Marketing Science Institute Working Paper Series 2025

Report No. 25-128

From Prompt to Product: Reimagining Visual Search with Generative AI

Jan Ole Krugmann and Jochen Hartmann

From Prompt to Product: Reimagining Visual Search with Generative AI

February 17, 2025

Abstract

Traditional e-commerce searches rely on keywords and filters, which often fail to fully capture consumer preferences comprehensively, leading to suboptimal product search results. This research proposes a generative AI-enabled search system that creates visual representations of consumers' desired products and matches them with the best options from tens of thousands of available products. Three experimental studies affirm that exposing consumers to AI-generated visualizations of their textual product descriptions increases both purchase intentions and design satisfaction with the displayed product matches. Prompt adherence–the extent to which consumers believe the input they provide is accurately represented and aligned with the results–functions as a mediating mechanism, though this mediation occurs only if *actual* prompt adherence is high. Therefore, the effectiveness of AI-powered visual search systems depends on their ability to generate accurate visualizations and display accurate product matches. A 2×2 experiment further disentangles these effects, revealing that while both visualization accuracy and product match accuracy independently enhance purchase intentions and design satisfaction, visualization accuracy exerts a stronger influence. These findings deepen theoretical insights into online product searches by demonstrating how visual feedback helps consumers feel heard, while also providing practitioners with actionable insights to increase purchase likelihood using generative AI-powered visual search systems.

Keywords:
Keywords: generative AI, product search, e-commerce, image analytics, prompt adherence

Functioning as a foundation for the digital economy, search systems enable virtually every element of it, from online shopping to content discovery (Bughin et al. 2011). For online retailing, effective search functions strongly influence conversion rates, by minimizing consumer search costs and streamlining the path to purchase (Jiang and Zou 2020; Ngwe, Ferreira, and Teixeira 2019; Ursu, Seiler, and Honka 2024; Yoganarasimhan 2020; Zhou et al. 2024). In contrast, ineffective search systems create significant customer frustration, especially when they fail to capture users' nuanced preferences accurately. Notably, 80% of online apparel shoppers report such dissatisfaction and identify it as a key barrier to purchase (Balchandani, Rölkens, and D'Auria 2024). The limited effectiveness of traditional search tools, like filters and search bars, seems largely due to their inability to deal with semantic variations, synonyms, and spelling errors (Nigam et al. 2019), though in parallel, vastly expanded product assortments exacerbate the complexity associated with purchase decisions (Chernev, Böckenholt, and Goodman 2015; Natan 2024) and the time required to find products (Jiang and Zou 2020).

Considering these salient challenges, new technologies such as generative AI could be the key to revolutionizing and improving online search systems. Fashion executives recognize that product search and discovery represent compelling use cases for generative AI (Balchandani, Rölkens, and D'Auria 2024), though to be effective, such uses must be defined carefully and strategically. When it comes to designing a good search system, prior research that prioritizes consumer needs provides valuable theoretical insights. In particular, it identifies two core concepts. First, effective search systems must allow consumers to articulate an ideal point, defined as their personal, mental representation of the perfect product offering, including both its desired attributes and their relative importance (Carpenter and Nakamoto 1989). Such systems thus need tools that actively solicit consumer preferences and also allow users to express those preferences freely and without significant constraints. Second, consumers require control over information flows during the search process. To address this need, search systems must present relevant options in such a way that users can readily align their preferences with those available choices, as well as choose which products are displayed, in what order, and for how long (Alba et al. 1997; Ariely 2000). By facilitating both the articulation of ideal points and control over the information flow, effective search systems promise to enhance both user satisfaction and retail performance.

In contrast, traditional search systems limit consumers' ability to articulate their preferences or control the information flow. For example, conventional search input methods, such as static filters and search bars, limit the ways that users can express their preferences and lack any feedback loop capabilities for confirming whether the consumer's preferences have been captured accurately. As a result, an expression gap emerges, between consumers' nuanced mental preferences and the outputs generated by the search systems. To address such gaps, we propose leveraging advanced generative AI to inform a search system that we have designed to combine interactive preference articulation with transparent, visual feedback. To explain how the novel product search process enabled by our proposed system improves both consumers' ability to specify their ideal point and the accuracy of the results, consider a hypothetical example of consumers looking for sneakers online. Their search process begins when they describe their ideal product design by writing their own text input (e.g., "Design an off-white sneaker with pastel-colored design elements"). With a generative text-to-image model, our proposed system generates a visual representation of this description.[1] The design visualization represents the ideal point, which the system uses to find visually similar products in the retailer's database, on the basis of image similarity matching. Finally, the system provides two outputs: (1) a text-guided, AI-generated visualization, together with transparent feedback that details how the user's preferences were interpreted, and (2) a selection of products that match the user's ideal point.

To test our concept and the potential of AI-generated visual feedback to enhance product search, we developed a proof-of-concept system. With a pilot study (N = 83 participants from Prolific, $M_{\text{age}} = 35$ years; 50% women), we compared this proposed system against

[1]The search system is agnostic with regard to which AI model are used. In this article, we establish robust findings for both Stable Diffusion 3 and DALL-E.

