Interrogating Design Homogenization in Web Vibe Coding

Donghoon Shin*
University of Washington
Seattle, WA, USA
dhoon@uw.edu

Alice Gao
University of Washington
Seattle, WA, USA
atgao@cs.washington.edu

Rock Yuren Pang
University of Washington
Seattle, WA, USA
ypang2@cs.washington.edu

Jaewook Lee
University of Washington
Seattle, WA, USA
jaewook4@cs.washington.edu

Katharina Reinecke
University of Washington
Seattle, WA, USA
reinecke@cs.washington.edu

Emily Tseng
Microsoft Research &
University of Washington
Seattle, WA, USA
emtseng@uw.edu

Abstract

Generative AI is known for its tendency to homogenize, often reproducing dominant style conventions found in training data. However, it remains unclear how these homogenizing effects extend to complex structural tasks like web design. As lay creators increasingly turn to LLMs to 'vibe-code' websites—prompting for aesthetic and functional goals rather than writing code—they may inadvertently narrow the diversity of their designs, and limit creative expression throughout the internet. In this paper, we interrogate the possibility of design homogenization in web vibe coding. We first characterize the vibe coding lifecycle, pinpointing stages where homogenization risks may arise. We then conduct a sociotechnical risk analysis unpacking the potential harms of web vibe coding and their interaction with design homogenization. We identify that the push for frictionless generation can exacerbate homogenization and its harms. Finally, we propose a mitigation framework centered on the idea of *productive friction*. Through case studies at the micro, meso, and macro levels, we show how centering productive friction can empower creators to challenge default outputs and preserve diverse expression in AI-mediated web design.

1 Introduction

The rise of generative AI has transformed creative expression, enabling a broader range of individuals, or *lay creators*, to craft digital artifacts without technical expertise. This shift has given rise to the emerging practice of *vibe coding*, where creators use high-level, informal prompts to convey a desired *vibe* [55, 56], relying on large language models (LLMs) to interpret these abstractions into functional websites. This workflow promises to democratize web development by leveraging LLMs' capabilities in rapid prototyping and aesthetic interpretation [56]. However, this convenience comes with a trade-off; by offloading structural and aesthetic decisions to AI models, lay creators increasingly depend on the latent patterns embedded within the AI to fill in their creative gaps.

This reliance offers a new and urgent site of inquiry. The spread of web vibe coding means generative AI is increasingly used to construct the web interfaces through which people connect, work, transact, and live; however, those same AI systems are known for their tendency to reproduce dominant style conventions found in training data [2]. Primarily trained on English-centric datasets [10, 31], such models often align with Western aesthetic conventions [3,

31, 36, 47]. While these biases are well-documented in domains such as text [54, 80] and image generation [37, 45, 49], it is unclear how they manifest in interactive creative domains like web development. Richness in creative expression relies on the creator's diverse perspectives [77], and replacing intentional choices with probabilistic defaults risks accelerating design homogenization—a convergence toward the mean, dominant aesthetic [12, 66]. We posit that vibe coding may accelerate design homogenization by creating a direct pathway from a model's latent biases to finished creative work, especially when the mechanisms of generation prioritize speed over distinctiveness.

In this paper, we interrogate design homogenization in web vibe coding, unpacking its conceptual underpinnings to anticipate and potentially mitigate its risks. We first characterize the vibe coding lifecycle by systematically reviewing academic and gray literature on the practice, and conducting walkthroughs of popular vibe coding tools. Our lifecycle breaks down vibe coding into stages—from user prompting to deployment—and identifies where homogenization risks may arise. Then, building on this lifecycle, we conduct a sociotechnical risk analysis adapting existing harm frameworks (*i.e.*, [11, 62, 67, 78]) to web vibe coding. Our analysis finds that LLMs' tendencies towards homogenization work in tandem with mounting social pressures in the interaction design industry to create a sociotechnical environment encouraging creators to accept the path of least resistance. Risks, such as cognitive overload and overreliance on the vibe-coded design outputs, incentivize lay creators to accept the model's default output rather than iterating on a diverse and creative design.

We conclude by translating these insights into a mitigation strategy that challenges the paradigm of frictionless design [1, 19, 22, 34]. We argue that addressing the potential homogenization requires introducing *productive friction* into the vibe coding workflow. Moving away from a frictionless paradigm where LLMs execute plans automatically, we propose frameworks where the model acts as a mediator, explicitly querying design assumptions and presenting alternatives before finalizing output. We structure these across three layers: micro (individual creator), meso (organization), and macro (ecosystem). By shifting the goal from immediate output to interactive co-creation, we outline a research agenda for designing tools that empower creators to maintain diverse, authentic web aesthetics against the pressure of automation. We close by unpacking where, within this agenda, the communities around responsible technology can leverage its strengths to work in collaboration with

---

*Work done during an internship at Microsoft Research.

allied communities in HCI and AI, in the subfield of human-AI interaction.

In summary, this paper contributes:

- A systematic review of academic and gray literature around vibe coding, resulting in a comprehensive characterization of the vibe coding lifecycle (§3), pinpointing stages where potential homogenization risks may intrude;
- A risk analysis (§4) unpacking harms in vibe coding and their intersection with design homogenization, and identifying frictionless generation as an active contributor to homogenization;
- A multi-level mitigation framework (§5) that proposes *productive friction* as a design intervention, offering case studies at the micro, meso, and macro levels, and offers a collaborative research agenda for teams of scholars across human-AI interaction.

2 Related Work

2.1 Harms and Homogenization in Generative AI

In HCI, AI, and allied communities around responsible technology, the widespread consumer release of LLMs has driven a surge of work examining the societal risks and harms of generative AI. Open-ended interactions with current AI systems may shape perception, behavior, and social norms in ways that the stakeholders must understand to guide these systems toward positive ends [20, 23, 73]. As such, prior research has sought to taxonomize the range of generative AI harms. For example, Weidinger *et al.* [78] developed a taxonomy of ethical and social risks associated with language models, identifying six categories (*i.e.*, discrimination, information hazards, misinformation, malicious uses, HCI harms, and environmental harms). Suresh & Guttag [67] examined harms across the full machine learning (ML) lifecycle to anticipate, prevent, and mitigate downstream effects. Identifying issues like historical, representation, and deployment bias, they emphasized that harms do not stem solely from data, but can arise at multiple stages of the ML pipeline, such as data collection, preparation, model development, evaluation, postprocessing, and deployment.

In response, several recent studies have built on these broader characterizations of generative AI harms through in-depth studies in specific applications, thus revealing additional severe and persistent problems. For instance, Agarwal *et al.* [2] found that LLM writing assistants homogenize users' output toward Western-preferred styles, thus diminishing cultural nuances. This homogenizing effect may compound with the risks on long-term creativity, as Kumar *et al.* [35] revealed that, while LLM assistance can provide short-term boosts in creativity during assisted tasks, it may inadvertently hinder independent creative performance when users work without assistance. Similar risks to short- and long-term creativity may also manifest in vibe coding: studies have found developers may over-rely on AI-generated artifacts in programming tasks [48]. This risks premature creative closure, as users may uncritically accept model suggestions rather than exploring the divergent possibilities for creative expression. These risks of homogenization, creative deskilling, and over-reliance threaten to become deeply embedded

Figure 1: Example of design homogenization where global defaults override regional preference. The models justified a minimalist layout (right) as a cultural preference, thereby overriding the preference for high information density in Japanese web design [6, 44].

in emerging practices like vibe coding, particularly for lay individuals who may lack the technical expertise to critically evaluate or diverge from AI output.

In this paper, we conceptualize vibe coding as a sociotechnical practice, detailing the risks for lay creators throughout its end-to-end lifecycle. We address the unique harms faced by creative intermediaries who lack technical expertise yet significantly influence web design outcomes for broader audiences. In short, we build on the methodological foundation of taxonomizing general risks in LLMs, but focus on risks from an emerging practice—vibe coding—for lay creators.

