

The Human Touch(point): Recommendations for Thoughtful AI Feature Design

Sierra Shell 🝺

Product Design, SAS Institute Inc., Cary, NC, USA (sierrashell1991[at]gmail.com)

Abstract: This article examines the evolving responsibilities of designers in an era of explosive AI growth. AI is a horizontal technology affecting nearly all industries, and designers must position themselves as stewards of the "human in the loop" to balance technological capabilities with human needs. Three core principles can guide ethical AI design: First, AI features should solve specific user problems rather than being implemented for novelty or marketing purposes. Second, strategic friction can serve as a beneficial design element when deployed at consequential decision points, encouraging users to engage thoughtfully with AI-generated content. Third, robust user feedback mechanisms should be prioritized to ensure continuous improvement based on real-world usage. Generative AI should facilitate — never replace — human expertise to avoid the centralization of ideas and displacement of creativity. Drawing from industry examples, the article demonstrates that successful human-AI collaboration depends not on technological sophistication alone, but on thoughtful design that empowers users as active participants rather than passive consumers of AI outputs. Eight recommendations are provided to ensure that the three core principles discussed are incorporated into a product's design.

Implications for research: This article raises issues that suggest fruitful research areas. The recommendation to compare a potential AI-driven design solution to a non-AI solution (Section 2) suggests the development of a conceptual framework for systematically assessing and justifying aspects of AI specification. Such a framework should be compact enough to integrate into R&D processes for ground-level impacts. The principle of thoughtful friction in AI product design (Section 3) is introduced in relation to frictive AI elements of granular typographic detail such as point size and positioning. Empirical research could determine which typographic and interface characteristics provide friction for contemporary users, to what degree different options do so, and what thresholds exist for users' emotional and cognitive responses. Further research could examine whether such frictive elements effectively facilitate trust calibration for potentially

@: ARTICLE> DISPATCHES>

Cite this article:

Shell, S. (2025). The human touch(point): Recommendations for thoughtful AI feature design. *Visible Language* [early view], 1–12. https://www.visible-language.org/journal/ issue-59-2-dfi-the-human-touchpoint-recommendations-forthoughtful-ai-feature-design/ First published online May 22, 2025.

© 2025 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

https://visible-language.org/journal/

Visible Language Consortium:

University of Leeds (UK) University of Cincinnati (USA) North Carolina State University (USA) erroneous AI-generated content. Finally, the recommendation to integrate feedback mechanisms into AI-driven products (Section 4) suggests research aimed at understanding the current state of AI feature feedback collection in popular products, both in terms of mechanism and frequency.

Keywords: AI design; design industry; ethical AI; friction; product design; responsible AI

1. Introduction

In their viral lecture *The AI Dilemma*, Center for Humane Technology co-founders Tristan Harris and Aza Raskin (2023) describe how the introduction of a new technology necessitates a new class of responsibilities. They are not the first to suggest as much, and digital designers are familiar with the emergence of novel and disruptive tech. As a horizontal technology, AI is already influencing almost every industry and digital service. Companies are quick to evangelize their "GenAI" (generative AI) strategies, sometimes incorrectly labeling broadly defined AI features as generative AI to leverage its current popularity. Disentangling buzzwords and hype from the potential substantive improvement AI offers is the job of technologists collectively. However, for the designer, AI escalates both the power and consequences of design choices.

The role of designers in the creation of AI features, products, and services is to be stewards of the "human in the loop." To that end, product designers, visual designers, user researchers, and others in the field of human-computer interaction must thoroughly understand a user's problem and how AI features specifically and uniquely address it. They should use friction intentionally to highlight important decisions and other moments that require careful human consideration. Finally, to minimize organizational risk, designers should include a user feedback mechanism alongside GenAI features to ensure they understand users' experiences. Most importantly, user feedback should be prioritized over marketable but superficial new features.

2. AI Features Should Solve a User's Problem

Problem definition is nothing new to the world of design. Experts like those in the Nielsen/Norman Group have shared guidance on using "how might we" statements to ideate on the right problem (Rosala, 2021) and crafting user need statements when defining the problem using design thinking methodology (Gibbons, 2019). A designer should seek, above all, to help a user solve their problem in the simplest, most intuitive way possible. The first step in doing so successfully is accurate problem definition.