Figure 1: PROPOSED PROOF-OF-CONCEPT SEARCH SYSTEM VS. TRADITIONAL SYSTEM

A

Proof-of-concept search system

B

Search system of a leading US online retailer

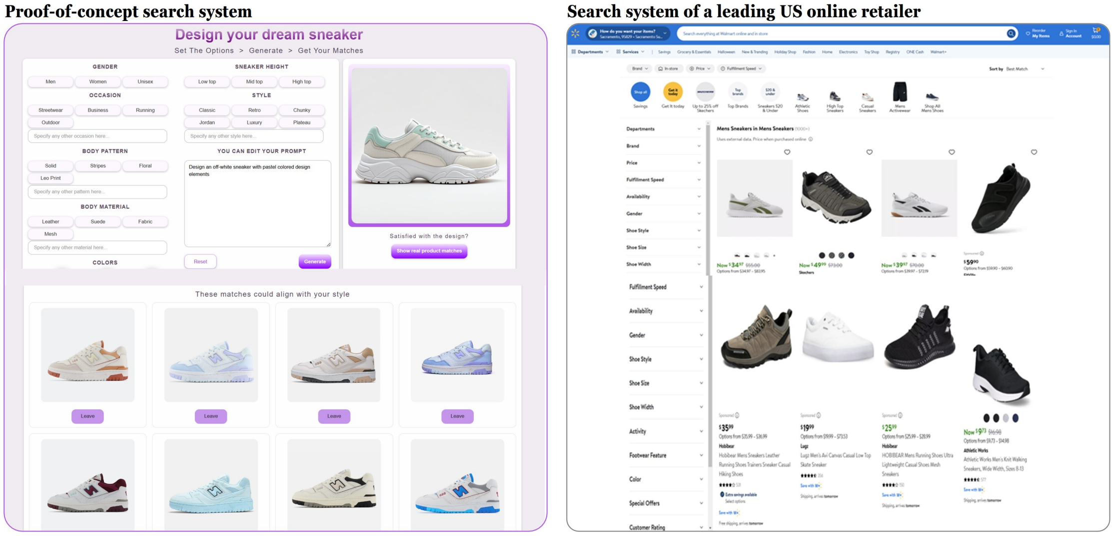

Notes: Panel A presents the proposed proof-of-concept search system, which generates a design visualization on the basis of the user prompt and retrieves visually similar product matches. Panel B displays a traditional, filter-based search system, including the search bar and product matches, from the site of a leading U.S. online retailer.

the existing online search system available from a leading U.S. online retailer (see Figure 1). (Web Appendix A details the experimental setup.) In both conditions, participants had to prompt the system to find a sneaker. The pilot study findings reveal significant improvements in several critical consumer search metrics among participants using the proposed system (Hong, Thong, and Tam 2004; Jiang and Zou 2020). First, purchase intentions for the displayed product matches increased by 27% (p = .007). Second, their search service evaluations improved by 16% (p = .014). Third, search time decreased by 51% (p = .003). These notable gains prompted our further efforts to determine precisely why the AI-generated design visualizations improve product search outcomes, so that we can comprehend the most effective system design and the mechanisms that underlie such effectiveness.

With three experimental studies, we accordingly investigate and determine that perceived prompt adherence, or the extent to which consumers believe their product specifications are accurately represented and aligned with the resulting product matches (Hartmann, Exner, and Domdey 2024; Saharia et al. 2022), is key to improved product search outcomes. The

Study 1 findings establish that receiving AI-generated design visualizations significantly enhances purchase intentions and design satisfaction with the product matches, in an effect mediated by perceived prompt adherence. Study 2 offers a causal test of this mechanism, such that we manipulate the mediator through moderation using actual prompt adherence. The effects persist only in the high actual prompt adherence condition, such that the effectiveness of AI-powered visual search systems fundamentally depends on their ability to understand and translate consumer preferences into accurate design visualizations and product matches. In Study 3, we clarify further that visualization accuracy and product match accuracy independently enhance design satisfaction and purchase intentions, and visualization accuracy emerges as the stronger driver.

With these findings, our research makes three main contributions. First, we advance consumer search theory by integrating information flow and ideal point theory with AI design principles, then use their integration to propose a novel product search approach. Second, we offer robust evidence that AI-generated design visualizations significantly improve important consumer search outcomes. Specifically, visualizations increase design satisfaction and purchase intentions toward actual products, matched from a database of more than 23,000 real products. Furthermore, our research identifies perceived prompt adherence as the mediating mechanism that drives these effects, thereby offering novel insights into the psychological processes that underlie consumer responses to AI-powered search systems. Third, our findings point to the substantial potential business value that retailers could realize by integrating design visualizations into their commercial search systems.

THEORETICAL DEVELOPMENT

Information Flow Control and Ideal Points in Online Search

Consumers navigating online product searches face a fundamental challenge: effectively specifying and communicating their preferences to the search system. This requirement is particularly challenging in settings marked by vast assortments, where information overload already hinders decision-making (Natan 2024; Payne, Bettman, and Johnson 1991). If consumers have control over the information flow though, they can find products that actually match their preferences (Alba et al. 1997; Hoffman and Novak 1996). In particular, determining which information is presented, for how long, and in which sequence increases consumers' ability to match their preferences with available products (Ariely 2000). This process also assumes that consumers possess a mental image of a perfect product that combines desired attributes, in accordance with their relative importance (Carpenter and Nakamoto 1989). When consumers have such a well-established ideal point, their decision-making becomes more straightforward, and they also likely express stronger preferences for selected products in large assortments (Chernev 2003).

