganization of time and attention around visible production.

This irony is not simply a matter of workload. Designers are not necessarily doing more work overall; rather, they are doing different work, under different temporal conditions. Under common conditions of adoption, AI can shift designers from extended cycles of framing and reflection toward shorter cycles of generation, evaluation, and selection. While these activities are part of design, they do not substitute for the sensemaking required to redefine goals, negotiate values, or reframe problems as situations evolve. Over time, this shift can make higher-level thinking harder to sustain, even as output increases.

The irony becomes particularly salient in organizational contexts where productivity gains are tracked and rewarded. When speed and volume become proxies for effectiveness, the cognitive work of framing risks becoming invisible and undervalued. Designers may still be responsible for making sense of complex situations and for ensuring the coherence and appropriateness of outcomes, but they are afforded less time and space to do so. In this way, AI does not eliminate the need for higher-level thinking; rather, it can constrain the conditions under which such thinking is possible.

This analysis does not deny that AI can support valuable forms of exploration or reduce friction in design work. Rather, it highlights a patterned trade-off: under common conditions of adoption, acceleration at the level of artifact production can compress the time and attention available for framing and sensemaking. Recognizing this irony is essential for evaluating AI's impact on design practice and for designing tools and organizational practices that genuinely support higher-level design cognition rather than crowding it out.

4.2 Myth 2: Explainability Restores Understanding and Control

A common response to concerns about AI accelerating design work is the claim that explainability can restore understanding and control. If designers can see why an AI system produced a particular output—what factors mattered, what assumptions were made, or how alternatives were weighted—then they can remain appropriately "in the loop," calibrate trust, and exercise judgment. From this perspective, explainability is framed as the primary safeguard against loss of agency: even if AI generates quickly, designers can slow down, inspect the rationale, and decide how to proceed.

This belief is reinforced by both regulatory and organizational pressures. Designers are increasingly expected to justify decisions using evidence that appears systematic, transparent, and defensible. Explainable AI promises to meet these demands by turning opaque model behavior into legible reasons, often presented through textual rationales, feature attributions, or performance summaries. In design contexts, such explanations are frequently positioned as a way to preserve reflective oversight while benefiting from AI's speed and scale.

Irony: explanation reorganizes sensemaking without restoring control. The abstraction underlying this myth treats understanding as something that can be delivered through post-hoc explanation, rather than as a situated process of interpretation, framing, and alignment carried out over time. In practice, explanation does not function as a simple aid to understanding when it is introduced as a substitute for ongoing sensemaking rather than as a support for it. Instead, it often introduces a new representational layer that designers must interpret, interrogate, and reconcile with their own evolving understanding of the design situation. Rather than answering the questions designers most urgently need to ask—What is happening in this context? What matters here? What should we do next?—explanations typically offer post-hoc rationales tied to the model's internal logic. Designers must then assess whether those rationales are relevant, complete, or aligned with the values, constraints, and priorities of the situation at hand. In the case of generative AI systems, explanations are often delivered through fluent natural language, which can amplify perceived confidence and coherence without corresponding increases in accuracy or appropriateness, making trust calibration more difficult rather than easier [9].

Empirical studies of AI transparency in design illustrate this tension. Research on "designerly understanding" shows that designers do not primarily seek exhaustive or technical explanations of model behavior; instead, they look for contextual insight that supports exploration, comparison, and reasoning about trade-offs [23]. Similarly, tools such as fAllureNotes demonstrate that helping designers understand model limitations and failure cases can improve reasoning and decision-making, but doing so requires additional effort in annotation, synthesis, and reflection [25]. More broadly, empirical work on explainable AI suggests that explanation needs are highly role-dependent, that many stakeholders need to understand how systems fail and mislead as much as how they work, and that explanation alone is often insufficient without access to other people and

opportunities for self-explanation and exploration [16]. These activities are cognitively demanding and time-consuming, particularly when designers are already operating under tight deadlines.

From a design cognition perspective, this matters because explanation does not replace sensemaking—it competes with it. Designers must decide how much weight to give an AI explanation, whether to trust it, how to communicate its implications to stakeholders, and when to override it. Each of these decisions involves judgment under uncertainty. When explanations are introduced without corresponding changes to workflow, representations, or organizational expectations, they can slow designers down in unproductive ways or be ignored altogether. In either case, the promise of restored control is undermined.

This irony is amplified by temporal pressure. Design work unfolds in a flow of activity where time does not pause for reflection on every output. Explanations that require designers to stop, interpret, and validate AI reasoning can disrupt this flow, especially when multiple AI-generated artifacts must be assessed in parallel. Under such conditions, explanations risk becoming noise—another stream of information that designers must manage—rather than a genuine support for understanding. The result is not enhanced control, but an increased burden of interpretive and justificatory work layered on top of already compressed timelines.

Importantly, this critique does not suggest that explanation is unnecessary or undesirable. Rather, it highlights a mismatch between what explanation is often assumed to provide and what designers actually need to sustain judgment. In design, understanding is not solely about knowing why a particular output was produced; it is about seeing how that output fits within a broader, evolving frame of goals, constraints, and values. Explanations that focus narrowly on model reasoning can therefore miss the mark, offering clarity at the wrong level of abstraction.

As a result, explainability can inadvertently reinforce the very dynamics it is meant to counteract. When AI explanations are treated as sufficient evidence of understanding, organizations may feel justified in increasing throughput or tightening timelines, assuming that designers can rely on explanations to maintain oversight. In practice, this can further erode the conditions under which framing and sensemaking occur. The irony, then, is that explanation—introduced to preserve designer control—can become another source of cognitive load and coordination work, making it harder rather than easier for designers to exercise reflective judgment.

4.3 Myth 3: AI Reduces Cognitive Load

A closely related promise in discussions of AI-assisted design is that AI reduces cognitive load. By generating options, summarizing information, and narrowing the space of possibilities, AI is expected to make design decisions easier and less mentally demanding. In this view, AI lightens designers' cognitive burden by handling information-intensive or repetitive aspects of the work, allowing designers to operate more efficiently and with less effort [20, 51].