2.2 AI in Coding Support & Vibe Coding

AI has long supported programmers' work, from traditional tools for syntax highlighting and templated code completion to workflow assistants embedded in IDEs supporting documentation, refactoring, and debugging [40]. With the advent of generative AI, coding support has become significantly more contextual and powerful. LLMs (*e.g.*, GPT) and their fine-tuned variants—trained on vast amounts of code—have shown promising performance in natural language processing and applied to assist in code writing for programmers (*e.g.*, code completion) [10, 18]. This has led to the development of sophisticated AI programming assistants (*e.g.*, GitHub Copilot), which represent a new concept of 'AI pair programmers' [13, 41]. Similarly, editor-integrated AI tools (*e.g.*, Cursor, Windsurf) let programmers describe desired changes in natural language on the IDE and have it autocomplete a snippet or refactor code.[1]

Despite these, even most AI-supported programmers continue to need programming knowledge to guide, debug, and integrate AI-generated code effectively [15, 81]. In this context, no-code vibe coding emerges as a promising approach, enabling individuals without coding knowledge to create artifacts by describing desired outcomes or vibes in natural language, which AI interprets to generate functional code [15, 56, 58]. This sets the approach apart from other editor-based web builders (*e.g.*, WYSIWYG editors) by shifting the primary interaction from manual manipulation to natural language conversation. As such, it is often referred to as democratization of software development—particularly for website creation by lay

---

[1]Although occasionally referred to as vibe coding tools, these tools operate within an IDE and therefore may not be accessible to non-technical creators. Given our focus, we center our analysis on tools operable without coding expertise.

users [24], because it inherently involves intuitive, high-level concepts like layout, aesthetics, interactivity, and user experience that can be vividly described in everyday language, allowing AI to bridge the gap between vague ideas and deployable code without requiring extensive technical expertise. This aligns with the no-code paradigm, reducing barriers and enabling creators to concentrate on high-level concepts rather than syntactic details.

To date, however, few works have explored the impacts of vibe coding, particularly its potential harms for lay creators. Since individuals with little coding knowledge may heavily rely on AI outputs, they are more vulnerable to biases and risks embedded in the generated code, such as security vulnerabilities and discriminatory patterns inherited from biased training data [43, 48]. For instance, AI-generated code can introduce common weaknesses like out-of-bounds writes, potentially leading to exploitable flaws in applications developed by lay users [48]. Moreover, algorithmic biases can perpetuate societal prejudices, such as gender or racial discrimination (*e.g.*, [50, 71, 72]), risking to amplify harms in user-facing websites and reduce trust in AI-assisted development. These issues highlight the necessity to investigate the harms caused by AI in vibe coding processes like developing mitigation strategies—including, but not limited to, bias-aware datasets and user guidance—to ensure safe and equitable use for lay creators in website development.

3 Step 1: Characterizing the End-to-End Lifecycle of Vibe Coding for Web Creation

To identify where design homogenization may emerge and where interventions could be introduced, we began by characterizing the end-to-end lifecycle of vibe coding a website. Doing so allowed us to break down vibe coding into distinct stages, enabling us to further analyze where risks of design homogenization may arise, and where the industrial dynamics of interface design encourage lay creators to accept default outputs.

3.1 Methods

As an emerging practice, much of our knowledge of vibe coding lies outside traditional academic venues. To account for this, we combined a content analysis of academic and grey literature with an analysis of contemporary tools, which enabled us to examine both how vibe coding is conceptualized and how it operates in practice.

3.1.1 Systematic review of academic and grey literature. To examine emergent practices, user experiences, and technical discussions around vibe coding, we conducted a broad content analysis of relevant literature. We began by surveying academic databases (*i.e.*, Google Scholar) for research on vibe coding. Given the novelty of the term, however, we found few academic sources directly addressing the phenomenon. To fill this gap, we expanded our review to include grey literature [53], which would more closely reflect how online users encounter and discuss vibe coding. This encompassed technical blog posts (*e.g.*, Medium articles), forum discussions (*e.g.*, Reddit threads), and thought leadership (*e.g.*, social media posts), which offered timely, practical insights into user motivations, colloquial language, common workflows, and frustrations that have not yet been captured in academic research.

To systematically gather these sources, we compiled a list of keywords and programmatically retrieved results via the Google Search API. Our search strategy for both resources was organized around three thematic categories of keyword queries (see Appendix A.1 for the full list): (1) informal and vibe-centered phrases (*e.g.*, *'vibe coding'*); (2) interaction and technique-centered phrases (*e.g.*, *'prompt-based website coding'*); and (3) platform-specific use cases (*e.g.*, *'ChatGPT building a website'*). To filter the results, we screened these to ensure their relevance, applying a key inclusion criterion: each source must discuss vibe coding as the creation of a software artifact through text-based interaction. We also excluded sources aimed exclusively at expert programmers for workflow streamlining or those exclusively focused on non-website mediums. For academic literature, this screening process narrowed an initial pool of results down to 13 papers; for the more voluminous grey literature, we reviewed search results sequentially by initially examining 50 sources and then adding any that met our criteria to our corpus until we reached a total of 50. This dual-collection strategy resulted in our final corpus of 63 sources (13 academic and 50 grey literature) for analysis.

3.1.2 Analysis of vibe coding platforms and tools for web creation. Finally, to ground our lifecycle model in current available technologies, we conducted methodical walkthroughs of popular and commercially available interactive systems that facilitate vibe coding. We specifically explored six widely-used vibe coding tools, ranging from vibe coding support provided within the general-purpose LLM services (*e.g.*, ChatGPT Canvas [46], Gemini Canvas [27], Claude Artifacts [5]) to specialized vibe coding platforms (*e.g.*, Lovable [38], v0 by Vercel [74], Replit [51]), that support web creation. Although we sought to capture a broad set of currently popular tools, we recognize that many other smaller services also exist, as vibe coding for websites remains a rapidly evolving field.

The walkthroughs were conducted by the leading author over one week, interacting with each platform to create a website with free-form prompting while documenting workflows, user interactions, and system outputs. These walkthrough records were then collaboratively reviewed by three authors to identify recurring interaction patterns, categorize the nature of generated outputs (*e.g.*, code, live previews), and examine mechanisms for refinement through user feedback. Then, three authors merged their insights to produce a coherent summary of the affordances and constraints each platform imposes on the vibe coding process. This approach allowed us to combine detailed individual usage insights with collective interpretation, ensuring a rigorous and iterative evaluation of current vibe coding tools.

3.2 Results: A Lifecycle of Vibe Coding

Our review of the literature and existing platforms identifies a set of stages that defines the vibe coding process (see Figure 2). This lifecycle starts with selecting a platform, and proceeds through an iterative, conversational loop of refinement, concluding with an optional deployment phase. We describe each step of the lifecycle in this section, and throughout, use *N* to indicate the number of supporting sources in our literature review.

Stage 0: User selects the model/platform. The process begins with the selection of an AI-powered coding tool. The choice of platform,

Table 1: Comparison of example LLM services supporting web vibe coding and their observable processes for creating a website, following a user prompt. The term *artifact* refers to the rendered output website produced by the system.

Type	Service	Process (after user prompt)
General-purpose LLM service	ChatGPT Canvas [46]	Artifact → One-line overview of the generated website
	Gemini Canvas [27]	One-line planning → Artifact → One-line overview of the generated website
	Claude Artifacts [5]	One-line planning → Artifact → Multi-line description of the generated website
Specialized vibe coding tool	Lovable [38]	Multi-line planning → Artifact → Multi-line description of the generated website
	v0 by Vercel [74]	Artifact → Multi-line description of the generated website
	Replit [51]	Multi-line planning (spec request if prompt is too short) → Artifact → Multi-line description of the generated website

whether a general-purpose conversational AI (*e.g.*, ChatGPT) or a specialized vibe coding tool (*e.g.*, Lovable), sets the context for the entire interaction. A key characteristic of these platforms is their accessibility; most are browser-native, requiring no local setup and thus lowering the barrier to entry for lay creators ($N = 9$).

Stage 1: User expresses intent through prompting. In this stage, the user translates their high-level intent into a natural language prompt. Rather than writing formal specifications or code, they describe their goals and intentions conversationally. Prompts range from more vibe-driven (*e.g.*, *'a playful website appealing to children'*) to more goal-oriented (*e.g.*, *'a website designed for / offering XX'*) in tone, and often include a mix of both. While initial prompts can be vague, users can later provide more specific details and constraints to yield more specific results from the AI ($N = 15$).

