Tigo Ponce de León

MAAD 14109 - Prof. Mankin

December 13th, 2025

Statistical Aesthetics: The Homogenization of Visual Language in AI-Generated User Interfaces

---

Essential Paper Details

AI (LLM) Prompt:
Provide me the code (in the coding language that you deem best) for me to eventually push a website to the internet. This website should be a recipe generator. Users should be able to upload ingredients that they have in stock in their kitchen and the website should be able to generate recipes given these inputs. Make the UI and UX as good as you possibly can.

AI (Website Builder) Prompt:
Make me a recipe generator website. Users should be able to upload ingredients that they have in stock in their kitchen and the website should be able to generate recipes given these inputs. Make the UI and UX as good as you possibly can.

Gemini Site:
https://tigo-poncedeleon.github.io/gemini-recipe-generator/

ChatGPT Site:
https://tigo-poncedeleon.github.io/chatgpt-recipe-generator/

Lovable Site:
https://kitchen-magic-maker-51.lovable.app

Base44 Site:
https://culinary-ai-9a870b49.base44.app

Grok Site:
https://grok-recipe-generator.onrender.com/

Claude:
https://claude.ai/public/artifacts/5137ed3e-4168-48dd-877e-26f4158399c2

Table Comparison

	ChatGPT	Gemini	Claude	Lovable	Base44	Grok
Layout Structure	Two cards: one for input and one for output	Big text as focal point, traditional nav bar, and find recipe search input	Minimal, vertically stacked	Generic food image with big, ad-like text, input search lower on page	Two cards: one for input and one for output, recipe generator logo at top	Hyper minimalist: "What's in your fridge?" + text box
Grid Logic	18 pixel padding between items	Responsive column grid	Loose but consistent margins	No components next to each other, more like a slide show scrolling down	Component cards leaning to left hand of page to give space for recipes on right side	Text + search bar + enter button
Color Palette	Black and blue gradient background with orange and yellow gradient highlights	Creamy white background with orange gradient highlights	Minimal neutral white and orange tones	Image with produce is biggest aesthetic feature	White background with gradient orange highlights	So much white space with black text and gray enter button
Material Metaphor	Flat, card-based	Soft shadows, depth	Near-flat	Elevated produce	Subtle elevation on component cards	Almost none
Typography	Sans-serif, system-like	Rounded modern sans	Clean, editorial sans	Playful rounded sans	Rounded, modern sans	Default web sans
Component Style	Buttons, cards, inputs	Cards, chips, CTA buttons	Buttons, cards, inputs, and preference options	Bare generate recipe component	Simple components	Basic HTML forms
Interaction Feedback	Hover effects on orange elements	Loading state animation when finding recipes	Subtle loading animation	Subtle loading animation	Ingredients change color to orange and hover and recipe image zooms in on hover	Functional only

Logo Design	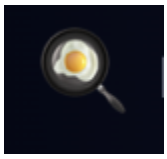					None
Overall Aesthetic Register	"Modern SaaS"	"Friendly productivity"	"Quiet utility"	"Image-driven"	"Professional dashboard"	"Functional prototype"

Essay

Generative artificial intelligence systems are now capable of producing entire user interface designs for "vibe-coded" websites, a development that raises fundamental questions about creativity, originality, and convergence in the post-AI web landscape. While popular discourse often frames AI-generated design as novel or disruptive, closer scrutiny suggests a different dynamic at work: one of repetition, standardization, and statistical norm reproduction. This study examines six functionally identical recipe-generator web interfaces, each produced by a different generative AI platform—Google Gemini, OpenAI's ChatGPT, Anthropic's Claude, Lovable AI, Base44, and xAI's Grok. All six systems were given the same design prompt (available to read in its entirety above): to build a recipe generator web application that accepts ingredient input and returns a generated recipe. This controlled setup allows the resulting interfaces to be treated as a comparative corpus of designed artifacts. Notably, the generators were intentionally given very little frameworks; they were given free reign to design the UI and UX to what they conceive as "best." Also, they weren't constrained to any one coding language or framework.

At its core, the central research question guiding this inquiry is not whether these AI systems can produce usable interfaces—they pretty clearly can—but rather how their designs converge, and

what that convergence reveals about the nature of AI-mediated design. Specifically, this paper asks: how do different generative AI systems, when given the same design task, arrive at strikingly similar visual and interaction patterns, and what does this convergence tell us about the statistical and archival logic underpinning generative design models?