This promise aligns with familiar models of human–computer interaction in which cognitive load is treated as a quantifiable resource

that can be reduced through better interfaces, automation, or decision support. Many AI tools explicitly target this goal, advertising features such as automatic synthesis, suggestion ranking, or smart defaults intended to simplify designers' choices. Empirical studies often report local reductions in effort for specific subtasks—such as drafting, variation generation, or initial ideation—lending further plausibility to the claim.

Irony: AI redistributes cognitive load without reducing the demands of judgment. The abstraction underlying this myth treats cognitive load as a property of individual tasks that can be reduced through delegation or automation, rather than as an emergent property of interpretation, coordination, and judgment across a design process. Rather than eliminating cognitive effort, AI in design can relocate cognitive load into new forms of work. Designers must evaluate the relevance and quality of AI-generated outputs, diagnose mismatches between suggestions and evolving intentions, and decide when to trust, modify, or reject what the system provides. This diagnostic and evaluative work is cognitively demanding, particularly because it must be performed under conditions of uncertainty and time pressure.

Research on AI-supported design consistently documents this shift. Designers report spending substantial effort learning how to prompt systems effectively, interpreting why certain outputs appear, and managing variability across generations [13, 20, 43]. Studies comparing different workflow configurations further show that while AI can reduce effort associated with producing artifacts, it can simultaneously increase effort devoted to comprehension, synthesis, and decision-making about AI outputs [51]. In effect, cognitive load is not removed; the composition of cognitive work is reorganized, shifting from generative activity toward oversight, interpretation, and coordination. This redistribution can take a polarized form: AI reduces effort for relatively routine or low-complexity tasks while increasing the cognitive demands associated with high-stakes, ambiguous, or exceptional cases—precisely where judgment is most critical [1, 38].

This redistribution matters because not all cognitive work is interchangeable, and not all forms of effort contribute equally to making sense of the design situation. Evaluating AI outputs requires designers to reason about context, intent, and consequences without always having access to the assumptions embedded in the system. Unlike traditional design activities—such as sketching or prototyping—where effort invested in making artifacts often directly contributes to understanding the problem space, diagnostic work on AI outputs can be detached from designers' own framing processes. Designers may expend substantial mental effort assessing artifacts that were not generated through their own inquiry, which can make it harder to build coherent mental models of the design situation.

This irony is especially evident in collaborative settings. When AI-generated artifacts become shared reference points, designers must also perform additional coordination work: explaining outputs to teammates, justifying decisions to stakeholders, and reconciling AI suggestions with organizational constraints. What appears as reduced individual cognitive load at the point of generation can

therefore translate into increased collective cognitive burden downstream. As with explainability, the cognitive work introduced by AI is often less visible and less formally acknowledged, even as it becomes central to getting work done [14].

This analysis does not suggest that AI inevitably increases cognitive load overall. In many cases, AI does reduce effort for specific subtasks. The irony lies in the assumption that such reductions straightforwardly translate into cognitive relief for designers. In practice, AI tends to shift effort toward forms of judgment, diagnosis, and coordination that are cognitively demanding and difficult to compress. When organizational expectations adjust to AI-enabled speed, designers may find themselves managing more information, more alternatives, and more decisions within the same—or shorter—timeframes. AI changes the composition of cognitive work in design, often increasing the demands placed on sensemaking and judgment even as it accelerates production. Recognizing this redistribution is essential for evaluating AI's impact on design practice and for designing tools and organizational practices that support, rather than undermine, designers' cognitive capacities.

4.4 Myth 4: AI Simplifies and Stabilizes Design Workflows

AI is often presented as a way to simplify and stabilize design workflows, treating design as a process that can be standardized and reliably executed through integrated pipelines. By automating steps, standardizing outputs, and integrating seamlessly into existing tools, AI systems are expected to reduce friction, smooth handoffs, and make design processes more predictable. In this view, AI helps replace ad hoc or messy practices, enabling teams to move more efficiently from ideation to delivery.

This promise is reinforced by how AI tools are frequently introduced in practice. Many systems are embedded directly into familiar environments such as design and prototyping tools, offering features that appear to slot cleanly into existing workflows—automatic layout suggestions, prompt-driven ideation panels, or AI-assisted content generation. Research and tool demonstrations often emphasize how such integrations can reduce context switching and simplify coordination by keeping designers “in flow” [24, 27]. At a surface level, these integrations can indeed make workflows appear more coherent and easier to manage. This apparent coherence is reinforced by the fact that many AI-supported workflows perform well under the constrained conditions of demonstrations or early adoption, creating the impression that extending these pipelines to the full range of contingencies encountered in practice will be straightforward [47].

Irony: apparent workflow simplification often introduces hidden coupling and brittleness. The abstraction underlying this myth foregrounds stable handoffs and standardized sequences, while backgrounding the ongoing coordination, interpretation, and adaptation required to sustain workflows as conditions change. Rather than stabilizing design workflows, AI integration can introduce new forms of coupling between tools, representations, and organizational expectations. This reflects a familiar substitution logic in automation: AI is introduced to replace or stabilize work as it is imagined, while the realities of work as done continue to

evolve, creating gaps that designers must bridge through adaptation and workaround [19, 21, 41]. While individual steps may become easier or faster, the overall workflow can become more tightly interdependent and less transparent. Designers must learn how AI systems behave in different contexts, anticipate their failure modes, and adapt their practices accordingly. These dependencies are often implicit, making it harder to diagnose problems when workflows break down.

Empirical studies of AI-supported design workflows illustrate this dynamic. Research examining AI as a design material shows that designers must continually negotiate between AI outputs, design intentions, and organizational constraints, often relying on informal boundary objects and workarounds to maintain coherence across teams [52]. Similarly, studies of tightly coupled AI–UI workflows report that while AI features can reduce friction in isolated tasks, they also require designers to manage new forms of coordination, versioning, and interpretation when AI-generated artifacts interact with existing design systems and collaboration practices [13, 27].