Companies often seek to increase the marketability of their products and services by implementing the novel technology *du jour*. Of 8,000 business leaders surveyed for Cisco's AI Readiness Index, 98% feel increased urgency to use AI within their organization ("Cisco 2024 AI Readiness Index," 2024). However, feeling pressure to utilize AI does not mean its implementation will inevitably be helpful. If used unnecessarily, it can introduce needless risk, requiring compliance with complex laws and standards. By understanding a user's problem, we can ensure an AI-driven solution is thoughtfully crafted and demonstrably better than one that does not employ AI.

A simple exercise can allow designers to feel confident that their AI use is thoughtfully addressing a user's problem: brainstorm a solution or task flow that does not incorporate AI at all. What might have been created 10 years ago without today's machine intelligence? Answering this question can help ensure that designers are able to clearly, quickly, and easily justify an AI-driven solution. Note that a designer's justification should be tied to the problem space itself, not to general characteristics of AI. In other words, the description of AI's advantage should not consist only of broad characteristics of AI.

▶ **Recommendation 1:** Make sure an AI-driven solution is clearly and demonstrably better than one without AI by creating and comparing to a non-AI design solution.

There is no doubt that AI has the power to improve product or service experience when used carefully. Such features can help a human user's expertise shine, or remove burdensome barriers. Consider GenAI. Generative AI is a subcategory of broad AI that creates new content, whether text, images, videos, reports, data, or other assets. GenAI is often successful when implemented in two contexts.

First, GenAI is a great tool for introducing productivity improvements like summarization, automation, and basic analysis. AI can accomplish mundane organizational tasks and tackle low-hanging analytical fruit. Thus, a user with deep expertise has more time to use their special skillset. GenAI can provide a shortcut that helps a user take advantage of their own expertise earlier in their process, removing the busy work.

Second, GenAI can grant access to skillsets a user may not otherwise have. For example, a small business may use Midjourney to create images or content for a marketing campaign. Without time or expertise in visual design or illustration, the business may have to rely on stock photos to accompany their product or service, getting lost among their competitors. But the specificity with which an image can be generated could empower them to stand out when putting their product, message, and brand out in the world.

Recommendation 2: Use GenAI features to facilitate and complete mundane tasks, allowing users to leapfrog to tasks requiring their own expertise.

▶ **Recommendation 3:** Use GenAI features to provide access to expertise users would not otherwise have.

In these examples, GenAI *facilitates* a user's expertise. It never replaces it. AI features that override a user's talent threaten the human in the loop, displace creativity, and lead to a centralization of ideas. The use of AI as a means of removing human ingenuity in favor of a quick and uninspired version of a human-created asset should be avoided whenever possible. The centralization of output is noticeable. Education researchers who examined student work noted telltale signs of the use of chatbot assistance, including increased use of Latin terms and consistent, unnecessary vocabulary (Morrison, 2024). The prevalence of its use has created a distinct essay style detectable by many teachers and professors.

Furthermore, humans are still much better at producing creative assets for a given context. AI does not have the intuition to understand how an asset fits into a larger project or how it would be received by diverse end users. Lived human experience and a nuanced understanding of the problem area give human creators an advantage over machine intelligence when context matters most.

The use of AI systems comes with real-world costs for organizations (and for the planet), so such features should be introduced carefully. Especially with regulations like the EU's AI Act in play, using AI necessitates a thorough understanding of its risks, potential harms, and compliance burden. AI should not be thrown at every problem.

However, sometimes designers have no ability to influence a product or service's requirements. Incorporating AI may be a company's goal. When required to include an AI feature where there is no distinct need, one should consider how it can fit seamlessly into the interface. The feature should be placed appropriately in the visual hierarchy and treated as a help mechanism, not something created for a demo or marketing materials. Thoughtfully considering and designing AI features, and clearly communicating their specific value-add, will serve any design professional well.

Recommendation 4: When encouraged to incorporate an AI feature for its own sake, think of it as a help mechanism tucked neatly into an appropriate place in the visual hierarchy. Features need not be displayed at the top level for marketing purposes.