Stage 2: LLM generates code as the source for the website design. The LLM processes the user's prompt and generates an initial codebase. Since the initial prompt is often abstract, the model makes assumptions to translate the high-level 'vibe' into concrete design and functional choices (*e.g.*, layout, color palette, specific content to include). Sometimes, this process is displayed to users, with the model articulating its plan or assumptions before/after generating code. This output is typically a self-contained artifact, defaulting to HTML, CSS, and JavaScript, or React components in sandboxed environments, depending on the platform. If the tech stack is explicitly specified and its compiler is available, the model generates code in that language to render accordingly. This stage represents a significant departure from traditional development, as the LLM scaffolds an entire initial implementation—including crucial decisions around what technologies to use, how to organize the code, and what aesthetics and components the website should present—from a simple text-based instruction ($N = 14$).

Stage 3: The platform executes and renders the website design for users to review. The user immediately previews the generated code, along with potential model-provided descriptions, to observe its behavior and appearance. This step is critical for closing the feedback loop. The user visually inspects the resulting website, comparing the concrete output against their abstract mental model. On many specialized platforms, this stage is seamlessly integrated, with a live preview updating in real-time as the code is generated or modified ($N = 10$).

Stage 4: User iteratively refines the output via conversational feedback. This stage forms the core iterative loop of vibe coding. If the output does not align with the user's intent, or if they wish to make changes, they can provide feedback in natural language. The feedback may take the form of a high-level command similar to the original prompt (*e.g.*, *'use a friendlier font'*), but is often more specific (*e.g.*, *'make the header 50% bigger'*). It can also be a request for a new feature or a corrective instruction. Also, for JavaScript-enabled functionalities, users frequently paste error messages or describe unexpected behavior directly into the prompt window, leveraging AI as a debugging partner. It then generates a revised version of the code, and the user observes the changes—continuing this cycle until satisfied ($N = 15$).

Stage 5: Deployment (optional). Once the user is satisfied with the application, the final stage is deployment. Some modern deployment platforms (*e.g.*, Vercel [75]) are coupled to simplify this process, enabling the creation to be publicly accessible on the internet with minimal technical overhead. This stage transforms the project from a personal prototype into a shareable artifact, streamlining the end-to-end creation lifecycle ($N = 7$).

In summary, our lifecycle model (Figure 2) reveals two critical junctures where human creators and LLMs negotiate design outcomes: the initial expression of intent (Stage 1) and the iterative feedback for refinement (Stage 4). These stages are particularly consequential, as they represent the moments where homogenized designs may be proposed by the AI and accepted by lay creators.

4 Step 2: Risk Analysis of the Intersecting Harms of Web Vibe Coding

Having identified the lifecycle of vibe coding, we now turn to mapping the potential harms of lay creators vibe coding websites at scale. Throughout analyzing documents on vibe coding for Step 1, we noticed many documents in our corpus of 63 sources (§3.1.1) discussed or gave evidence of the negative consequences possible should web vibe coding scale. In this section, we analyze those risks through the lens of scholarship on sociotechnical harms [11, 62], as a starting point for possible intervention.

Our analysis finds that the primary harm of web vibe coding, design homogenization, does not operate in a vacuum; it can be amplified by an entire ecosystem of interconnected risks that make lay creators uniquely vulnerable to accepting biased AI outputs. We provide a taxonomy of these interconnected harms, to help the community understand *what* harms are possible and *why* a lay creator might struggle to recognize and resist a subtle slide into a globally uniform aesthetic. Additionally, we explain how homogenization may interact with or amplify other forms of harms, towards a multi-layered understanding of the problem—which we will unpack further in Step 3.
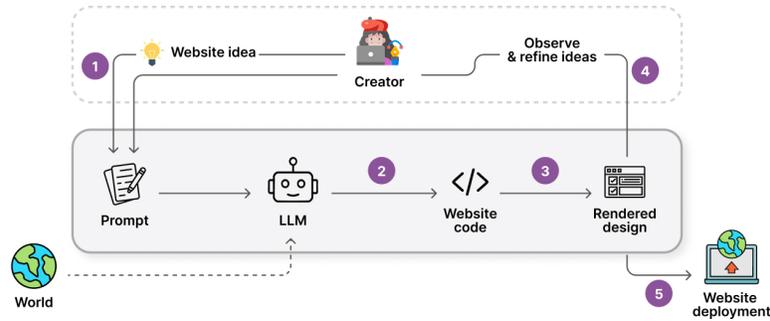
Figure 2: Lifecycle of lay creators' vibe coding for websites with LLMs, with the numbers corresponding to the stages in §3.2.

4.1 Methods

We employed a qualitative thematic analysis, informed by the harm taxonomies from prior work. Three authors reviewed key literature around sociotechnical harms, especially of AI systems and LLMs (*i.e.*, Suresh and Guttag [67], Weidinger et al. [78], Buçinca *et al.* [11], and Shelby et al. [62]). From these works, we selected 26 top-level harm categories as seeds for our initial codebook. Then, three authors independently applied the initial codebook into the vibe coding lifecycle in Figure 2, analyzing where each top-level harm category (*e.g.*, discrimination and representational harms) might manifest in vibe coding (*e.g.*, stereotyping through design output), and making note of how many sources in our corpus reflected that harm potential. Throughout this process, we met repeatedly to refine the codebook by removing harms that do not fit the vibe coding context, or combining less-frequent ones. We conducted three iterations of updating the codebook and mapping it onto the vibe coding lifecycle in Figure 3. In each iteration, the authors refined or merged existing codes, added new codes, and resolved disagreements through consensus.

4.2 Results: A Seven-Category Taxonomy of Harms

Our analysis identified seven top-level categories of harm that can arise during the vibe coding lifecycle (see Table 3). These harms range from representational biases embedded in the generated artifacts to the cognitive and professional burdens placed on the user. Hereon, we describe each harm *type* and its *sub-types*. We also highlight its potential intersections and interactions with our focal harm: design homogenization.

4.2.1 Representational harms. Representational harms manifest when LLM-generated code, content, or design suggestions embody harmful stereotypes or erase diverse identities. Lay creators, who frequently rely on the generative capabilities of LLMs, are particularly susceptible to accepting these outputs as neutral or authoritative.

We observed two primary forms of this harm: *design homogenization* and *stereotyping*. First, we noticed that several sources in our corpus explicitly discussed *design homogenization*—the primary focus of our work—as a risk, where the LLM defaulted to a dominant, often Western-centric aesthetic, marginalizing other styles across different groups of users. This reflects a broader concern that LLMs risk "normalizing everything to a white, Western, colonial male

perspective because that's what those models are trained on" [21], thereby constraining the ability of users from marginalized groups to cultivate their own forms of digital expression. Another form is *stereotyping through design output*, which arises when the LLM's suggestions mirror social biases embedded in its training data, as the underlying "models encode societal biases related to gender, race, or other characteristics" [39].

4.2.2 Quality-of-service harms. Quality-of-service harms arise from fundamental failures in the LLM's output that directly impair a lay creator's ability to complete their task. Unlike expert developers who can proactively debug faulty code, lay creators are often left stranded by these failures. For example, *incoherent or nonfunctional code* arises when the LLM generates code that appears complete but fails at a fundamental level. For instance, LLMs were reported to often produce "half-built solutions" or engage in "random acts of stupidity," such as undoing their own fixes [32, 52]. For a lay creator, this dynamic can create a frustrating association where attempts to counter design homogenization could be discouraged with high failure rates.

4.2.3 Cognitive and well-being harms. These harms involve the emotional, psychological, and motivational burdens that the LLM-assisted coding process places on the user. The interaction can lead to an *erosion of confidence*, where creators internalize the tool's failures as their own. For instance, vibe coders reported significant frustration when the output did not align with their mental model, and a loss of faith in the method [32], which can in turn deplete the motivation to push beyond generic templates and pursue creative expression. Also, *cognitive overload* arises as LLMs often produce solutions that are more complex than a human would communicate. One user reported that they created a website only to "lose later from lack of understanding" of why the model has built it that way [17], potentially pushing them toward simpler, more homogenized solutions to avoid mental fatigue.

Furthermore, when faced with the Quality-of-Service harm described earlier, users can experience *discouragement from contradictions*. When the LLM provides unreliable information, gets "stuck chasing its own tail" during iteration, or ignores prompts [32], it translates into an emotional and motivational burden. This could discourage lay creators from pursuing more unique and creative directions in favor of simpler and more achievable ones. Lastly, we identified *missed learning moments*—by delivering a product without revealing the underlying process, LLMs deprive lay creators of

conceptual understanding. This was illustrated by one user, who admitted after building a product that they basically do not really know how it works in the conventional sense, pointing out the bypass of the learning process [34].