From a systems perspective, this is important because simplification at the surface can mask growing complexity beneath [48]. As AI becomes embedded in design workflows, teams may come to rely on its outputs as stable inputs for downstream activities—reviews, stakeholder discussions, or implementation—without fully understanding the assumptions or conditions under which those outputs were generated. When AI behavior changes, produces unexpected results, or fails to generalize to new contexts, these hidden dependencies can surface abruptly, creating brittle points in the workflow that are difficult to anticipate or repair.

This irony is compounded by organizational pressures for consistency and efficiency. Once AI-supported workflows are perceived as “working,” they can quickly become institutionalized as standard practice. Designers may be encouraged—or required—to follow AI-augmented pipelines that privilege speed and uniformity, even when those pipelines constrain opportunities for reframing or local adaptation. Over time, workflows that appear stable can become resistant to change, reducing designers' ability to respond flexibly when design situations evolve.

This critique does not imply that workflow integration is inherently problematic. Well-designed integrations can reduce unnecessary friction and support coordination across tools and teams. The problem is the assumption that simplification at the level of individual steps automatically yields robustness at the level of the system. In practice, AI-assisted workflows can trade visible simplicity for hidden coupling, shifting the burden of resilience onto designers' judgment and improvisation when things do not go as planned. The practical implication is that workflows should be assessed not only for local efficiency, but for how well they remain adaptable, interpretable, and repairable as design conditions change.

4.5 Myth 5: AI Democratizes Design

AI is widely framed as a democratizing force in design. By lowering barriers to producing plausible artifacts—layouts, visuals, copy, or interaction concepts—AI is expected to make design capabilities accessible to a broader range of people. In this narrative,

non-designers can participate more fully in design work, while designers are freed from technical bottlenecks and routine production tasks. Expertise, from this perspective, becomes less centralized, and design becomes more inclusive and participatory.

This promise resonates with long-standing aspirations in design and HCI to broaden participation and reduce gatekeeping. Empirical studies of generative AI in design contexts often report that AI tools enable rapid exploration and allow users with varying levels of formal training to produce outputs that appear “design-like” [20, 22]. In collaborative settings, AI-generated artifacts can serve as common reference points, enabling stakeholders to contribute ideas and preferences without mastering specialized tools or techniques.

Irony: expertise is redistributed rather than eliminated. The abstraction underlying this myth treats democratization as expanded access to producing artifacts, while backgrounding how expertise in framing, judgment, and accountability is exercised and distributed in practice. Rather than removing the need for design expertise, AI often reconfigures where expertise resides and how it is exercised. Under common conditions of adoption, this redistribution can entail reduced opportunities to develop and exercise core design judgment, alongside misplaced responsibility as designers remain accountable for outcomes increasingly shaped by AI-generated suggestions [1, 36]. As AI systems take on greater responsibility for producing artifacts, designers’ work shifts away from direct creation toward evaluation, interpretation, and governance of AI outputs. Designers become responsible for diagnosing when outputs are appropriate, identifying subtle mismatches with context or intent, and anticipating downstream consequences that may not be visible in the generated artifacts themselves.

Empirical studies of AI-supported design illustrate this redistribution. Research on designers engaging AI as a design material shows that effective use often depends on tacit knowledge about when to trust outputs, how to probe system behavior, and how to translate between AI-generated representations and organizational constraints [13, 52]. Similarly, work examining AI-enabled participation highlights that while AI can broaden involvement, it can also introduce new forms of anchoring, constraint, or deference to generated artifacts, shifting the burden of critical judgment onto designers who must manage these dynamics [22].

From a design cognition perspective, this matters because design expertise is not limited to producing artifacts; it lies in the capacity to frame problems, negotiate values, and exercise judgment under uncertainty within evolving conditions of practice. When AI-generated outputs appear polished and complete, the expertise required to assess their appropriateness can become less visible, even as it becomes more consequential. Designers may find themselves accountable for outcomes shaped by AI suggestions, while having less explicit authority over the process that produced them.

This irony is amplified in organizational contexts where AI outputs are treated as neutral or authoritative starting points. Non-designers may reasonably assume that AI-generated artifacts reflect best practices or objective standards, making it harder for designers to challenge or reframe early proposals. As a result, democratization at the level of artifact production can coexist with increased concentration of responsibility at the level of judgment and accountability.

This analysis does not deny the possibility that AI can support more inclusive forms of participation in design. AI can help surface perspectives, lower technical barriers, and facilitate engagement. The problem is the assumption that broader access to artifact production translates directly into shared design judgment. In practice, AI often shifts the locus of expertise from making to evaluating, from producing to governing, and from visible craft to less visible forms of sensemaking and responsibility. Participation may broaden even as the cognitive and ethical demands placed on designers intensify, concentrating responsibility in less visible ways.

4.6 Myth 6: AI Makes Design More Objective and Less Biased

AI is widely presented as a way to make design decisions more objective and less biased. By grounding judgments in data, statistical inference, and model-driven analysis, AI is expected to reduce the influence of individual subjectivity and inconsistency. In this narrative, AI offers designers access to evidence-based insights—about users, preferences, performance, or outcomes—that promise more neutral, defensible, and fair design decisions.

This promise is particularly compelling in organizational contexts where designers are asked to justify decisions to diverse stakeholders. AI-generated outputs are often accompanied by quantitative signals, confidence scores, or explanatory rationales that resemble familiar forms of analytical evidence. Research on AI-supported design and responsible AI frequently frames transparency and trade-off communication as mechanisms for supporting such objectivity, suggesting that clearer articulation of model behavior can help designers reason more rigorously about competing goals and values [23, 53].

Irony: value judgments are displaced rather than removed. The abstraction underlying this myth treats objectivity as a property of tools and data, rather than as an achievement of situated reasoning that requires ongoing negotiation of values in context. Rather than eliminating subjectivity or bias, AI in design can relocate value judgments into less visible places. Once presented, AI-generated outputs can also function as cognitive anchors, shaping what designers attend to next and increasing confirmation effects, even when those outputs are incomplete, misaligned, or inappropriate for the context [9]. Choices about what data to collect, how outcomes are defined, which metrics are optimized, and how trade-offs are balanced all embed normative assumptions. When AI systems surface outputs as seemingly neutral recommendations, these assumptions can become obscured, making it harder for designers to recognize and interrogate the values shaping the results.