However, expression gaps often arise between consumers' ideal point and the search system's ability to identify and articulate it. This gap complicates the identification of relevant products, particularly those for which subjective or creative preferences are salient, as in design-driven product categories like fashion. Traditional filter- or menu-based search systems aim to help consumers articulate their preferences by narrowing the assortments (Häubl and Trifts 2000; Morales et al. 2005), ideally toward options that come closer to the consumer's ideal point. For example, with menu-based prompting (Baty and Lee 1995), users interact with predefined categories or set filters to specify their preferences. Such tools helpfully reduce cognitive effort and create more focused consideration sets (Parra and Ruiz 2009), but they also limit preference articulation and creativity by offering a predefined set of choices (menus or filters) or keywords (search bars).

The emergence of generative AI has revealed a potentially transformative solution to the expression gap. With text-to-image diffusion models, generative AI gathers consumer input, then defines related design visualizations. Sisodia, Burnap, and Kumar (2024) demonstrate the use of generative AI to uncover and generate human-interpretable visual designs for watches, which then supports a process that leads to ideal point designs reflecting consumers' preferences. To start, the derived visualizations provide visual feedback that helps consumers articulate their mental ideal points more accurately. Then in a second step, consumers can use the design visualization to narrow the search results toward more and more relevant options. Because highly precise, personalized search results tend to increase both click-through and purchase rates (Yoganarasimhan 2020; Yuan et al. 2024; Zhou et al. 2024), search systems that prioritize visually well-matched items outperform others, particularly when consumers are seeking specific products (Park and Sela 2020). By matching ideal point visualizations with the product assortment, our proposed method delivers visually similar product matches that represent personalized, precise search results. Thus, the use of design visualizations bridges two fundamental gaps that hinder online search: the expression gap between consumers' mental ideal points and their ability to articulate them and thee matching gap between expressed preferences and available products. We leverage such reasoning to predict:

H_1**:** Product search that integrates exposure to design visualization (through generative AI) prompts greater design satisfaction and purchase intentions toward actual product matches.

The Underlying Mechanism: Consumers' Need to Feel Heard

A critical driver of consumer satisfaction is their fundamental need to feel heard–in our study context, their need to perceive that the search system accurately understands and reflects their preferences (Yin, Jia, and Wakslak 2024). Prior recommendation systems research also highlights the importance of explainability for building trust and acceptance among consumers (Gedikli, Jannach, and Ge 2014; Marchand and Marx 2020). Furthermore, systems that mimic human decision behavior help consumers make better choices (Aksoy et al. 2006). We propose that design visualizations created by generative AI meet all these criteria, in that they provide explicit feedback, in the form of a visual representation of the predicted ideal point. Such a representation demonstrates whether the system has understood their preferences and reinforces the feeling of being heard. As noted, such a feeling reflects users' perceptions of the accuracy and quality of the system outputs (Yin, Jia, and Wakslak 2024).

In turn, the concept of prompt adherence, defined as the degree to which the system's output aligns with the user's input, has emerged as critical to evaluations of the perceived effectiveness of generative text-to-image models like FLUX, DALL-E, Imagen, or Ideogram (Hartmann, Exner, and Domdey 2024; Ideogram 2004; Imagen-Team-Google 2024; Betker, James et al. 2023; Saharia et al. 2022). Also referred to as prompt following (Hartmann, Exner, and Domdey 2024) or image-text alignment (Saharia et al. 2022), prompt adherence offers a perceptual measure of how well the system understands and represents consumers' intentions, as indicated by their input.[2] We posit that receiving design visualizations should enhance perceived prompt adherence, compared with systems without visual feedback. By making the system's interpretation of the user's input visible and tangible to that user, the visualizations constitute clear evidence of whether ideal point preferences have been understood. Hartmann, Exner, and Domdey (2024) confirm the relevance of prompt adherence measures, by using the concept to evaluate how well creative AI outputs align with human briefing inputs; they thus showcase its applicability for assessing AI systems' interpretative accuracy. Accordingly, we propose that perceived alignment between user input and system output determines the effectiveness of AI systems.

Specific to our research context, we use perceived prompt adherence as a tangible measure of a search system's "listening ability." Consumers who receive AI-generated visual feedback should perceive that their input and intentions have been understood, regardless of the overall performance of the search system, because the visualization functions as an explicit confirmation of understanding. Greater perceived prompt adherence (i.e., sense that the input aligns with the output) also should increase design satisfaction and purchase intentions, by reinforcing the fit between stated preferences and search results. That is, a

[2] In its attempt to establish such a measure, Google uses the question "Considering the text above, which image better captures the intent of the prompt?" to solicit evaluations of the performance of its Imagen 3 model (Imagen-Team-Google 2024).

strong perceived fit between input preferences and product matches significantly enhances satisfaction and purchase intentions (Merle et al. 2010; Yoo and Park 2016). We predict particularly that the mechanism driving satisfaction and purchase intentions toward actual product matches is this perceived alignment between consumer input and system outputs, as gauged by our measure of perceived prompt adherence. Formally, we predict the following mediation effect:

H_2: Receiving an AI-generated design visualization increases perceived prompt adherence, which then increases both design satisfaction and purchase intentions toward the product matches.