In this paper, I argue that these systems do not invent new interface aesthetics. Instead, they operationalize a statistical consensus of contemporary product design conventions. The resulting interfaces are visually competent, legible, and broadly usable, yet culturally homogenized and strikingly boring. Their likeness, and oftentimes exact sameness, reflects not creative authorship, but the power of dominant design archives (frameworks, templates, design systems, and platform guidelines) that define what "good design" looks like at scale. Because the training datasets and internal parameters of these models are proprietary and opaque, this study adopts an output-based inference approach: for every AI, I inputted the same prompt then deployed the code they gave me to the web through Github Pages. Here, I treat the AI-generated interfaces themselves as empirical evidence of the aesthetic archives embedded within the models. Through close analysis of the outputs, we can infer the otherwise hidden design assumptions the systems have internalized.

Situating this inquiry within design theory and media studies, the paper combines three methodological lenses: visual genealogy, archive theory, and empirical interface analysis. Visual genealogy traces the historical lineages of web and product design aesthetics that inform contemporary interfaces. Archive theory, drawing from thinkers such as Derrida and Foucault, frames AI training data as an aesthetic archive that preserves certain norms while excluding others. Empirical interface analysis, grounded in close personal observation, provides a systematic way to document and compare the concrete visual and interactive features of each interface. Together, these approaches allow AI-generated UIs to be understood not as neutral outputs, but as culturally situated artifacts that reflect dominant design ideologies.

Methodologically, this study treats each of the six recipe-generator websites as a primary source, akin to case studies in architecture or graphic design history. All six sites were generated through first-hand experimentation by me using each platform's native tools or APIs. No manual human coding, styling, or post-generation modification was applied. Each site implements the same core functionality: an input mechanism for ingredients and a generated recipe output. This creates a tightly controlled corpus in which differences can plausibly be attributed to the AI systems themselves rather than divergent requirements.

The analysis proceeds through close reading of visual and interaction design elements, including layout, whitespace, color, typography, UI components, and affordances, which can be seen described in specificity in the table above. Close observation here means careful description of observable attributes, trying to erase subjectivity wherever possible. Ultimately, the goal is not to judge aesthetic quality based on my personal taste, but to identify recurring patterns and shared assumptions. Comparative analysis then highlights points of convergence and minor deviation across the six interfaces.

Across the corpus, the most striking feature is the degree of visual and structural similarity. All six interfaces employ a clean, centered layout with a linear top-to-bottom flow. Typically, a header or title appears at the top, followed by an input area for ingredients, a primary call-to-action button, and a results section below. The layout logic is unmistakably that of a single-column responsive grid, occasionally widening into a modest two-column arrangement on large screens. None of the designs attempt asymmetry, experimental navigation, or unconventional spatial organization. Instead, they converge on what has become the default architecture of modern web applications.

This layout convergence reflects the influence of responsive design frameworks and corporate guidelines that emphasize clarity, scalability, and device-agnostic structure. Generous margins and centered content suggest adherence to best practices commonly practiced in industry.

The repeated use of card-like containers for displaying recipe results further reinforces this lineage. Cards, introduced and popularized through Material Design, have become a dominant pattern for grouping content. Their presence across all six interfaces indicates how deeply this component has been naturalized within the design archive.

Whitespace plays a similarly uniform role across the designs. All six interfaces use abundant negative space to separate elements, establish hierarchy, and create a sense of visual calm. Input fields, buttons, and output sections are isolated through padding rather than decorative borders. This treatment reflects a widely accepted contemporary belief that good design is uncluttered and restrained. Whitespace functions as a structural tool rather than a stylistic flourish, guiding the user's attention and reinforcing information hierarchy. The consistency of whitespace usage across independent AI systems underscores how strongly minimalist aesthetics have been encoded as normative.

Color usage further illustrates this statistical convergence. Each interface adopts a light background (typically white or off-white) with dark text for readability. A single accent color is used sparingly to highlight primary actions such as the "Generate Recipe" button or section headers. Fascinatingly, almost all of the websites decide to use orange gradients without me instructing them explicitly to do so. Likewise, half of the designs use the exact same chef top hat logo, which is too great of a number to be done coincidentally. Likely, these different LLMs are all operating from an incredibly similar, if not the exact same, dataset.