Empirical work on transparency and trade-offs in AI-supported design illustrates this displacement. Studies show that designers must grapple with inherent conflicts—such as accuracy versus fairness, efficiency versus inclusion, or consistency versus contextual sensitivity—and that these conflicts cannot be resolved through optimization alone [53]. Other research demonstrates that designers seek forms of understanding that allow them to explore how model behavior changes across contexts and user groups, rather than static explanations that imply a single correct answer [23]. Tools that support such exploration can improve reasoning, but

they also make explicit the persistence of value-laden judgment rather than eliminating it.

From a design cognition perspective, this matters because design quality depends on the explicit negotiation of values in context. When AI-generated outputs are treated as objective baselines or authoritative starting points, the work of articulating and debating values can be short-circuited. Designers may still be responsible for ensuring that outcomes are appropriate, ethical, and context-sensitive, but they must now do so against artifacts that appear to carry epistemic authority. Challenging or reframing these outputs can require additional justification, particularly in organizations that privilege data-driven decision making.

This irony is amplified by institutional and accountability pressures. When AI outputs are used as evidence of rigor or fairness, designers may be asked to explain deviations from model recommendations rather than to defend the underlying design rationale. In such settings, objectivity becomes a property attributed to the tool rather than a quality of the reasoning process. The burden of judgment is not reduced; it is reassigned, often becoming more difficult to exercise and defend.

This analysis does not deny that AI can help surface patterns, reveal biases, or support more systematic consideration of trade-offs. AI can play a valuable role in making certain assumptions visible and in expanding the space of considerations designers must attend to. The irony lies in assuming that these contributions render design decisions objective or value-neutral. In practice, AI changes how values are encoded, negotiated, and justified, without removing the need for human judgment. AI can redistribute where and how value judgments are made, often making them less visible and more difficult to contest. Recognizing this displacement is essential for understanding how epistemic authority and responsibility are reorganized in AI-assisted design, and for ensuring that design remains a practice of reflective, accountable judgment rather than deference to seemingly neutral outputs.

Taken together, the myths and ironies outlined above suggest diagnostic signals that warrant attention as AI becomes embedded in design work. These include rapid gains in output accompanied by fewer opportunities for reframing; growing reliance on explanations, metrics, or rankings that require additional interpretation and justification; shifts of cognitive effort from making toward evaluating, coordinating, or defending outputs; workflows that appear stable in demonstrations but depend on increasing informal workarounds in practice; and situations in which responsibility for outcomes remains human while epistemic authority is attributed to AI-generated artifacts. Attending to these signals does not resolve the ironies identified here, but it can help designers, educators, and toolmakers anticipate how design judgment is being reorganized as AI adoption proceeds.

5 Discussion

The six myths and corresponding ironies articulated in this paper describe recurring patterns in how AI is adopted, evaluated, and institutionalized in design practice. Each myth captures a plausible abstraction that foregrounds certain benefits while backgrounding others; the corresponding ironies emerge as omitted cognitive and

coordinative work reasserts itself under common organizational and temporal conditions. Taken together, these patterns suggest that the most consequential impacts of AI-assisted design are found in how AI reorganizes the conditions under which framing, sensemaking, articulation, and judgment are sustained.

This discussion synthesizes these patterns, situates them within ongoing scholarly conversations, and outlines implications for AI tool design, design practice, and design education. In doing so, it highlights diagnostic signals, evaluative questions, and practical implications that support reasoning about how AI adoption reorganizes design work as it becomes embedded in practice.

5.1 Why the Myths Persist

Here, myths are treated not as misconceptions to be corrected but as recurring narratives that organize adoption, evaluation, and justification. This framing resonates with existing trends in AI-in-design research that document paradoxes and tensions in co-creative work (e.g., speed vs. reflection; control vs. autonomy) [30]. The contribution here differs in focusing on myths as *adoption narratives*—stories that coordinate expectations across stakeholders and make AI legible as a solution to organizational demands—rather than as abstract tensions inherent to co-creation. In practice, these narratives persist because they are *instrumental*: promises of higher-level thinking, reduced cognitive load, workflow simplification, democratization, and objectivity align closely with managerial logics of efficiency, scalability, and accountability.

A second reason myths persist is that they are *locally and temporarily plausible*. Empirical studies show genuine gains in particular moments of practice: faster generation of options, more rapid production of plausible artifacts, or helpful assistance in early ideation [2, 20, 24]. These local benefits provide strong experiential evidence that appears to validate broader claims. Yet other work shows that as AI becomes embedded in workflows, the downstream consequences—evaluation burden, increased coordination, pressure to keep pace with AI’s generative capacity—become more salient [13, 14, 51]. This mismatch in timescales—immediate gains versus cumulative reconfiguration of work—helps explain why myths remain compelling even in communities that recognize AI’s limitations.

A third reason these myths remain compelling lies in how AI capabilities are demonstrated and evaluated. As Woods [47] observes, technological optimism is often fueled by the presumption that because a capability can be demonstrated under some conditions, extending it to handle the full range of complexities that arise and change over real-world life cycles will be straightforward. Contemporary AI systems are particularly effective at producing convincing demonstrations, performing impressively on curated tasks, controlled inputs, or narrow problem framings. These demonstrations create the impression that remaining challenges are matters of scale or refinement rather than qualitatively different demands. In practice, however, design work unfolds across shifting goals, evolving constraints, multiple stakeholders, and unanticipated breakdowns. The gap between what AI systems demonstrate under idealized conditions and what is required to support design judgment across this

Table 1: Recurring myths of AI-assisted design, the common narratives that sustain them, and the corresponding ironies that emerge in practice.