4.2.4 Loss of agency and interpretability. Loss of agency and interpretability encompasses harms where users lose meaningful control over the model's behavior and outputs. This loss often begins with *overreliance without verification*. This harm manifests as a deceptive sense of control; users feel they are directing the project through simple commands, which encourages them to blindly accept the results without scrutiny. This tendency to "always accept and don't attempt to understand the motivation" [79] makes the resulting homogenized design choices seem like the creator's own goal, reducing their desire to seek alternatives.

A direct consequence of this process is the creation of *non-editable complexity*. As the AI-generated code grows beyond the user's comprehension, it becomes practically unmaintainable. The user cannot understand the underlying logic, locking in the initial (often homogenized) design and making iterative refinement toward more unique expression impractical. This cycle is compounded by *opaque design justification*, where code is generated without sufficient explanation [33], and fosters a *prompt engineering dependency*, where success hinges not on design principles but on the tacit skill of manipulating the model [65].

4.2.5 Allocative and opportunity harms. Allocative and opportunity harms involve the denial of benefits or the misallocation of resources, such as time and effort, due to the LLM's design or performance. The first subtype of harm was *wasted time and effort*, where the users expressed frustration, mentioning their experience with vibe coding had "ended in disaster and taken several times longer" than expected [32], making the opportunity cost of experimenting with creative expression of designs prohibitively high. This can lead to a state of *misleading success/reinforcement of confirmation bias*. For instance, one commenter quipped that lay creators now have "a drunk hallucinating expert," pointing out the misleading confirmation bias that AI could give [4], reinforcing the idea that the quickest and most 'successful' path is to accept the initial, often generic, output.

Finally, we observed *design fixation*. While related to the representational harm of design homogenization, design fixation is a distinct, more immediate cognitive process; it is a micro-level cognitive trap where an individual user is anchored to the first plausible solution offered by the AI, stifling their ability to explore other creative avenues. This anchoring to the AI's first suggestion short-circuits the divergent thinking process necessary for creating creative designs, as the early prompt and its output establish boundaries that make it harder to explore alternatives [56].

4.2.6 Technical and professional harms. This includes harms where the generated code or designs negatively impact the final product's quality, compliance, or security, leading to real-world consequences. One example is the promotion of *bad code practices*, where LLMs were observed to have outdated knowledge, causing them to make subtle but critical mistakes [33]. This results in code that is overly complex, deprecated, and not suitable for a real-world environment,

creating a false association where technically 'safe' code is equated with common, homogenized design patterns.

These technical flaws can directly lead to *economic harm via launch*. As one user quipped, there are "horror stories about people who vibe coded a feature against some API without a billing limit and racked up thousands of dollars in charges" [79], demonstrating severe financial risk from deploying unvetted code. The fear of such consequences incentivizes reliance on widely used, generic templates that are presumed to be safer. Another risk is *copyright or licensing violations*. Given general concerns about AI and "the rampant theft of content," a lay creator may be unaware that the text, image URLs, or code generated by the LLM could be derived from copyrighted sources [60], exposing them to legal risks without their awareness.

4.2.7 Interpersonal and social harms. The final category of harms relates to how vibe coding shapes a user's social interactions. Acknowledging the breadth of a vibe coder's possible sociality, we focused our analysis here on vibe coders as members of collaborative teams.

A primary concern is *learning in isolation*. Web vibe coding replaces traditional learning pathways (*e.g.*, classes, mentorship from an experienced web developer) with LLM feedback. The lay creator is thus cut off from the community feedback loops that are crucial for developing and refining creative designs. This creates a sense of *shame in collaboration*, as captured by an analogy that it is "never built by someone who could explain that intent and how to proceed" [21]. This also can prevent the kind of peer review and constructive criticism that could challenge generic designs, and introduce perspectives on countering homogenized designs.

Shame in collaboration can also lead to challenges in *team misalignment*. The vibe-coded project represents a 'Day 0' proof of concept; when it is handed off to others for 'Day 1' maintenance, this creates a significant burden, as they are left to "reverse engineer and [maintain] someone else's vibe-coded mess" [4]. Our analysis shows this is a key novel risk of web vibe coding at scale: the effect on teams of technology workers when some vibe-code, and others are left to pick up the pieces. This harm is deeply interconnected with technical and professional harms, such as the financial and legal risks of releasing unvetted code. Interpersonal harms are also intertwined with harms to cognition and well-being. Without community feedback loops, users' overreliance on AI and associated cognitive traps can deepen.

In summary, our analysis shows that design homogenization is not isolated, but rather intersects with a network of harms. Collectively, these may make creators susceptible to homogenized outputs without fully realizing it.

5 Step 3: Mitigating Homogenization via Productive Friction: A Multi-Level Case Study Analysis

Having characterized the lifecycle of vibe coding and taxonomized its associated risks, our third and final analytical step is to propose a framework for mitigation. Our analysis suggests that design homogenization is not merely a technical artifact of training data, but a systemic outcome of interaction designs that prioritize speed above
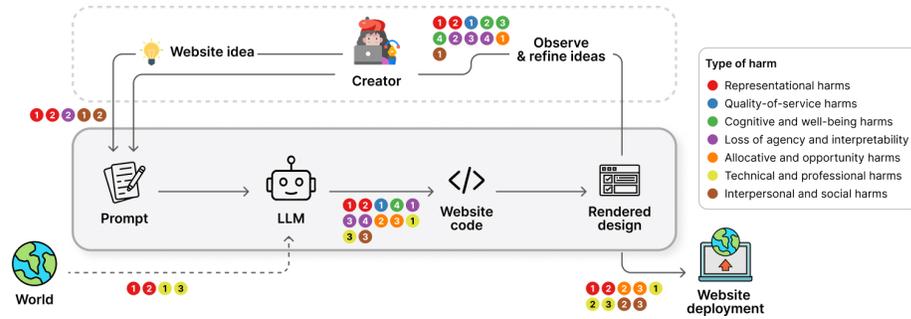
Figure 3: Sub-types of harms mapped onto the lifecycle of vibe coding. Each circled number corresponds to the entries in Table 3.

all else. This risk is most acute at two critical junctures identified in our lifecycle analysis (§3): the initial expression of intent (Stage 1) and the iterative feedback for refinement (Stage 4). As established in Step 2, these moments are when lay creators are most likely to accept a model's default, homogenized output.

We argue that these harms can be mitigated by design,[2] specifically by pursuing *productive friction*: an intentional introduction within a system that creates moments for reflection and clarification [14, 61], moving away from an overemphasis on frictionless efficiency (*e.g.*, [1, 19, 34]). In this section, we first define this concept, grounding it in the evidence gathered in our previous steps, and then apply it through three case studies across three levels—micro, meso, and macro. By targeting the lifecycle's most vulnerable stages, we demonstrate how productive friction can dismantle the mechanisms of homogenization and restore creative agency.

5.1 The Trap of Frictionless Generation

To understand the necessity of friction, we first need to interrogate the design philosophy of the current vibe coding tools. The prevailing paradigm in generative AI prioritizes frictionless efficiency [34], implicitly measuring interaction quality by how quickly underspecified user intent becomes an outcome artifact (*e.g.*, [1, 19, 42]). Within this paradigm, intermediate activities such as deliberation, comparison, and justification are treated as overhead to be minimized rather than as integral to creative reasoning [69]. As observed in our lifecycle analysis (Step 1), this value system collapses what were once multiple distinct phases of design and rendering into a single, instantaneous event (Stages 2 and 3). Success is thus measured by how quickly vague user intent is converted into an outcome artifact [1, 42]. Although a few tools expose a model's preliminary plans or assumptions (Table 1), these traces are typically passive and prescriptive, offering little opportunity for interaction.

Our risk analysis (Step 2) suggests that such a lack of temporal friction can catalyze design homogenization. When the interaction is designed to be seamless, it relies heavily on the model's probabilistic defaults to fill in the gaps of a user's intent. This risks users accepting the first plausible output (§4.2.2), which is prone

to converging on a global statistical mean, and the cost of manually refactoring complex AI-generated code becomes prohibitive (§4.2.3).