Moving on, the spatial aesthetic employed by these designs also converge around flatness with subtle depth cues. Backgrounds are flat, textures are absent, and depth is suggested only through faint shadows on cards or buttons. This reflects the post-skeuomorphic design paradigm established in the early 2010s, when textured realism gave way to flat, abstracted interfaces. Both Apple's and Google's design systems played a central role in this transition, and the AI-generated

interfaces demonstrate how thoroughly this aesthetic has been internalized. None of the systems attempt alternative visual languages such as brutalism, retro skeuomorphism, or ornamental design.

Typography across the six interfaces is similarly uniform. All designs rely on modern sans-serif fonts, likely system defaults or widely used web fonts. Serif fonts, decorative typefaces, or expressive typographic gestures are entirely absent. Hierarchy is established through size and weight: large titles at the top, medium-weight section headers, and smaller body text for instructions. Line spacing and alignment follow standard readability conventions. These typographic choices mirror the guidance found in platform style guides, which emphasize legibility, consistency, and restraint. Typography, like color and layout, is treated as an invisible infrastructure rather than a site of expression.

Also, UI components reinforce this pattern of convergence. Input fields are rectangular, lightly bordered, and padded, often with placeholder text indicating expected input. Buttons are solid, rounded rectangles in the accent color, sized generously to meet touch-target guidelines. Output content is displayed in cards with clearly delineated sections for ingredients and instructions. Navigation elements are minimal or nonexistent, reflecting the single-purpose nature of the app. No interface introduces novel components or experimental interaction patterns. Instead, all rely on a shared library of familiar UI primitives.

In the same vein, interaction affordances and feedback mechanisms further demonstrate adherence to mainstream UX heuristics. Buttons look clickable, input fields invite typing, and loading or result states provide feedback that the system is responding. While the specific animations or transitions may vary slightly depending on platform defaults, the underlying interaction logic is consistent. The designs reflect an implicit understanding of principles such as visibility of system status, error prevention, and immediate feedback. These are not innovative choices but established norms derived from decades of usability research and platform documentation.

When viewed holistically, the six interfaces exhibit overwhelming convergence. Differences exist, but they are superficial: a slightly different accent color, some buttons more rounded, or the inclusion of an image upload feature in one case. These variations occur within a tightly bounded aesthetic space. No system breaks from the shared template to propose a radically different interface logic or visual language. The result is six interfaces that are pretty much easily interchangeable, both visually and experientially.

This obvious convergence is significant because the AI systems in theory operate independently, trained on slightly different datasets and built by competing organizations. But with that said, their similarity suggests that the dominant design norms of contemporary digital culture exert a powerful gravitational pull. Generative AI, when asked to design an interface without highly specific constraints, gravitates toward the statistical center of its training data. The most probable solution, or the one most consistent with thousands of examples labeled implicitly as "good design", wins out and looks the same.

From an archival perspective, this experiment reveals how AI models function as curators of aesthetic memory. The training data acts as an archive that preserves certain design ideologies while marginalizing others. Corporate design systems, startup templates, and popular frameworks dominate this archive, shaping what the AI learns as normative. Through output-based inference, we can see how these archives manifest in the interfaces themselves. The designs tell us less about the creative agency of AI and more about the consolidation of design culture over the past decade and the increasing homogenization of visual language on the web, which is becoming ever more widespread with many developers relying heavily on AI-produced code.

This has broad implications for the future of design: on one hand, such convergence can be beneficial in that familiar interfaces reduce cognitive load, improve accessibility, and allow users to transfer knowledge across apps; on the other hand, the widespread adoption of generative design

tools risks accelerating aesthetic homogenization and lack of imagination. If AI-generated interfaces become the default starting point for design, the visual diversity of digital culture will surely narrow further, reinforcing a monoculture defined by a small set of hegemonic standards.

Ultimately, this study treats AI-generated user interfaces not as technological novelties, but as cultural artifacts. They are mirrors reflecting the dominant values, assumptions, and aesthetic preferences embedded in contemporary design archives. The creativity of generative AI lies not in invention, but in recombination, sampling from what already exists and assembling it into statistically coherent forms. By closely reading these outputs, we gain insight into the collective design memory that AI systems have learned, and into the power structures that define what counts as "good design" in the age of machine generation. Contrary to what many say, design is not dead; and, in fact, creative human designers are needed more than ever to breathe life back into the web.