OVERVIEW OF STUDIES

In three controlled studies, we test these hypotheses. The studies pertain to the effects of receiving a design visualization on design satisfaction and purchase intentions for an actual product match, and they also test for the predicted underlying mediation through perceived prompt adherence. To establish causality, we systematically manipulate actual prompt adherence (Study 2). Furthermore, we carefully disentangle the distinct effects of design visualization accuracy and product match accuracy (Study 3). Table 1 offers an overview of the study objectives, designs, and key variables.

REDESIGNING VISUAL CONSUMER SEARCH

To test our hypotheses, we develop a generative AI-powered visual search system that translates consumer input into personalized design visualizations and identifies the most visually similar product within a database. This system integrates advanced natural language processing, image generation, and image similarity matching techniques; it thus represents a novel approach to product search. We depict its key components and workflow in Figure 2, showing that the process starts with a user interface designed to capture consumer design preferences. Users express their preferences in natural language, entered into free-text fields.

Table 1: OVERVIEW OF STUDIES

	Study 1	**Study 2**	**Study 3**
Main objective	Test the base effect (H_1) and mediation effect (H_2)	Investigate how actual prompt adherence influences the effects on the dependent variables and the mediation	Disentangle the independent effects of design visualization accuracy versus product match accuracy
Manipulation	Design visualization versus no design visualization (control)	Actual prompt adherence: high versus low Design visualization versus no design visualization	Design visualization accuracy: high versus low Product match accuracy: high versus low
Dependent variables	*Design satisfaction* with actual product match *Purchase intentions* toward actual product match	*Design satisfaction* with actual product match *Purchase intentions* toward actual product match	*Design satisfaction* with actual product match *Purchase intentions* toward actual product match
Mediator	*Perceived prompt adherence*	*Perceived prompt adherence*	-
Moderator	-	*Actual prompt adherence*	-
Study setup	Two-factorial between-subjects design, lab setting; 136 Prolific participants	Four-factorial between-subjects design, lab setting; 228 Prolific participants	Four-factorial between-subjects design, lab setting; 279 Prolific participants

As Figure 2 shows, user inputs get transformed into structured textual prompts through an input conversion function, which helps ensure consistency across all prompts. The prompts also are optimized for generative text-to-image models by applying a large language model (LLM) and in-context learning, through the application of exemplary best-practice prompts.[3] For example, OpenAI's gpt-4-turbo API (OpenAI 2024) refines input prompts to generate precise design instructions for the text-to-image model (Brown et al. 2020). These optimized prompts then provide the input for a state-of-the-art generative text-to-image model. For this research, we use DALL-E 3 (Betker, James et al. 2023) in Studies 1 and 2 and Stable Diffusion 3 (Esser et al. 2024) in Study 3. The models produce high-resolution, photorealistic product images; we instructed them to generate a side view of the product design against a white background. Because of the purposefully establish modular nature, our proposed visual search system is agnostic to the AI model applied.

[3]A hypothetical example might read as follows: *Input* "Form: low ankle running sneaker; Sole: white mesh sole with black highlights; bodyMaterial: Fabric body material; bodyPattern: Grey crocodile-like sneaker body pattern with light red accents; laces: dark grey; style: Sportive style for men". Then the *LLM-optimized prompt* would read: "Create an image of a low ankle, sportive style men's running sneaker featuring a white mesh sole with black highlights. The body of the sneaker should be made of fabric material, designed with a grey crocodile-like pattern and light red accents. It should be complemented with dark grey laces."

Figure 2: METHODOLOGICAL WORKFLOW FOR THE GENERATIVE AI-POWERED VISUAL SEARCH SYSTEM

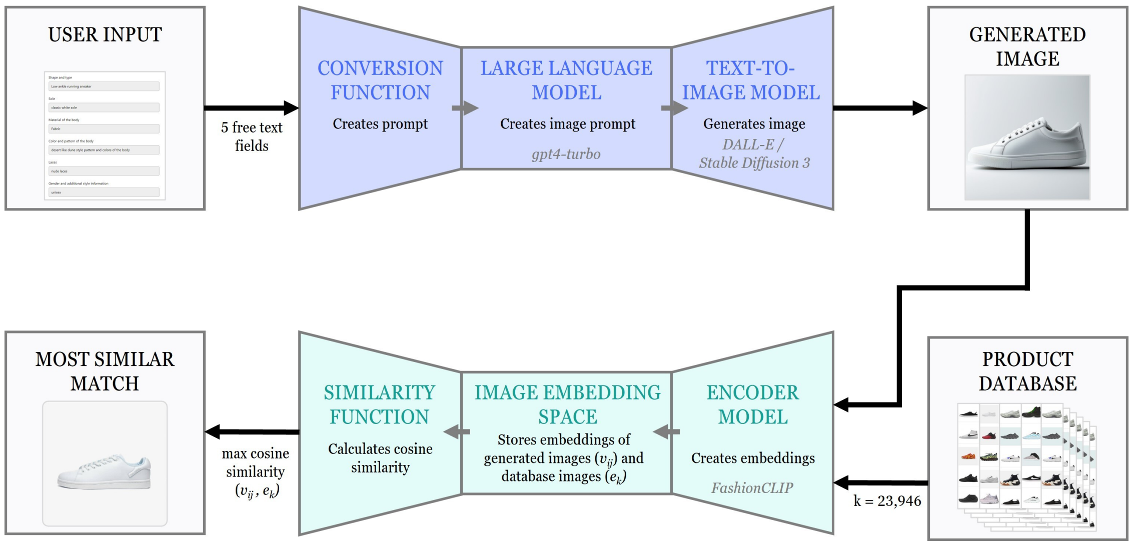

Notes: In this workflow, users provide input, which gets processed through a conversion function and a large language model (e.g., GPT-4 turbo) to generate an image prompt. A text-to-image model (e.g., DALL-E, Stable Diffusion 3) then produces a design visualization that can be encoded into an embedding space. The system matches the embedding with products in a database, according to their cosine similarity, and returns the most visually similar product match as output for the user.