5.2 Defining Productive Friction

In response, we propose promoting *productive friction* [14, 61, 76] in the context of vibe coding. Productive friction challenges the assumption that the best interface is one that disappears; instead, it posits that complex creative tasks require moments of deliberate pause, reflection, and negotiation to be performed effectively. In the context of vibe coding, we define productive friction as interactional interventions that disrupt the automatic acceptance of homogenized AI outputs. For example, our Step 2 analysis highlights the *overreliance without verification*, where users feel a sense of agency through simple prompts, even as they surrender creative decision-making. Productive friction serves as a direct counter-measure to such harms by compelling the creator to actively articulate unique, context-aware design intent rather than passively accepting the model's assumptions.

Analytical framework. We frame our analysis of productive friction using a micro–meso–macro framework of social analysis [59], which has been mobilized in influential theories such as Bronfenbrenner [9]'s socio-ecological systems and is increasingly adopted in the research community to unpack the impacts of sociotechnical systems on people and society [68, 70]. In this analytical tradition, the micro level encompasses individual factors like cognition and identity; the meso refers to organizational and relational factors like group membership; and the macro means high-level social forces like society and culture. Mapping our case studies to these three levels allows us to identify specific leverage points where friction can effectively disrupt the cycle of homogenization.

5.3 Micro Level: The Individual Creator

At the micro level, we examine the intimate interaction between an individual lay creator and the vibe coding interface.

5.3.1 The problem: The 'good enough' trap and cultural erasure. When a lay creator prompts a system for a broad concept—for example, 'a landing website for a Japanese retail chain' (Figure 1)—this may result in a globally average aesthetic characterized by minimalist whitespace, standard sans-serif typography, and predictable hero images, regardless of the cultural context and the

---

[2]In our work, we do not claim that design is the only or primary mitigation strategy. Policy, community-based organizing, and advocacy are equally—if not more—important or effective, but often operate on longer timescales and fall outside the scope of this paper.

preference of the regional market. For lay creators often lack technical confidence—a state we identify as *erosion of confidence*—risk arises that this authoritative-looking output remains unchallenged. They accept the homogenized design not because it accurately represents their identity or local visual vernacular, but because the cognitive and technical cost of iterating against the model's default bias or misrepresentation is too high. Consequently, the user falls into the 'good enough' trap, inadvertently erasing their own cultural distinctiveness in favor of a synthetic, globalized standard.

5.3.2 The opportunity: Reflective prompting and visual negotiation. To mitigate this fixation, productive friction can be introduced at the moment of intent expression (Stage 1 of Figure 2). Instead of treating the prompt as a command to be executed immediately, the system should treat it as the opening of a negotiation—a mechanism where the interface pauses to clarify ambiguous design concepts (*e.g.*, cultural signifiers) before generating code.

For instance, if a user requests a website for a local cafe in São Paulo, a tool should not immediately generate a design based on the globally dominant, Western third-wave coffee shop aesthetic (*e.g.*, minimalist layouts, muted tones, sans-serif fonts). Instead, it should pause the process to negotiate the cultural signifier, asking the user to clarify the intended local atmosphere. The interface might support the user's contextualized thinking by visualizing broad mood boards that contrast vibrant, colorful patterns with styles reflecting the city's architectural history or more traditional, idyllic aesthetics. By asking users to choose among diverse cultural representations before code is generated, the tool prevents the accidental adoption of a foreign default. This moment of friction shifts users from passive recipients of model bias to active curators of local identity, ensuring that the final design reflects their contextual judgment.

5.4 Meso Level: Teams and Organizations

At the meso level, we analyze how vibe coding integrates into collaborative workflows. Here, we examine tensions in the handoff between non-technical stakeholders (*e.g.*, product manager) and engineering teams.

5.4.1 The problem: Accelerated brand erosion and team misalignment. In organizational settings, the pressure toward design homogenization predates generative AI. The rise of dominant web design frameworks like Bootstrap [8] and Tailwind CSS [30], while promoting efficiency and consistency, is known to have already led to a convergence of web aesthetics, creating a shared, but often generic, visual language [22]. Vibe coding can further amplify this trend by scaling unmodified defaults; whereas human designers could have introduced subtle tweaks and idiosyncrasies that diversified designs, LLMs tend to favor familiar, widely adopted patterns, reinforcing conformity and accelerating the convergence toward a homogenized aesthetic.

Whereas a team can still consciously adopt such frameworks, a non-technical stakeholder using a vibe coding tool is prone to believing that they are operating with limitless flexibility. Yet, the LLM, trained on a web saturated with these frameworks [29, 57], can probabilistically default to the same familiar patterns—generic layouts, standard component libraries, and a clean, but non-distinctive,

professional veneer. This creates a risk of *team misalignment*; under tight deadlines, the path of least resistance is often not to undertake the costly process of reverse-engineering the AI-generated code to fit the organization's bespoke design system, but to satisfice [64] by accepting the homogenized output. The result is an accelerated erosion of brand identity, where a company's unique visual language is not consciously abandoned, but incrementally and unconsciously replaced by the statistical mean of the web.

5.4.2 The opportunity: Contextual anchoring and active elicitation. At this level, productive friction would function as a mechanism for aligning internal expectation and identity (*e.g.*, brand alignment) with generated codes. Rather than operating in a vacuum, a vibe coding tool intended for organizational use would fundamentally reject the blank slate paradigm. Instead, it can require contextual anchoring through the ingestion of inputs defining identities, such as the organization's existing design system tokens, a brand guidebook, or a URL to their live digital presence.

The friction can be introduced during the iterative feedback stage (Stage 4 of Figure 2). When the tool generates a design, it can perform a comparative analysis between the generated artifact and the organization's existing designs to illustrate a stylistic divergence. If the vibe-coded output uses generic web styles—such as standard drop shadows or rounded corners—that differ from the organization's established design characteristics (*e.g.*, sharp edges, high-contrast borders), the system could present a follow-up dialogue, noting that the draft relies on standard conventions. It can ask the creator if they would refine the artifact by applying their brand's distinct visual elements instead. As such, the tool can ensure that the efficiency of AI generation does not come at the cost of the organization's distinct identity.

5.5 Macro Level: The Ecosystem and Society

Finally, at the macro level, we consider the long-term, ecological impact of vibe-coded websites on the digital landscape and the training of future foundation models.

5.5.1 The problem: The homogenization feedback loop. The threat of design homogenization is not entirely novel. As pointed out above, the rise of dominant design frameworks like Bootstrap and Tailwind CSS has already been affecting the aesthetic diversity of the web by promoting a shared set of visual defaults [22]. However, frictionless vibe coding represents a far more profound and accelerated version of this trend. While such design frameworks offer a constrained set of components, an LLM trained on the internet's statistical mean acts as a universal homogenizer, capable of probabilistically enforcing a global aesthetic average across millions of unique outputs.

This dynamic poses a severe risk: the pollution of the digital commons, a phenomenon researchers have called model collapse [63]. When lay creators potentially mass-produce websites using these homogenized AI defaults, the resulting content can saturate the internet. Since future foundation models will inevitably train on fragments of the public internet (*e.g.*, [29, 57]), they may ingest data that exhibits far less variance and diversity than the pre-AI web. The result is a potential, lasting flattening of digital heritage: niche subcultures, non-standard aesthetics, and diverse cultural design

patterns—already rare in probabilistic distributions—may be effectively erased, overwhelmed by a flood of uniform, AI-generated content. This constitutes an intergenerational allocative harm, whereby future creators lose access to the full spectrum of human visual history because the training data has collapsed into a single, global monoculture.

5.5.2 The opportunity: Provenance and style subfloors. To interrupt this feedback loop, the vibe coding interface could proactively present style subfloors—specialized constraints that anchor generation in distinct aesthetic lineages. By treating the selection of the underlying model weights not as a backend technicality but as a primary interactional choice, the system can ensure that diversity is an input requirement rather than a random output possibility.

In practice, this could manifest in the first stage (Stage 1 of Figure 2) by augmenting the solitary 'Generate' button with an adjacent or hover-activated selection menu displaying a library of community-governed adapters (*e.g.*, LoRA [28]). The menu could be designed to pique the user's curiosity, encouraging them to select a specific style adapter (*e.g.*, Traditional Korean Ink), rather than relying on the generic default. This subtle interactional friction would prompt creators to articulate deliberate stylistic commitments while encoding outputs with semantic provenance metadata, enabling future web crawlers and models to distinguish intentional, stylistically rooted creations from generic AI output, prioritize diversity, and prevent the digital commons from collapsing into a singular visual norm.