Myth	Common Narrative	Irony in Practice
AI frees designers to focus on higher-level thinking	AI automates routine or low-level work, allowing designers to devote more time to strategy, framing, and concept development.	Time for framing and sensemaking is compressed; freed capacity is recaptured as increased throughput.
AI explanations restore understanding and control	Explainable AI enables designers to inspect rationales, calibrate trust, and remain “in the loop.”	Explanations reorganize sensemaking rather than restore control, adding interpretive and justificatory burdens.
AI reduces cognitive load	AI simplifies decisions by generating options, summarizing information, and narrowing choices.	Cognitive effort is redistributed rather than reduced, shifting from generating toward diagnosing, evaluating, and coordinating outputs.
AI simplifies and stabilizes design workflows	AI standardizes processes and integrates into existing tools, making workflows more predictable and efficient.	Apparent simplification introduces hidden coupling and brittleness; workflows become more interdependent and harder to adapt as conditions change.
AI democratizes design	AI lowers barriers to producing design artifacts, broadening participation and reducing reliance on expert designers.	Expertise is redistributed rather than eliminated; designers take on greater evaluative responsibility, often with reduced visible authority.
AI makes design more objective and less biased	Data-driven AI outputs provide neutral, evidence-based guidance that reduces subjectivity and bias.	Value judgments are displaced rather than removed, becoming embedded in data, metrics, and models and harder to contest.

broader range of situations helps explain why early abstractions persist and why their limitations are systematically underestimated.

Beyond the persuasive power of demonstrations, a further factor sustaining these myths is what Woods describes as *conceptual lag*: the tendency for interpretive frames and analytic concepts to remain anchored to earlier, simpler forms of work even as systems grow in scale, coupling, and architectural complexity [48]. As AI systems are integrated into design practice, they are often understood through familiar task-level or substitution-based lenses that fail to capture how work is reorganized across people, tools, and time. This lag makes early abstractions especially durable, as new capabilities are interpreted using conceptual resources that are poorly suited to reasoning about brittleness, coordination, and lifecycle dynamics. As a result, accumulating evidence of downstream complexity does not necessarily undermine the myths that motivated adoption, but is instead absorbed into existing frames that treat such effects as peripheral or exceptional.

Another reason myths persist is that they align with longstanding ideals of rationalization and optimization that are deeply embedded in technical organizations. Much AI discourse assumes that better inference yields better decisions under stable criteria, whereas design theory emphasizes that criteria, constraints, and values are themselves negotiated and revised through framing and sensemaking [4, 6, 32, 40]. These myths are compelling in part because they make optimization-oriented narratives of AI compatible with organizational demands for defensible, “evidence-based” decision-making.

Finally, the persistence of myths is consistent with longstanding observations in automation and sociotechnical systems: new capabilities are rarely adopted solely because they work in isolation; they spread because they fit organizational needs and narratives, and because early benefits are visible while compensating costs are delayed or externalized [1, 3, 38, 41]. In design contexts, this implies that myths are not anomalies to be dispelled but recurring consequences of introducing partial abstractions into interpretive, evolving forms of work.

5.2 AI, Design Work, and Design Cognition

Much current work in AI-assisted design is structured around tools, tasks, or workflow phases (e.g., ideation systems, prompt-integrated prototyping tools, or collaborative co-creation interfaces) [27, 35, 44]. This literature has been valuable for identifying benefits and challenges in situated use. The synthesis here extends this work by examining how AI reshapes design work more broadly—using design cognition as one analytic reference point for understanding how designers frame situations, make sense of evolving constraints, and exercise judgment over time.

Several recent papers already point in this direction. Work on “designerly understanding” emphasizes that designers seek transparency not as technical explanation but as contextual insight that supports exploration and trade-off reasoning [23]. Studies show that supporting designers in understanding limits and failure modes can improve reasoning but requires additional structured reflection and interpretive work [25]. Workflow studies demonstrate that

different configurations of AI integration can produce different patterns of cognitive activity and perceived agency, underscoring that “AI in the loop” is not a single condition but a family of sociotechnical arrangements [51]. The myths-and-ironies account brings these observations into a more integrated picture: AI does not remove framing, sensemaking, or judgment, but changes their timing, visibility, and fragility.

This framing also helps connect AI-in-design research to the broader automation lineage. The patterned dynamics described here—e.g., recapture of “freed capacity” as throughput, redistribution of effort from production to evaluation, difficulty of monitoring and diagnosis, and task-complexity polarization—are direct descendants of classic ironies of automation [1, 9, 38]. However, in design contexts, these dynamics take distinct forms: the core challenge is less about real-time supervisory control and more about sustaining interpretive, value-laden work—framing, sensemaking, and judgment—under accelerating output and evolving constraints. In this sense, my argument complements human factors and productivity-focused accounts of AI-related ironies by specifying what is at stake in design work—and why the losses are often subtle, distributed, and difficult to defend.

Finally, the analysis here reinforces and extends a growing trend to treat AI as a “design material” that must be communicated, negotiated, and stabilized through representations and boundary objects [10, 11, 52]. It shows how those representational practices interact with cognitive and operational resources: explanation is not “free,” workflow integration is not “neutral,” and democratization changes what expertise *is*—governance, diagnosis, accountability—rather than simply lowering skill thresholds.

5.3 Implications for AI Tool Design

Existing literature on AI-assisted design includes promising directions for improvement—e.g., metacognitive support, workflow scaffolding, transparency mechanisms, and collaborative prompting practices [12, 14, 23, 27]. My analysis does not aim to replace these efforts; rather, it offers a lens for evaluating whether particular interventions support the forms of cognitive and coordinative work on which design practice depends.

First, the argument here suggests that AI tool design should prioritize *representational support for sensemaking* rather than post-hoc explanation alone. This aligns with calls for designerly transparency [23] and with tools that foreground failure modes and limits [25]. The key implication is that explanations should be evaluated by whether they help designers answer design-relevant questions—what is happening here, what matters, what might happen next—without imposing additional interpretive burden that competes with framing under time pressure and other organizational constraints. The fluency of natural language output makes this especially important: coherent explanations can increase perceived authority even when they do not improve understanding [9].

Second, tools should explicitly support *governance and judgment work*, not only generation. Several studies indicate that, under common adoption conditions, AI can shift designers toward evaluation and coordination [13, 51]. Toolmakers should therefore support

activities such as comparison across alternatives and framings, documentation of decision rationales and uncertainties, and communication of trade-offs to stakeholders. Work on communicating inherent algorithmic trade-offs provides a concrete precedent for this direction [53].