Next, the generated visualization for each user i and prompt j is encoded into a 512-dimensional real-valued vector representation v_{ij} using FashionCLIP, a model that already has been fine-tuned on more than 700,000 fashion images from Farfetch (Chia et al. 2022), such that it possesses good ability to extract fashion-specific visual characteristics from images. The product database, which contains $k = 23,946$ side-view sneaker images (Hashemi 2022), similarly gets processed according to comparable vector embeddings, using the same model. To match consumers' ideal designs with the most visually similar product, the system iterates over all sneakers in the database, calculating cosine similarity between the embedding of the generated visualization (v_{ij}) and the embedding of each database product (e_k), as follows:

$$\text{cosine\_similarity}(v_{ij}, e_k) = \frac{v_{ij} \cdot e_k}{\|v_{ij}\| \|e_k\|}.$$

.

By ranking the products according to their similarity with the database image embeddings, the system can identify the product with the highest similarity score, max(cosine_similarity(v_{ij}, e_k)), as the best match. The system's dual outputs, of AI-generated visualizations and corresponding product matches, offer transparent, explicit feedback and enrich the search experience.

For our studies, we select sneakers as the product category, due to its broad appeal and relevance across many demographics. Sneakers are omnipresent in daily fashion and use and accordingly appear frequently in marketing research (Krause et al. 2023; Merle et al. 2010; Fuchs and Schreier 2023). To facilitate participants' detailed and comprehensive preference articulations, we deconstruct sneaker designs into six core elements that they can specify individually: (1) shape and type, (2) sole, (3) body material, (4) body shape and pattern, (5) laces, and (6) gender and additional style details.

STUDY 1: BASE EFFECT AND MEDIATION

With Study 1, we test both H_1 and H_2. To examine the base effect, we compare the design satisfaction and purchase intentions expressed toward actual product matches by the treatment group, which receives an AI-generated design visualization in the search process, and the control group, which does not receive a design visualization. In a mediation analysis, we also test the predicted underlying mechanism, namely, perceived prompt adherence.

Participants

We recruited 150 participants via Prolific with a relatively balanced gender distribution ($M_{\text{age}} = 36$ years; 40% women); they received monetary compensation for their participation. The application of predefined exclusion criteria (which are the same for all three studies), including failing an English proficiency screener, attention checks, and system usage verification, led us to exclude 14 of them, leaving a final sample of 136 participants for the analysis. All participants resided in the United States.

Design and Measures

We employed a two-factorial between-subjects design with random assignment to one of two groups: a treatment group that received a design visualization created by generative AI and the respective actual product match or a control group that received only the actual product match. Figure 3 provides an overview of the user process flow.

Figure 3: EXPERIMENTAL USER FLOW: FROM DESIGN SPECIFICATION TO DESIGN VISUALIZATION AND PRODUCT MATCH

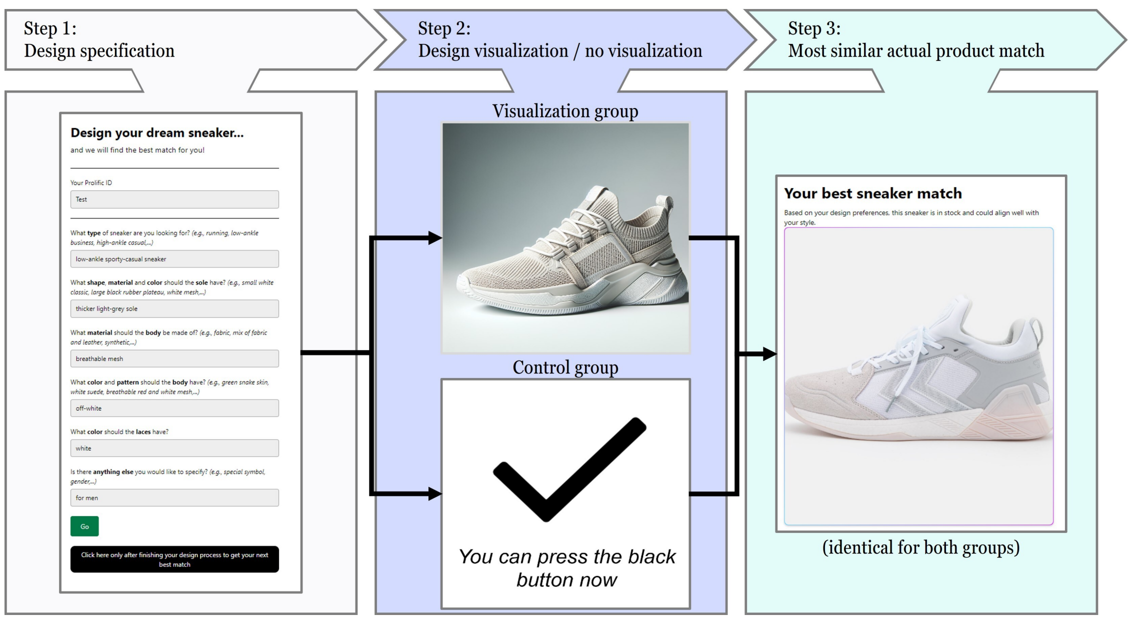

Notes: In Step 1, participants specify their ideal sneaker design by providing text inputs describing attributes such as type, sole material, body material, and color. In Step 2, the system generates an AI-created design visualization based on the specified attributes. Only participants in the treatment group see this design visualization; those in the control group do not. In Step 3, the system matches the generated design to the most visually similar product in the database and displays the actual sneaker match to participants in both groups. The blue background in Step 2 signals the design visualization process; the teal background in Step 3 corresponds to the product matching step that we depicted in Figure 2.