3. Friction Can Be a Friend

For many designers, friction is a dirty word. Friction generally refers to anything within a user's experience that makes it more difficult to complete a task. For example, consider a website whose call-to-action button lies just below the scroll line. A user must scroll slightly in order to click the button as desired by the site's owner. This is annoying friction — the additional requirement of scrolling likely reduces the frequency with which users complete the task (here, clicking a button).

Sometimes friction is employed intentionally because it works. Consider how frequently a user is bombarded with a dialog upon arrival to a retail site, requesting personal information in exchange for a discount. Exiting from these windows is arduous, as they generally use a tiny exit icon that is difficult to click. The window stops a user in their tracks, but can result in a higher number of individuals providing their information. While not exactly a dark pattern, the frictive UI element is designed to stop a user from accomplishing their task for the sake of the site's owner, not their user.

In the world of AI tools and features, friction finds a fundamental purpose. Helping users stop and consider their actions when using AI tools can support a healthier relationship with AI. This is especially true when it comes to decisioning tools. For companies that create, deploy, and manage their own AI models to enhance decision making, friction is key for safety and risk reduction.

As an example, SAS Institute Inc. recently published an open-source guide for trustworthy AI governance called the Trustworthy AI Life Cycle Workflow (hosted on the SAS GitHub page; SAS Institute Inc., 2025). The workflow (Figure 1) outlines steps for evaluating and deploying a more trustworthy AI system, based on the U.S. National Institute of Standards and Technology's (NIST) recommendations, standards, and best practices. Using the guide, teams can produce documentation outlining how the organization has done its due diligence to ensure a model is fair and minimizes harm.

After its initial release, internal reviewers in the Fraud domain at SAS were consulted for feedback. While the workflow was generally well-received, several comments described "barriers" to a seamless operation, and remarked how "extensive" and "big" the workflow was. Reviewers understandably wanted to streamline and shorten the process, making certain steps easier for users and decreasing the overall time to model deployment.

The team spent significant time discussing this feedback. However, many of the included steps erected intentional barriers; for example, rather than simply asking if a model uses personally identifiable information (PII), the workflow outlines steps to identify such variables, asks users to explore whether their inclusion is necessary, requires documentation describing why they must be included in the model, and requests a description of any resulting risk. To simplify the workflow, it would be possible to automate the step (throw a flag if "yes" is selected) and assume the model is high-risk and should be subject to the appropriate regulations. But requiring a user's active thought may result in fewer unnecessary uses of sensitive information.



Figure 1. The Trustworthy AI Life Cycle Workflow is based on a flow of the same name orchestrated in SAS© Workflow Manager [low-resolution product export, only intended to demonstrate complexity].

Friction can and should be employed beyond the decisioning space. One might consider adding friction during points of high consequence for the user. Think about intentional "pause points" early on when creating a task flow. Areas utilizing AI assistance deserve special scrutiny. Points of high consequence include actions that affect multiple people or the organization as a whole, like deploying a model. They may also include a final revision or submission step, or communication with a large group of individuals. For instance, certain social media sites have attempted to reduce harm by asking users to consider kinder messaging when posting a comment. Some companies like TikTok are creating content reminding the user that they have been online for a long time and should go "touch some grass." Areas of high consequence are moments where a human's "humanness" matters, or where results of an incorrect or undesirable outcome would have real-world consequences.

Recommendation 5: Add frictive elements during decision points of high consequence or those involving sensitive data. Adding friction around AI assistance features is particularly important.

As it pertains to GenAI, friction often includes interrupting the presentation of generated content with a reminder that the content is created by an AI model. Indeed, frictive features may need to be more interruptive than users are accustomed to. Google's AI Overview description, for example, labels its content with a title indicating an "AI Overview," yet the font size of the label is smaller than that of the entry (Figure 2). The AI Overview appears by default above other results, at the top of the content hierarchy. AI Overview content often features highlighted text, bold text, or both. All of these aspects create a seamless interaction where a user's eye is drawn to the summary, even if they prefer standard search content. Such is the power of Google's frictionless design.