Third, workflow integration should be treated as a *sociotechnical intervention*, not merely a UX feature. Close coupling between AI and design tools can reduce local friction [27], but it can also stabilize work-as-imagined workflows that diverge from work-as-done as contexts change [19, 41]. The apparent success of a system under demonstration or early deployment conditions is often taken as evidence of readiness for broader deployment, with the assumption that extending it to handle the full range of complexities that emerge and change over a system’s life cycle will be straightforward [47]. In practice, this presumption underestimates the degree to which real design workflows involve shifting goals, evolving constraints, coordination across roles, and unanticipated breakdowns—precisely the conditions under which tightly coupled systems tend to become brittle. This suggests a design agenda oriented toward adaptability: supporting multiple rhythms of work, preserving space for reframing, and avoiding premature closure around AI-generated artifacts.

These implications can be operationalized as evaluative questions in design reviews and workflow assessments, rather than treated as requirements to be “solved” through tool features alone:

- **Turnaround and brittleness:** How much work is required to adapt the tool to a materially different design context (new domain, constraints, stakeholders), rather than a parametric variation of the same context?
- **Articulation burden:** What new coordination, justification, or verification work does the tool introduce, and who is responsible for it?
- **Framing slack:** Does the workflow preserve temporal space for reframing and reflection, or does it compress work into rapid generation and selection?
- **Authority and uncertainty:** How does the system represent limits and uncertainty, and how are these conveyed when outputs become shared boundary objects in teams?

5.4 Implications for Design Education and Practice

Recent studies already suggest that effective AI-assisted design requires new skills: prompting literacy, understanding limitations, managing uncertainty, and coordinating interpretations across teams [13, 14, 20]. The analysis here reframes these skills not as merely “new tool skills,” but as extensions of core design capabilities: framing, sensemaking, and judgment under evolving constraints. Recent HCI education research suggests that students already navigate these questions in practice, weighing when and how GenAI supports learning, how it affects pride and ownership, and where it risks short-cutting understanding rather than deepening it [15]. Educators are simultaneously experimenting with GenAI in UX studios, reporting changes in brainstorming and visualization workflows alongside concerns and barriers to curricular integration [45].

For design education, this implies that AI training should not be isolated as a technical module or tool tutorial. Students should be taught to recognize the myths that surround AI adoption and to anticipate the corresponding ironies. In practical terms, this means designing assignments that preserve and assess framing work (not only output), that require explicit articulation of assumptions and trade-offs, and that treat explainability as an object of critique rather than a guarantee of understanding. One practical mechanism is to treat AI disclosure not merely as compliance, but as a reflective design artifact that documents where AI influenced framing, alternatives considered, and judgment calls. Recent work in HCI design education argues for “AI contribution statements” and process-aware citation models precisely to support transparency and assessment by making students’ reasoning legible in AI-mediated workflows [37]. Instruction should also attend to the ways AI can redistribute responsibility and narrow opportunities to exercise certain forms of practice, even when artifact production becomes easier [1, 36].

A common pattern illustrates this shift: in teams where generative tools are introduced to accelerate early exploration, more concepts may be produced, but meetings can shift toward evaluating and justifying AI-generated artifacts under time pressure, with less opportunity to revisit framing decisions. In such situations, the apparent democratization of idea generation coincides with increased responsibility for sensemaking and accountability resting on a smaller number of designers.

For professional practice, this lens offers a vocabulary for negotiating AI adoption within organizations. Many of the harms described are not caused by AI “misuse” in a narrow sense, but by organizational dynamics that equate speed and volume with effectiveness and that treat AI outputs as objective baselines. Making framing and judgment visible—and defensible—becomes a central task for design leadership. The myths-and-ironies lens can support conversations about when to slow down, when to privilege exploration over throughput, how to document and communicate uncertainty, and how to distribute accountability when AI-generated artifacts shape decisions.

More broadly, this discussion complements trends in responsible AI and co-creative systems research by clarifying what it would mean to make AI supportive of design practice rather than merely productive. The challenge is not whether AI belongs in design, but how to configure sociotechnical arrangements—tools, workflows, evaluation practices, and educational approaches—that remain responsive as early abstractions harden and interpretive work is reorganized rather than eliminated.

6 Conclusion

This paper examined AI-assisted design through the lens of myths and ironies, arguing that many of the promises accompanying AI adoption—higher-level thinking, reduced cognitive load, simplified workflows, democratization, and objectivity—are best understood as partial abstractions of design work that become authoritative as technologies are institutionalized. As design practice resettles around these abstractions, the cognitive and coordinative work

they background—framing, sensemaking, articulation, and judgment—reasserts itself, often in less visible and more difficult-to-defend forms. The resulting ironies are not failures of implementation, but predictable consequences of introducing new capabilities into forms of work that are interpretive, evolving, and resistant to stabilization.

By situating these dynamics within design theory and the longer lineage of automation research, I have argued that AI’s most consequential effects in design are rarely confined to individual tools or tasks. Instead, they emerge through shifts in timing, visibility, authority, and responsibility—how work is evaluated, whose judgments count, and how accountability is distributed when AI-generated artifacts shape decisions. The contribution of this paper is to make these dynamics more legible: to offer a vocabulary for reasoning about AI adoption that moves beyond productivity metrics toward a clearer understanding of how design work and judgment are reorganized under persistent messiness.