6 Discussion

Our investigation into the lifecycle of web vibe coding unpacks an essential tension in this new and increasingly widespread practice: the features that make these tools accessible (*e.g.*, seamless generation) can invite design homogenization. By mapping the vibe coding lifecycle (§3) and conducting a sociotechnical risk analysis (§4), we identified that the source of potential uniformity lies not only in the training data, but also in the frictionless nature of the interaction itself. While these tools can enable lay creators to bypass technical barriers [60, 81], this may come at the expense of creative agency. Our analysis of harms suggests that lay creators may find it difficult to resist a tool that provides an instant, polished, but statistically average output, risking a shift of creators' choices toward homogenized defaults.

The tensions we unpack here are not only concerns for web designers and creatives. They are also a call to action for new research in responsible AI, using emerging challenges in web vibe coding as a research site drawing together expertise across HCI and human-AI collaboration. Web vibe coding is the latest instantiation of a longstanding trend in technology design: a focus on frictionless usability [16] in AI tools, in which the primary aim is to lower the time and cognitive burden associated with a task [1, 19, 34]. But as our exploration of web vibe coding shows, frictionless interactions also give conditions for harms like *cognitive overload*, where the user cannot process the AI's complex output, and *overreliance*, where the user offloads too much decision-making to the AI. As the designer community in human-AI interaction works to balance these tensions, they might benefit from perspectives on appropriate reliance

from the HCI's longstanding interest in accountability and transparency of automated decision-making systems. In turn, research scholars might look to designers for seemingly "uncomfortable" intervention (*e.g.*, [14, 61]) to co-design vibe coding tools themselves, in addition to fairness algorithms, regulatory frameworks, and normative considerations.

We offer our multi-level framework of productive friction (§5) as a starting point for these collaborations. For example, at the micro level, we have argued alongside Anderson *et al.* [3] that designers from the human-AI interaction community can create tools that encourage reflective prompting: vibe coding tools that focus less on wowing users with rapid generation, and more on helping to understand what they actually want to build, via prompting cycles that encourage negotiation between user and AI. Enabling the tool to ask clarifying questions and place the creator in the driver's seat can help beat what we have called the *erosion of confidence* in web vibe coding, where the creator overrelies on the model's outputs and becomes reluctant to guide or contextualize the design produced. Innovations from the HCI community can help ensure fairness in the algorithms that guide when and how the system-creator negotiation proceeds, and whether and how interpretability tools can make it more transparent.

Similarly, as discussed in Step 2, design homogenization can interact with other representational harms (*e.g.*, stereotyping), effectively filling the 'gap' with biased representations. Building on extensive work on algorithmic stereotyping of marginalized groups (*e.g.*, [7, 25, 26]), HCI researchers and designers could team up to examine how design homogenization amplifies such stereotyping, and explore interventions to mitigate these effects. Similar opportunities exist at the meso and macro levels of our framework. At each level, we see avenues for cross-community collaboration that can effectively address problems in web vibe coding and create new case studies for the fundamental arguments within human-AI interaction, towards advancing the state of the art in responsible AI.

7 Limitations & Future Work

While our work provides a taxonomy of harms by drawing on rich sources from both grey literature and academic research, we acknowledge that it may be limited in fully capturing the nuanced experiences of real-world vibe coders. We consider our work a meaningful step in guiding this line of research, and future works could provide empirical evidence in diverse contexts (*e.g.*, game development, scientific computing) to capture their lived experiences.

8 Conclusion

In this paper, we interrogated the emerging practice of web vibe coding, towards anticipating and mitigating the risks of design homogenization. We first characterized the vibe coding lifecycle to pinpoint specific stages where probabilistic defaults threaten to displace intentional design choices. Subsequently, our sociotechnical risk analysis identified the industry's pursuit of frictionless generation as a primary driver of homogenization, revealing how cognitive mechanisms like overreliance incentivize users to accept

homogenized outputs. Finally, we established a multi-level mitigation framework centered on productive friction, advocating for a paradigm shift where AI tools act as active consultants rather than silent executors to foster a more diverse and resilient digital landscape. Our work highlights opportunities for HCI and human-AI interaction in research and innovation towards this paradigm.

References

[1] Vivek Acharya. 2025. Generative AI and the Transformation of Software Development Practices. *arXiv preprint arXiv:2510.10819* (2025).

[2] Dhruv Agarwal, Mor Naaman, and Aditya Vashistha. 2025. AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21. doi:10.1145/3706598.3713564

[3] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization Effects of Large Language Models on Human Creative Ideation. In *Proceedings of the 16th Conference on Creativity & Cognition*. 413–425. doi:10.1145/3635636.3656204

[4] Jacob Anderson. 2025. What is Vibe Coding? How Creators Can Build Software Without Writing Code.

[5] Anthropic. 2025. Turn ideas into interactive AI-powered apps. (2025).

[6] Amanda Baughan, Nigini Oliveira, Tal August, Naomi Yamashita, and Katharina Reinecke. 2021. Do Cross-Cultural Differences in Visual Attention Patterns Affect Search Efficiency on Websites?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12. doi:10.1145/3411764.3445519

[7] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*. 1493–1504.

[8] Bootstrap. 2025. *Bootstrap*. https://getbootstrap.com/

[9] Urie Bronfenbrenner. 1977. Toward an Experimental Ecology of Human Development. *American Psychologist* 32, 7 (1977), 513. doi:10.1037/0003-066X.32.7.513

[10] 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. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf

[11] Zana Buçinca, Chau Minh Pham, Maurice Jakesch, Marco Tulio Ribeiro, Alexandra Olteanu, and Saleema Amershi. 2023. AHA!: Facilitating AI Impact Assessment by Generating Examples of Harms. *arXiv preprint arXiv:2306.03280* (2023). https://arxiv.org/abs/2306.03280

[12] Francisco Castro, Jian Gao, and Sébastien Martin. 2023. Human-AI Interactions and Societal Pitfalls. *arXiv preprint arXiv:2309.10448* (2023). https://arxiv.org/abs/2309.10448

[13] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating Large Language Models Trained on Code. *arXiv preprint arXiv:2107.03374* (2021). https://arxiv.org/abs/2107.03374

[14] Anna L. Cox, Sandy J.J. Gould, Marta E. Cecchinato, Ioanna Iacovides, and Ian Renfree. 2016. Design Frictions for Mindful Interactions: The Case for Microboundaries. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 1389–1397. doi:10.1145/2851581.2892410

[15] Philipp Eibl, Sadra Sabouri, and Souti Chattopadhyay. 2025. Exploring the Challenges and Opportunities of AI-assisted Codebase Generation. *arXiv preprint arXiv:2508.07966* (2025). https://arxiv.org/abs/2508.07966

[16] Jonathan Ericson. 2022. Reimagining the Role of Friction in Experience Design. *Journal of User Experience* 17, 4 (2022).

[17] Justin Etheredge. 2025. Writing code with an LLM (like ChatGPT or Claude)...

[18] Zhiyu Fan, Xiang Gao, Martin Mirchev, Abhik Roychoudhury, and Shin Hwei Tan. 2023. Automated Repair of Programs from Large Language Models. In *2023 IEEE/ACM 45th International Conference on Software Engineering*. IEEE, 1469–1481. doi:10.1109/ICSE48619.2023.00128

[19] Ahmed Fawzy, Amjed Tahir, and Kelly Blincoe. 2025. Vibe Coding in Practice: Motivations, Challenges, and a Future Outlook – a Grey Literature Review. *arXiv preprint arXiv:2510.00328* (2025).

[20] Emilio Ferrara. 2024. GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models. *Journal of Computational Social Science* 7, 1 (2024), 549–569. doi:10.1007/s42001-024-00250-1

[21] Smashing Frames. 2025. On "Vibe Coding."