Instead, they could improve their implementation by adding thoughtful friction — AI-generated text could be completely hidden at first until a user engages with a button, rather than providing several enticing lines that beg the user to reveal the rest of the content. The overview could be an opt-in feature, turned on by interested users intentionally. At the very least, Google could remind readers that content generated by AI can contain mistakes, so to take caution. Currently, the only warning they include is to remind their audience that "Generative AI is experimental" in small font below the expanded summarization (Figure 3). The term "experimental" can be interpreted in many ways and does not explicitly inform a user that content may be wholly incorrect.

Google follows a good heuristic — indicators of AI-generated content should be placed such that a user's eye moves to the label before the content itself. The mind should be primed to consume AI-generated content before it is consumed, as humans are likely to interpret or rely upon it differently. For left-to-right text, this suggests a warning placed



Figure 2. Google's AI Overview preview appears above traditional results and features bold text and highlighting (captured February 25, 2025).



Figure 3. Google provides users with the vague warning that "Generative AI is experimental" at the end of the entry (captured February 25, 2025).

at the top left of the summary. However, Google's small font and hierarchical placement of the summary itself counteract the effect.

▶ **Recommendation 6:** Include a warning about the potential inaccuracy of AIgenerated content, and design the area such that the eye will view the warning message before viewing the content.

The most important step is already an established best practice, though often ignored — a user should always know when interacting with or consuming content created by an artificial intelligence. However, to be a good steward of their users, designers should go beyond the bare minimum and build in time to consider how they would like to make use of AI in the context of their product or service.

4. Users Can Keep Designers in the Loop

AI tools often produce weird results, most noticeably GenAI systems. Many people can recall a funny story or example of when GenAI got something terribly wrong. The most advanced models still have limitations, even as AI is improving exponentially. Designers should plan for this reality from the outset.

It is impossible to foresee all of a GenAI feature's strange and unexpected results, but users can be relied upon to help keep product teams in the know. To this end, every AI feature should be accompanied by a clear user feedback mechanism *in situ*. When a user encounters generated content, they should have an immediate, accessible, obvious way to flag content, elaborate on pain points, or indicate when the AI has misunderstood their request. These feedback mechanisms should be baked into the design in the earliest stages.

 Recommendation 7: Every AI feature should be accompanied by an obvious user feedback mechanism in close proximity.

Relying on users helps ensure that designers are, themselves, also "humans in the loop." Many organizations integrating third-party base models have very little or no control over how such tools interact with users. For example, the voice and tone of generated text is baked in.

True, certain techniques can be used to modify output: one technique gives access to proprietary information a base model does not know using retrieval-augmented generation (RAG), allowing a model to include organization-specific information in a search context. Another technique recommends adding guiding text to help a user when engineering their prompt. Yet another includes filtering, applying a ruleset to system output. But generally speaking, organizations integrating with existing tools might be in the dark about all the ways the tool can behave. Meanwhile, a user still associates the integrated model's behavior with the product or service's brand, regardless of its origin. This is why feedback mechanisms are not a "Phase 2" feature. Without the ability to substantially modify a model's behavior, user feedback becomes the most valuable source of information about how the GenAI feature performs in real-world contexts.

While an organization may not have control over the exact outputs of an integrated third party LLM, its employees do oversee the product or service itself. Designers must advocate for user feedback mechanisms over potentially more marketable, dazzling features that demo well. While new AI capabilities may attract initial attention, only their effectiveness, reliability, and accuracy will keep people happy.

One way to help convince stakeholders to prioritize modifications driven by user feedback is to formalize their inclusion. Product teams could integrate a tag or flag into their feature management system that shows the requested change came from user feedback. These are high-priority items. Users' understanding of a product or service can impact retention, sales, branding, everything. Their opinion should be taken seriously.

Recommendation 8: Product teams must prioritize changes driven by user feedback over dazzling and easily marketable new features. They can do so through formalizing the user feedback integration process.