References

- [1] Lisanne Bainbridge. 1983. Ironies of Automation. *Automatica* 19, 6 (1983). doi:10.1016/B978-0-08-029348-6.50026-9
- [2] Li-Yuan Chiou, Peng-Kai Hung, Rung-Huei Liang, and Chun-Teng Wang. 2023. Designing with AI: An Exploration of Co-Ideation with Image Generators. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1941–1954. doi:10.1145/3563657.3596001
- [3] Ruth Schwartz Cowan. 2023. *More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave*. Plunkett Lake Press. Google-Books-ID: PzCwEAAQBAJ.
- [4] Nigel Cross. 1982. Designerly ways of knowing. *Design Studies* 3, 4 (1982), 221–227.
- [5] Sidney Dekker, David Woods, and Martijn Mooij. 2002. Envisioned practice, enhanced performance: The riddle of future (ATM) systems. *International Journal of Applied Aviation Studies* 2, 1 (2002), 23–32.
- [6] Kees Dorst. 2015. Frame Creation and Design in the Expanded Field. *She Ji: The Journal of Design, Economics, and Innovation* 1, 1 (2015), 22–33. doi:10.1016/j.sheji.2015.07.003
- [7] Kees Dorst and Nigel Cross. 2001. Creativity in the design process: co-evolution of problem–solution. *Design Studies* 22, 5 (Sept. 2001), 425–437. doi:10.1016/S0142-694X(01)00009-6
- [8] Joseph Dunne. 1999. Professional judgment and the predicaments of practice. *European Journal of Marketing* 33, 7/8 (1999), 707–720. doi:10.1108/03090569910274339
- [9] Mica R. Endsley. 2023. Ironies of artificial intelligence. *Ergonomics* 66, 11 (Nov. 2023), 1656–1668. doi:10.1080/00140139.2023.2243404
- [10] K. J. Kevin Feng, Maxwell James Coppock, and David W. McDonald. 2023. How Do UX Practitioners Communicate AI as a Design Material? Artifacts, Conceptions, and Propositions. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2263–2280. doi:10.1145/3563657.3596101
- [11] K. J. Kevin Feng, Q. Vera Liao, Ziang Xiao, Jennifer Wortman Vaughan, Amy X. Zhang, and David W. McDonald. 2025. Canvil: Designerly Adaptation for LLM-Powered User Experiences. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3706598.3713139
- [12] Frederic Gmeiner, Kaitao Luo, Ye Wang, Kenneth Holstein, and Nikolas Martelaro. 2025. Exploring the Potential of Metacognitive Support Agents for Human-AI Co-Creation. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1244–1269. doi:10.1145/3715336.3735785
- [13] Frederic Gmeiner, Humphrey Yang, Lining Yao, Kenneth Holstein, and Nikolas Martelaro. 2023. Exploring Challenges and Opportunities to Support Designers in Learning to Co-create with AI-based Manufacturing Design Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3544548.3580999
- [14] Yuanning Han, Ziyi Qiu, Jiale Cheng, and RAY LC. 2024. When Teams Embrace AI: Human Collaboration Strategies in Generative Prompting in a Creative Design Task. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing*