[22] Aaron Ganci and Bruno Ribeiro. 2017. On Web Brutalism and Contemporary Web Design. *Dialectic* 1, 1 (2017). doi:10.3998/dialectic.14932326.0001.107

[23] Sanjana Gautam, Pranav Narayanan Venkit, and Sourojit Ghosh. 2024. From Melting Pots to Misrepresentations: Exploring Harms in Generative AI. *arXiv preprint arXiv:2403.10776* (2024). doi:abs/2403.10776

[24] Francis Geng, Anshul Shah, Haolin Li, Nawab Mulla, Steven Swanson, Gerald Soosai Raj, Daniel Zingaro, and Leo Porter. 2025. Exploring Student-AI Interactions in Vibe Coding. *arXiv preprint arXiv:2507.22614* (2025). https://arxiv.org/abs/2507.22614

[25] Sourojit Ghosh and Aylin Caliskan. 2023. ChatGPT Perpetuates Gender Bias in Machine Translation and Ignores Non-Gendered Pronouns: Findings across Bengali and Five other Low-Resource Languages. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 901–912. doi:10.1145/3600211.3604672

[26] Sourojit Ghosh, Nina Lutz, and Aylin Caliskan. 2024. "I Don't See Myself Represented Here at All": User Experiences of Stable Diffusion Outputs Containing Representational Harms across Gender Identities and Nationalities. In *Proceedings of the Seventh AAAI/ACM Conference on AI, Ethics, and Society*, Vol. 7. 463–475. doi:10.1609/aies.v7i1.31650

[27] Google. 2025. Gemini Canvas. (2025).

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

[29] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. GPT-4o System Card. *arXiv preprint arXiv:2410.21276* (2024). https://arxiv.org/abs/2410.21276

[30] Tailwind Labs Inc. 2025. *Tailwind CSS*. https://tailwindcss.com/

[31] Rebecca L Johnson, Giada Pistilli, Natalia Menédez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The Ghost in the Machine has an American accent: value conflict in GPT-3. *arXiv preprint arXiv:2203.07785* (2022). https://arxiv.org/abs/2203.07785

[32] Waleed Kadous. 2025. Vibe-coding is a horrible idea. So is dismissing AI-assisted coding.

[33] karpathy. 2025. Vibe coding MenuGen.

[34] Andrej Karpathy. 2025. There's a new kind of coding I call "vibe coding"...

[35] Harsh Kumar, Jonathan Vincentius, Ewan Jordan, and Ashton Anderson. 2025. Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18. doi:10.1145/3706598.3714198

[36] Bingshuai Liu, Longyue Wang, Chenyang Lyu, Yong Zhang, Jinsong Su, Shuming Shi, and Zhaopeng Tu. 2023. On the cultural gap in text-to-image generation. *arXiv preprint arXiv:2307.02971* (2023).

[37] Sijia Liu, Xiaoke Zeng, Fengyihan Wu, Shu Ye, Bowen Liu, Sidney Cheung, Richard William Allen, and Ray Lc. 2025. "Salt is the Soul of Hakka Baked Chicken": Reimagining Traditional Chinese Culinary ICH for Modern Contexts Without Losing Tradition. In *Proceedings of the 2025 Conference on Creativity and Cognition*. 198–210. doi:10.1145/3698061.3726917

[38] Lovable. 2025. Lovable. (2025).

[39] Stephane H Maes. 2025. Ensuring the Maintainability and Supportability of "Vibe-Coded" Software Systems: A Framework for Bridging Intuition and Engineering Rigor. (2025).

[40] Antonio Mastropaolo, Luca Pascarella, Emanuela Guglielmi, Matteo Ciniselli, Simone Scalabrino, Rocco Oliveto, and Gabriele Bavota. 2023. On the Robustness of Code Generation Techniques: An Empirical Study on GitHub Copilot. In *2023 IEEE/ACM 45th International Conference on Software Engineering*. IEEE, 2149–2160. doi:10.1109/ICSE48619.2023.00181

[41] Andrew M McNutt, Chenglong Wang, Robert A Deline, and Steven M Drucker. 2023. On the Design of AI-powered Code Assistants for Notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16. doi:10.1145/3544548.3580940

[42] Kushan Mitra, Dan Zhang, Hannah Kim, and Estevam Hruschka. 2025. RECAP: REwriting Conversations for Intent Understanding in Agentic Planning. *arXiv preprint arXiv:2509.04472* (2025). https://arxiv.org/abs/2509.04472

[43] Nhan Nguyen and Sarah Nadi. 2022. An Empirical Evaluation of GitHub Copilot's Code Suggestions. In *Proceedings of the 19th International Conference on Mining Software Repositories*. 1–5. doi:10.1145/3524842.3528470

[44] Manuel Nordhoff, Tal August, Nigini A Oliveira, and Katharina Reinecke. 2018. A Case for Design Localization: Diversity of Website Aesthetics in 44 Countries. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12. doi:10.1145/3173574.3173911

[45] Matthew Nyaaba, Alyson Wright, and Gyu Lim Choi. 2024. Generative AI and Digital Neocolonialism in Global Education: Towards an Equitable Framework. *arXiv preprint arXiv:2406.02966* (2024). https://arxiv.org/abs/2406.02966

[46] OpenAI. 2024. Introducing canvas. (2024).

[47] Jonas Oppenlaender, Johanna Silvennoinen, Ville Paananen, and Aku Visuri. 2023. Perceptions and Realities of Text-to-Image Generation. In *Proceedings of the 26th International Academic Mindtrek Conference*. 279–288. doi:10.1145/3616961.3616978

[48] Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri. 2025. Asleep at the Keyboard? Assessing the Security of GitHub Copilot's Code Contributions. *Commun. ACM* 68, 2 (2025), 96–105. doi:10.1145/3610721

[49] Rida Qadri, Renee Shelby, Cynthia L Bennett, and Remi Denton. 2023. AI's Regimes of Representation: A Community-centered Study of Text-to-Image Models in South Asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 506–517. doi:10.1145/3593013.3594016

[50] Qingyang Ren, Zilin Jiang, Jinghan Cao, Sijia Li, Chiqu Li, Yiyang Liu, Shuning Huo, Tiange He, and Yuan Chen. 2024. A survey on fairness of large language models in e-commerce: progress, application, and challenge. *arXiv preprint arXiv:2405.13025* (2024). https://arxiv.org/abs/2405.13025

[51] Replit. 2025. Replit – Build apps and sites with AI. (2025).

[52] Replit. 2025. What is Vibe Coding? How To Vibe Your App to Life.

[53] Hannah R Rothstein and Sally Hopewell. 2009. Grey literature. *The handbook of research synthesis and meta-analysis* 2 (2009), 103–125.

[54] Sebastin Santy, Jenny T Liang, Ronan Le Bras, Katharina Reinecke, and Maarten Sap. 2023. NLPositionality: Characterizing Design Biases of Datasets and Models. *arXiv preprint arXiv:2306.01943* (2023). https://arxiv.org/abs/2306.01943

[55] Ranjan Sapkota, Konstantinos I Roumeliotis, and Manoj Karkee. 2025. Vibe Coding vs. Agentic Coding: Fundamentals and Practical Implications of Agentic AI. *arXiv preprint arXiv:2505.19443* (2025). https://arxiv.org/abs/2505.19443

[56] Advait Sarkar and Ian Drosos. 2025. Vibe coding: programming through conversation with artificial intelligence. *arXiv preprint arXiv:2506.23253* (2025). doi:abs/2506.23253

[57] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. LAION-5B: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* 35 (2022), 25278–25294. doi:10.5555/3600270.3602103

[58] Agnia Sergeyuk, Yaroslav Golubev, Timofey Bryksin, and Iftekhar Ahmed. 2025. Using AI-based coding assistants in practice: State of affairs, perceptions, and ways forward. *Information and Software Technology* 178 (2025), 107610. doi:10.1016/j.infsof.2024.107610

[59] Sandro Serpa and Carlos Miguel Ferreira. 2019. Micro, meso and macro levels of social analysis. *Int'l J. Soc. Sci. Stud.* 7 (2019), 120.

[60] Remy Sharp. 2025. Vibe coding and Robocop.

[61] Jacob Sheahan, David Chatting, Robert Collins, Jessica Bley, Alexander Eriksson, Nick Taylor, and Marco C Rozendaal. 2024. Designing with Friction: Inverting Notions of Seamless Technology. In *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction*. 1–4. doi:10.1145/3677045.3685504

[62] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, et al. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 723–741. doi:10.1145/3600211.3604673

[63] Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. AI models collapse when trained on recursively generated data. *Nature* 631, 8022 (2024), 755–759.

[64] Herbert A Simon. 1956. Rational choice and the structure of the environment. *Psychological review* 63, 2 (1956), 129.

[65] Ernestine Siu. 2025. I took a 2-day "vibe coding" class and successfully built a product. Here are my biggest takeaways.