Participants received task instructions (see Web Appendix B for detailed participant instructions for all studies), explaining that they could search for their dream sneaker using the interface of an online search system. Then they specified their preferred sneaker design, using six free-text fields in the interface that refer to the design of core sneaker elements, such as its sole, body material, and color. Participants from both control and treatment groups thus engaged in the exact same sneaker specification procedure, using identical interfaces (Step 1). After a short wait of approximately 20 seconds, constant across both groups, participants in the treatment group received a photorealistic design visualization created by DALL-E, and participants in the control group saw a message indicating that the search was complete (Step 2). Thus, the only difference was that the control group did not receive the design visualization created by the AI model, though this visualization still was created. Using the visualization (which participants in the control group never saw), we could perform the visual match with an actual product available in the database. Therefore, the actual product matches that both groups ultimately received were derived in the same way, using visual similarity calculations. Thus, in Step 3, participants in both groups received a product match that offered the highest visual similarity to their preferences.

Following the completion of the search process, participants evaluated their design satisfaction, purchase intentions, and perceived prompt adherence for the actual product match. They also provided demographic information. To measure design satisfaction with the actual product match, we used a seven-point Likert scale ("How satisfied are you overall with the sneaker match?"), adapted from Hildebrand, Häubl, and Herrmann (2014) and de Bellis et al. (2019). For purchase intentions, we included another seven-point Likert scale ("How likely is it that you would purchase your sneaker match?"), adapted from King, Auschaitrakul, and Lin (2022). Finally, we measured perceived prompt adherence on a seven-point Likert scale, using the item, "How accurately does your form input describe the sneaker match?" taken from Saharia et al. (2022).

Results

In support of H_1, treatment group participants who saw both an AI-generated design visualization and the actual product match ($M = 4.97$, $SD = 1.89$) report significantly higher purchase intentions toward the actual product match than the control group participants who only received the product match ($M = 4.21$, $SD = 2.21$, $t(134) = 2.22$, $p = .028$). Similarly, expressed design satisfaction with the actual product match is significantly higher in the visualization group ($M = 5.29$, $SD = 1.61$) compared with the control group ($M = 4.40$,

$SD = 1.12$, $t(134) = 2.75$, $p = .007$).

For the test of H_2, we apply our proposed conceptual mediation model (see Figure 4), with the treatment group (1 = visualization, 0 = control) as the independent variable, perceived prompt adherence as the mediator, and design satisfaction and purchase intentions as dependent variables. We conduct this mediation analysis using the Hayes Process Macro in R (Model 4), with bias-corrected bootstrapping ($n = 5,000$, 95% confidence interval [CI]), as recommended by Preacher and Hayes (2008), such that the mediation is successful if the CI does not include 0. We observe full mediation for design satisfaction: The indirect effect of perceived prompt adherence is significant (95% CI [.25, 1.56]), whereas the direct effect is not (95% CI [-.0766, .78]). That is, perceived prompt adherence (i.e., the extent to which consumers believe their product specifications have been accurately represented in and align with the resulting product matches) functions as a mediating mechanism. In further detail, exposure to the design visualization significantly increases perceived prompt adherence ($\beta_{\text{viz}} = .73$, $p = .024$), which significantly increases design satisfaction ($\beta_{\text{prompt ad.}} = .76$, $p < .001$). We find consistent results in the mediation analysis with purchase intentions as the dependent variable, and we report these findings in detail in Web Appendix 7.3.

Figure 4: MEDIATION ANALYSIS, STUDY 1

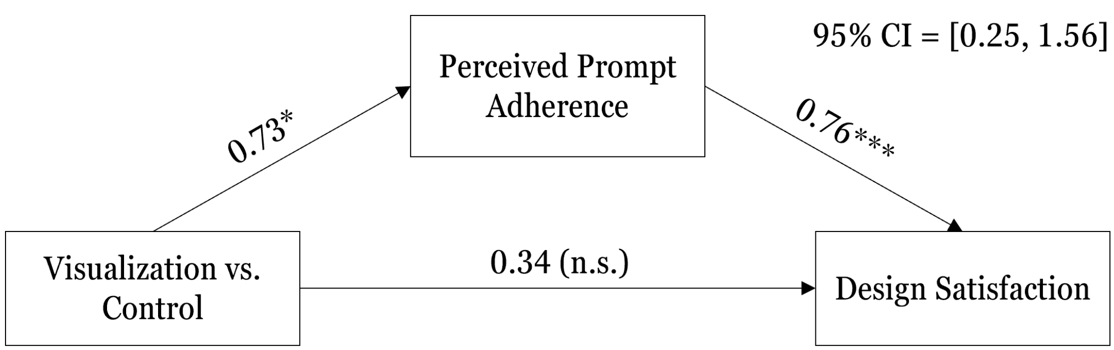

†$p < .1$; * $p < .05$; ** $p < .01$; *** $p < .001$.