- Systems (CHI '24). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3613904.3642133
- [15] Clint Heyer, Elisabet M. Nilsson, and Jens Pedersen. 2025. AI-assisted Learning in HCI Education: Opportunities and Dilemmas from a Student Perspective. In *Proceedings of the EduCHI 2025: 7th Annual Symposium on HCI Education*. ACM, Bloomington IN USA, 1–7. doi:10.1145/3742901.3742910
 - [16] Robert R. Hoffman, Shane T. Mueller, Gary Klein, Mohammadreza Jalaeian, and Connor Tate. 2023. Explainable AI: roles and stakeholders, desirments and challenges. *Frontiers in Computer Science* 5 (Aug. 2023), 1117848. doi:10.3389/fcomp.2023.1117848
 - [17] E. Hollnagel. 2003. The Role of Automation in Joint Cognitive Systems. *IFAC Proceedings Volumes* 36, 22 (Sept. 2003), 9–11. doi:10.1016/S1474-6670(17)37685-1
 - [18] Erik Hollnagel. 2015. Why is Work-as-Imagined Different from Work-as-Done? In *Resilient Health Care, Volume 2: The Resilience of Everyday Clinical Work*, Robert L. Wears, Jeffrey Braithwaite, and Erik Hollnagel (Eds.). Ashgate Studies in Resilience Engineering Series, Vol. 2.
 - [19] Erik Hollnagel and Sidney W. A. Dekker. 2024. The ironies of 'human factors'. *Theoretical Issues in Ergonomics Science* (Dec. 2024), 1–11. doi:10.1080/1463922X.2024.2443976
 - [20] Abidullah Khan, Atefeh Shokrizadeh, and Jinghui Cheng. 2025. Beyond Automation: How Designers Perceive AI as a Creative Partner in the Divergent Thinking Stages of UI/UX Design. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3706598.3713500
 - [21] P. Koopman and R.R. Hoffman. 2003. Work-arounds, make-work, and kludges. *IEEE Intelligent Systems* 18, 6 (Nov. 2003), 70–75. doi:10.1109/MIS.2003.1249172
 - [22] Tatiana Lau, Scott Carter, Francine Chen, Brandon Huynh, Everlyne Kimani, Matthew L Lee, and Kate A Sieck. 2024. Democratizing Design through Generative AI. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference (DIS '24 Companion)*. Association for Computing Machinery, New York, NY, USA, 239–244. doi:10.1145/3656156.3663703
 - [23] Q. Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. 2023. Designing Understanding: Information Needs for Model Transparency to Support Design Ideation for AI-Powered User Experience. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3544548.3580652
 - [24] Vivian Liu, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2023. DALL-E: Integrating Text-to-Image AI in 3D Design Workflows. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1955–1977. doi:10.1145/3563657.3596098
 - [25] Steven Moore, Q. Vera Liao, and Hariharan Subramonyam. 2023. fAllureNotes: Supporting Designers in Understanding the Limits of AI Models for Computer Vision Tasks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3544548.3581242
 - [26] Harold G. Nelson and Erik Stolterman. 2012. *The Design Way: Intentional Change in an Unpredictable World* (second ed.). The MIT Press, Cambridge, Massachusetts; London, England.
 - [27] Savvas Petridis, Michael Terry, and Carrie J Cai. 2024. PromptInfuser: How Tightly Coupling AI and UI Design Impacts Designers' Workflows. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 743–756. doi:10.1145/3643834.3661613
 - [28] Han Qiao, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2025. To Use or Not to Use: Impatience and Overreliance When Using Generative AI Productivity Support Tools. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3706598.3714103
 - [29] Peinuan Qin, Chi-Lan Yang, Jingshu Li, Jing Wen, and Yi-Chieh Lee. 2025. Timing Matters: How Using LLMs at Different Timings Influences Writers' Perceptions and Ideation Outcomes in AI-Assisted Ideation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3706598.3713146
 - [30] Zainab Salma, Raquel Hijón-Neira, and Celeste Pizarro. 2025. Designing Co-Creative Systems: Five Paradoxes in Human-AI Collaboration. *Information* 16, 10 (Oct. 2025), 909. doi:10.3390/info16100909
 - [31] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW Seriously: Supporting Articulation Work. In *Computer Supported Cooperative Work (CSCW)*, Vol. 1. 7–40.
 - [32] Donald A. Schön. 1983. *The Reflective Practitioner: How Professionals Think In Action*. Basic Books.
 - [33] Donald A. Schön. 1988. Designing: Rules, types and worlds. *Design Studies* 9, 3 (1988), 10.
 - [34] Donald A. Schön. 1992. Designing as a Reflective Conversation with the Materials of a Design Situation. *Research and Engineering Design* 3 (1992), 131–147. doi:10.1007/BF01580516
 - [35] Yang Shi, Tian Gao, Xiaohan Jiao, and Nan Cao. 2023. Understanding Design Collaboration Between Designers and Artificial Intelligence: A Systematic Literature Review. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (Sept. 2023), 1–35. doi:10.1145/3610217
 - [36] Prakash Shukla, Phuong Bui, Sean S Levy, Max Kowalski, Ali Baigelenov, and Paul Parsons. 2025. De-skilling, Cognitive Offloading, and Misplaced Responsibilities: Potential Ironies of AI-Assisted Design. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–7. doi:10.1145/3706599.3719931
 - [37] Prakash Shukla, Suchismita Naik, Ike Obi, Jessica Backus, Nancy Rasche, and Paul Parsons. 2025. Rethinking Citation of AI Sources in Student-AI Collaboration within HCI Design Education. In *Proceedings of the EduCHI 2025: 7th Annual Symposium on HCI Education*. ACM, Bloomington IN USA, 1–8. doi:10.1145/3742901.3742909
 - [38] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2024. Ironies of Generative AI: Understanding and mitigating productivity loss in human-AI interactions. doi:10.48550/arXiv.2402.11364 arXiv:2402.11364 [cs].
 - [39] Niya Stoimenova and Rebecca Price. 2020. Exploring the Nuances of Designing (with/for) Artificial Intelligence. *Design Issues* 36, 4 (Sept. 2020), 45–55. doi:10.1162/desi_a_00613
 - [40] Erik Stolterman. 2008. The Nature of Design Practice and Implications for Interaction Design Research. *International Journal of Design* 2, 1 (2008), 55–65.
 - [41] Barry Strauch. 2018. Ironies of Automation: Still Unresolved After All These Years. *IEEE Transactions on Human-Machine Systems* 48, 5 (Oct. 2018), 419–433. doi:10.1109/THMS.2017.2732506
 - [42] Anselm Strauss. 1988. The Articulation of Project Work: An Organizational Process. *The Sociological Quarterly* 29, 2 (June 1988), 163–178. doi:10.1111/j.1533-8525.1988.tb01249.x
 - [43] Mathias Peter Verheijden and Mathias Funk. 2023. Collaborative Diffusion: Boosting Designers' Co-Creation with Generative AI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3544549.3585680
 - [44] Wen-Fan Wang, Chien-Ting Lu, Nil Ponsa i Campanyà, Bing-Yu Chen, and Mike Y. Chen. 2025. AIdation: Designing a Human-AI Collaborative Ideation System for Concept Designers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–28. doi:10.1145/3706598.3714148
 - [45] Joshua Weinberg and Monica Chan. 2025. Experimenting with Generative AI Tools and their Implications: Insights from High School UX Educators. In *Proceedings of the EduCHI 2025: 7th Annual Symposium on HCI Education*. ACM, Bloomington IN USA, 1–6. doi:10.1145/3742901.3742912
 - [46] David Woods and Sidney Dekker. 2000. Anticipating the effects of technological change: A new era of dynamics for human factors. *Theoretical Issues in Ergonomics Science* 1, 3 (Jan. 2000), 272–282. doi:10.1080/14639220110037452
 - [47] David D. Woods. 2016. The Risks of Autonomy: Doyle's Catch. *Journal of Cognitive Engineering and Decision Making* 10, 2 (June 2016), 131–133. doi:10.1177/1555343416653562
 - [48] David D. Woods. 2024. Limits of Automata—Then and Now: Challenges of Architecture, Brittleness, and Scale. *Journal of Cognitive Engineering and Decision Making* 18, 4 (Dec. 2024), 394–401. doi:10.1177/15553434241240203
 - [49] David D. Woods and Erik Hollnagel. 2006. *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. CRC Press.
 - [50] Peter Wright, Andy Dearden, and Bob Fields. 2000. Function allocation: a perspective from studies of work practice. *International Journal of Human-Computer Studies* 52, 2 (Feb. 2000), 335–355. doi:10.1006/ijhc.1999.0292
 - [51] Xiaotong (Tone) Xu, Arina Konnova, Bianca Gao, Cindy Peng, Dave Vo, and Steven P. Dow. 2025. Productive vs. Reflective: How Different Ways of Integrating AI into Design Workflows Affect Cognition and Motivation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3706598.3713649
 - [52] Nur Yildirim, Alex Kass, Teresa Tung, Connor Upton, Donnacha Costello, Robert Giusti, Sinem Lacin, Sara Lovic, James M O'Neill, Rudi O'Reilly Meehan, Eoin Ó Loideáin, Azzurra Pini, Medb Corcoran, Jeremiah Hayes, Diarmuid J Cahalane, Gaurav Shivhare, Luigi Castoro, Giovanni Caruso, Changhoon Oh, James McCann, Jodi Forlizzi, and John Zimmerman. 2022. How Experienced Designers of Enterprise Applications Engage AI as a Design Material. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3491102.3517491
 - [53] Bowen Yu, Ye Yuan, Loren Terveen, Zhiwei Steven Wu, Jodi Forlizzi, and Haiyi Zhu. 2020. Keeping Designers in the Loop: Communicating Inherent Algorithmic Trade-offs Across Multiple Objectives. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, New York, NY, USA, 1245–1257. doi:10.1145/3357236.3395528