5. Closing the Loop

In the rapidly evolving landscape of AI, the organizations that thrive will not necessarily be those with the most advanced technology, but those that maintain the most effective human–AI collaboration. AI may not be the best solution to the problem at hand, and it can introduce potential risks and burdens like inaccuracy or compliance requirements. The three core principles outlined here can guide ethical AI design. (1) AI features should be introduced thoughtfully and intentionally, and only where their inclusion is better than a design without them. (2) Friction can encourage users to consume GenAI content thoughtfully and critically, especially when introduced in moments of high consequence. And (3), by including mechanisms for collecting feedback *in situ*, designers can understand how users perceive the AI feature and, in turn, how the feature informs customers' understanding of the organization's brand and reliability.

In summary, to ensure AI features are included with a high impact and minimal harm, be mindful of the eight recommendations, reiterated here:

- 1. Make sure an AI-driven solution is clearly and demonstrably better than one without AI by creating and comparing to a non-AI design solution.
- 2. Use GenAI features to facilitate and complete mundane tasks, allowing users to leapfrog to tasks requiring their own expertise.
- 3. Use GenAI features to provide access to expertise users would not otherwise have.
- 4. When encouraged to incorporate an AI feature for its own sake, think of it as a help mechanism tucked neatly into an appropriate place in the visual hierarchy. Features need not be displayed at the top level for marketing purposes.
- 5. Add frictive elements during decision points of high consequence or those involving sensitive data. Adding friction around AI assistance features is particularly important.
- 6. Include a warning about the potential inaccuracy of AI-generated content, and design the area such that the eye will view the warning message before viewing the content.
- 7. Every AI feature should be accompanied by an obvious user feedback mechanism in close proximity.
- 8. Product teams must prioritize changes driven by user feedback over dazzling and easily marketable new features. They can do so through formalizing the user feedback integration process.

Designers, as stewards of the human in the loop, must ensure that users remain empowered participants in this collaboration, not passive consumers of AI outputs. Through thoughtful application of AI, appropriate friction, and robust feedback mechanisms, product teams can create AI experiences that truly enhance human capabilities rather than merely dazzle or, worse, displace them.

6. Acknowledgments

Thank you to Rajiv Ramarajan, Cheryl Coyle, Zoe Rackley, and Matthew Peterson for assistance and review.

7. References

- Cisco 2024 AI Readiness Index. (2024). Cisco. https://www.cisco.com/c/m/en_us/solutions/ai/ readiness-index.html
- Harris, T., & Raskin, A. (2023, March 24). *The AI dilemma* [audio]. Center for Humane Technology. https://www.humanetech.com/podcast/the-ai-dilemma
- Gibbons, S. (2019, March 24). User need statements. NN/g. https://www.nngroup.com/articles/ user-need-statements/
- Morrison, N. (2024, March 27). The tell-tale signs students are using ChatGPT to help write their essays. *Forbes*. https://www.forbes.com/sites/nickmorrison/2024/03/27/ the-tell-tale-signs-students-are-using-chatgpt-to-help-write-their-essays/
- Rosala, M. (2021, January 17). How might we questions. *Nielsen Norman Group*. https://www. nngroup.com/articles/how-might-we-questions/
- SAS Institute Inc. (2025). SAS trustworthy AI life cycle. *GitHub*. https://github.com/sassoftware/ sas-trustworthy-ai-life-cycle

Author

Sierra Shell is a technologist and UX designer with experience in trustworthy AI design and governance, ethical design strategy, and data visualization design. She has nearly a decade of experience creating complex enterprise products, and currently works with a small team focused on expanding SAS's already robust trustworthy AI capabilities. She received her Master's of Digital Technology Policy from University College London, where she partnered with the British Standards Institution conducting and publishing research on the future of responsible AI standards. Her work required that she develop a deep understanding of common hurdles and disincentives organizations face in implementing AI governance strategies. This has led to a special empathy for organizations and users adopting AI tools and models. As a result, she has a drive to design streamlined solutions that facilitate effective governance and address business concerns intuitively and responsibly.

DISPATCHES FROM INDUSTRY is a *Visible Language* initiative to strengthen the design discipline by bridging industry and academia in the field of interface, experience, and communication design. The column gathers insights from industry practitioners that might inform and guide design practice, while suggesting high-leverage areas for research and scholarship in design. Authors are invited by the editorial board.