Discussion

The results of Study 1 support H_1, by showing that product search systems that include a design visualization feature significantly increase both design satisfaction and purchase intentions toward the actual product match. According to these findings, providing personalized visual feedback can effectively influence consumer attitudes and behaviors in response to search results. The findings in support of H_2 further demonstrate that these increases in design satisfaction and purchase intentions are mediated by perceived prompt adherence. This evidence implies that providing a design visualization makes consumers feel heard. Thus, Study 1 establishes both the base effect and the mediation mechanism. Yet it also leaves some important questions open. First, does perceived prompt adherence causally drive the observed increases in design satisfaction and purchase intentions? Second, might the effects be attributed to a confound, such as the mere presence of an additional AI-generated image? To address both these questions, in Study 2 we systematically manipulate actual prompt adherence, to establish causality and isolate the role of the visualization.

STUDY 2: PROVIDING CAUSAL EVIDENCE FOR MEDIATION

With this study, we investigate whether the increases in design satisfaction and purchase intentions that we observed in Study 1 truly are driven by enhanced perceptions of prompt adherence, by manipulating actual prompt adherence. That is, we manipulate the alignment between the user's input and the system's outputs in search of causal evidence of the mediating role of perceived prompt adherence.

Participants

For Study 2, we offered monetary compensation and recruited 300 participants from Prolific with a balanced gender distribution ($M_{\text{age}} = 37$ years; 50% women). Seventy-two participants were excluded due to the predefined screening criteria, leaving a final sample of 228 participants for the analysis. All participants were residents of the United States.

Design and Measures

In the 2×2 factorial between-subjects design, we assigned participants randomly to four groups, defined by prompt adherence and design visualization exposure. To manipulate actual prompt adherence, we varied the accuracy of the AI-generated design visualization (accurate vs. inaccurate).

Initially, we created a design visualization in the backend and derived the product matches on the basis of their visual similarity with the design visualization. Similar to Study 1, the backend search process was identical across all groups, regardless of whether the participants received the design visualization. In the low actual prompt adherence groups, participants received previously defined, inaccurate design visualizations that purposefully did not reflect their input. That is, the product matches they saw were inaccurate relative to their input but accurate with respect to the predefined design visualization. With this approach, we ensure consistent matching processes across all groups. We also confirmed that the matches were inaccurate with a manipulation check. Perceived prompt adherence serves a dual role: as a manipulation check to confirm the effectiveness of the prompt adherence manipulation and as a mediator that links design visualization exposure to consumer outcomes. Figure 5 illustrates the experimental design and the following, respective groups:

- *Group 1 - Visualization + high actual prompt adherence*: Participants received an accurate design visualization tailored to their input and the most visually similar product match from the database (consistent with the treatment group in Study 1).

- *Group 2 - Visualization + low actual prompt adherence*: Participants received a predefined, inaccurate AI design visualization that remained constant across all participants, regardless of their input. The matched product was also constant and represented the most visually similar product to the predefined design visualization.

- *Group 3 - Control + high actual prompt adherence*: Participants did not receive a design visualization but were shown the most visually similar product match based on their input (consistent with the control group in Study 1).

- *Group 4 - Control + low actual prompt adherence*: Participants did not receive a design visualization and were shown the same inaccurate product match as in Group 2.

Figure 5: EXPERIMENTAL GROUPS, STUDY 2

The task instructions asked participants to search for their dream sneaker, using the interface of an online search system. After interacting with the system, they evaluated their design satisfaction, purchase intentions, and perceived prompt adherence, using the scales from Study 1.

Results

Group 1, which received the accurate design visualization and high actual prompt adherence, exhibits the highest design satisfaction ($M = 4.42$, $SD = 2.01$) and purchase intentions ($M = 4.15$, $SD = 2.22$). These values are significantly higher than those for Group 3, which only received the product match, without visualization, and high actual prompt adherence (design satisfaction $M = 3.56$, $SD = 1.90$, $t(112) = 2.34$, $p = .021$; purchase intention $M = 3.39$, $SD = 1.96$, $t(112) = 1.93$, $p = .056$). In Groups 2 and 4, which experienced low actual prompt adherence, we find similarly low levels of design satisfaction and purchase intentions. The Group 2 participants (inaccurate design visualization) indicate a mean design satisfaction of $M = 2.20$ ($SD = 1.76$) and purchase intentions of $M = 2.20$ ($SD = 1.86$); those in Group 4 (no design visualization) offer a mean design satisfaction of $M = 1.93$ ($SD = 1.51$) and purchase intentions of $M = 1.90$ ($SD = 1.69$). Neither difference, in design satisfaction ($t(112) = .86$, $p = .389$) or purchase intentions ($t(112) = .90$, $p = .370$), is statistically significant between Groups 2 and 4. These findings thus rule out an alternative explanation, namely, that the mere presence of the additional image in the design visualization produces greater design satisfaction and purchase intentions. Instead, our results affirm that accurate visualizations and product matches function as the critical drivers of design satisfaction and purchase intentions.