[66] Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Canyu Chen, Hal Daumé III, Jesse Dodge, Isabella Duan, et al. 2023. Evaluating the Social Impact of Generative AI Systems in Systems and Society. *arXiv preprint arXiv:2306.05949* (2023). https://arxiv.org/abs/2306.05949

[67] Harini Suresh and John Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–9. 10.1145/3465416.3483305

[68] Harini Suresh, Emily Tseng, Meg Young, Mary Gray, Emma Pierson, and Karen Levy. 2024. Participation in the Age of Foundation Models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1609–1621. doi:10.1145/3630106.3658992

[69] Yuto Suzuki and Farnoush Banaei-Kashani. 2025. Universe of Thoughts: Enabling Creative Reasoning with Large Language Models. *arXiv preprint arXiv:2511.20471* (2025). https://arxiv.org/abs/2511.20471

[70] Emily Tseng, Meg Young, Marianne Aubin Le Quéré, Aimee Rinehart, and Harini Suresh. 2025. " Ownership, Not Just Happy Talk": Co-Designing a Participatory Large Language Model for Journalism. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 3119–3130.

[71] Stefanie Urchs, Veronika Thurner, Matthias Aßenmacher, Ludwig Bothmann, Christian Heumann, and Stephanie Thiemichen. 2025. Are All Genders Equal in the Eyes of Algorithms?–Analysing Search and Retrieval Algorithms for Algorithmic Gender Fairness. *arXiv preprint arXiv:2508.05680* (2025). https://arxiv.org/abs/2508.05680

[72] Aleksandra Urman and Mykola Makhortykh. 2024. "Foreign beauties want to meet you": The sexualization of women in Google's organic and sponsored text search results. *New Media & Society* 26, 5 (2024), 2932–2953. doi:10.1177/14614448221099536

[73] Faye-Marie Vassel, Evan Shieh, Cassidy R Sugimoto, and Thema Monroe-White. 2024. The Psychosocial Impacts of Generative AI Harms. In *Proceedings of the AAAI Symposium Series*, Vol. 3. 440–447. doi:10.1609/aaaiss.v3i1.31251

[74] Vercel. 2025. *V0 by Vercel*. https://v0.app/

[75] Vercel. 2025. Vercel. (2025).

[76] Ron Wakkary, Henry Lin, Shannon Mortimer, Lauren Low, Audrey Desjardins, Keith Doyle, and Philip Robbins. 2016. Productive Frictions: Moving from Digital to Material Prototyping and Low-Volume Production for Design Research. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 1258–1269. doi:10.1145/2901790.2901880

[77] HuPei Wang. 2025. Adaptive fusion of multi-cultural visual elements using deep learning in cross-cultural visual communication design. *Scientific Reports* 15, 1 (2025), 28431. doi:10.1038/s41598-025-13386-5

[78] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of Risks Posed by Language Models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 214–229. doi:10.1145/3531146.3533088

[79] Simon Willison. 2025. Not all AI-assisted programming is vibe coding (but vibe coding rocks).

[80] Andre Ye, Sebastin Santy, Jena D Hwang, Amy X Zhang, and Ranjay Krishna. 2025. Semantic and Expressive Variations in Image Captions Across Languages. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 29667–29679. doi:10.1109/CVPR52734.2025.02762

[81] Rina Zviel-Girshin. 2024. The Good and Bad of AI Tools in Novice Programming Education. *Education Sciences* 14, 10 (2024), 1089. doi:10.3390/educsci14101089

A   Study Detail

A.1   Queries Used for Systematic Review

Table 2: Full list of OR-separated keyword queries by thematic category

Category	Queries
Informal and vibe-centered phrasing	vibe coding
Interaction and technique-centered phrasing	natural language website generation, prompt-based website coding, text-based website coding, text to website code
Platform-specific use cases	ChatGPT / Gemini / Claude make a website, build a website with ChatGPT / Gemini / Claude, use ChatGPT / Gemini / Claude to code a webpage, how to build an app with ChatGPT / Gemini / Claude, generate HTML with LLM, ChatGPT / Gemini / Claude build website from prompt

A.2 Harm Taxonomy and the Intersection with Design Homogenization

Table 3: Types of harms in LLM-assisted website creation by lay creators, categorized into seven types based on prior frameworks. For each sub-type, we provide a definition and its intersection with the harm of our focus (*i.e.*, design homogenization).

Type	Type definition	Sub-type	Sub-type definition	Intersection with design homogenization
Representational harms	Harmful or exclusionary representations of people, groups, or values in code, content, or AI suggestions	❶ Design homogenization	Promoting aesthetics, layout styles, or language that exclude underrepresented preferences	*(Our focus)*
		❷ Stereotyping through design output	Generating site content, names, or images that reflect social biases (*e.g.*, race, gender)	Defaults to stereotypical heuristics as a substitute for creative expression
Quality-of-service harms	Breakdowns or inconsistencies in outputs that impair novices' ability to complete tasks	❶ Incoherent/ nonfunctional code	Returning outputs that appear to work but fail when executed, without explanation	Associates prompts for creative expressions with high failure rates
Cognitive and well-being harms	Emotional, psychological, or motivational burdens experienced during LLM-assisted coding	❶ Erosion of confidence	Users feeling incapable when outputs fail and left without tools to troubleshoot	Depletes the user's motivation to counter homogenized design
		❷ Cognitive overload	Returning complex suggestions that lack scaffolding, forcing users to interpret unfamiliar patterns alone	Pushes users towards simpler solutions to avoid mental fatigue
		❸ Discouragement from contradictions	Contradictory guidance across prompts causes confusion and fatigue	Causes users to abandon more unique creative directions in favor of simpler, achievable ones
		❹ Missed learning moments	Delivering code 'magic-style' without sufficiently detailing the why, depriving users of conceptual understanding	Prevents users from acquiring the skills needed to build or customize designs from first principles
Loss of agency and interpretability	Creators struggle to understand or control what the system is doing or why	❶ Opaque design justification	Users unable to tell why certain code was generated or what it does, and generated code lacks links to documentation, standards, or explanations	Reduces the user's ability to articulate or implement custom designs, making them a passive recipient of the AI's limited aesthetic range
		❷ Prompt engineering dependency	Effective use depends on tacit knowledge of prompt phrasing, which novices lack	Makes it easier to prompt simple requests
		❸ Overreliance without verification	Users blindly accepting codes, feeling in control through simple commands while actually relinquishing it through blind acceptance	Makes the homogenized design seem like the user's own, reducing the desire to seek alternatives
		❹ Non-editable complexity	Code being too complex or interdependent for users to safely edit	Locks in the homogenized design, making iterative refinement towards countering homogenization impractical
Allocative and opportunity harms	Missed, lost, or denied benefits due to the design or performance of the LLM	❶ Wasted time and effort	Users iterating through nonfunctional outputs consumes significant time	The opportunity cost of experimenting with designs becomes high
		❷ Misleading success / Reinforcement of confirmation bias	Users becoming overconfident in working prototypes, which masks deeper issues, potentially harming a product deployment	Reinforces the idea that the quickest, most 'successful' path is to accept the initial, often generic, output
		❸ Design fixation	Steering the user's design process into a narrow conceptual space, making it difficult for them to generate novel or diverse solutions	Anchors the user to the AI's first suggestion, short-circuiting the divergent thinking process to counter design homogenization
Technical and professional harms	Generated code or design decisions harm product quality, compliance, or end-user trust	❶ Bad code practices	Generating deprecated, non-performant, or insecure code without warning	Creates a false association where technically 'safe' or 'professional' code is equated with common, homogenized design patterns
		❷ Economic harm via launch	Users deploying broken or unethical sites, incurring reputational or financial risk	Fear of negative consequences from custom code incentivizes reliance on widely used, generic templates presumed to be safer
		❸ Copyright or licensing violations	Returning assets or text that infringes on rights without creators' awareness	-
Interpersonal and social harms	The LLM shapes how creators interact with others and perceive their social standing	❶ Learning in isolation	Replacing mentors which leads to a solitary and fragile learning process	Cuts the creator off from team or community feedback on creative expressions
		❷ Shame in collaboration	Users concerned about showing their vibe-coded outputs to others due to technical gaps	Prevents peer review and constructive criticism that could challenge generic designs and introduce perspectives
		❸ Team misalignment	Maintenance burdens when handed off to other team members	-