To test the causal mediation model, we also conduct a moderated mediation analysis with actual prompt adherence as a moderator (1 = high, 0 = low), visualization exposure as the independent variable (1 = visualization, 0 = control), perceived prompt adherence as the mediator, and design satisfaction and purchase intentions as the dependent variables (Figure 6). This analysis uses the Hayes Process Macro in R (Model 7) with bias-corrected bootstrapping ($n = 5,000$, 95% CI). It reveals, for design satisfaction, that the indirect effect is significant only when actual prompt adherence is high ($\beta_{\text{high}} = .92$, 95% CI [.36, 1.50]), not when it is low ($\beta_{\text{low}} = .10$, 95% CI [-.38, .60]). Thus, receiving the design visualization increases perceived prompt adherence ($\beta_{\text{vis}} = 1.03$, $p < .001$), which in turn increases design satisfaction ($\beta_{\text{prompt ad}} = .90$, $p < .001$), though only if actual prompt adherence is high, such that the visualization accurately reflects the consumer's desired product. The interaction between visualization exposure and actual prompt adherence also has a significant effect on perceived prompt adherence ($\beta_{\text{high}} = .92$, $p = .035$). We find similar significance levels in the moderated mediation analysis in which we use purchase intentions as the dependent variable (see Web Appendix D for details).

Figure 6: MODERATED MEDIATION ANALYSIS STUDY 2

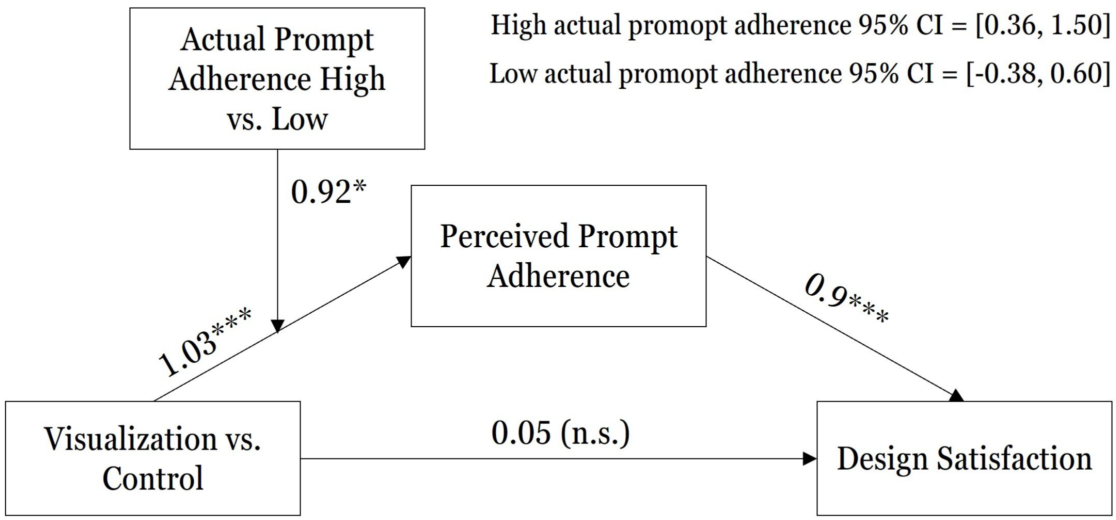

† $p < .1$; * $p < .05$; ** $p < .01$; *** $p < .001$

Discussion

Study 2 thus establishes the causal role of perceived prompt adherence in driving design satisfaction and purchase intentions. By manipulating actual prompt adherence, we show that the effectiveness of AI-generated visualizations depends critically on the system's ability to provide accurate search results, aligned with users' input. In our moderated mediation analysis, the design visualization increases perceived prompt adherence, which then enhances design satisfaction and purchase intentions, though only if actual prompt adherence is high. In other words, an accurate search system is essential for the visualization to exert the intended positive impact. The accuracy of the search system itself acts as a moderator, activating or deactivating the mediation effect. Understandably, inaccurate design visualizations and product matches lead to consistently lower evaluations, regardless of whether a design visualization is available or not. Thus, the Study 1 effects cannot be attributed merely to the presence of an additional AI-generated image. Instead, we note the relevance of user input-system output alignment for establishing accurate search systems.

As Figure 3 demonstrated though, our proposed search system consists of two main elements. Study 1 demonstrates the positive base effect of providing a design visualization, and Study 2 confirms it by manipulating the search system's overall accuracy. To comprehend the effects of design visualizations, we also need to clarify the distinct implications of both design visualization accuracy and product match accuracy. The former refers to how well the generated visualization reflects the user's input; the latter captures how closely the matched product aligns with the user's input. Together, these components determine the system's overall ability to translate user input into meaningful, accurate outputs, so from a system design perspective, understanding their independent roles is essential. Accordingly, with Study 3, we investigate whether the inclusion of the design visualization function benefits only accurate search systems or if less accurate systems can derive value from this feature too. By disentangling the roles of design visualization accuracy and product match accuracy, we seek to specify how each mechanism drives design satisfaction and purchase intentions.

STUDY 3: DISENTANGLING DESIGN VISUALIZATION AND PRODUCT MATCH ACCURACY

By isolating the distinct effects of accurate design visualizations and accurate product matches, in Study 3 we attempt to quantify the overall impact of including design visualizations in the product search process.

Participants

We recruited 302 participants via Prolific with a balanced gender distribution ($M_{\text{age}} = 38$ years; 38% women), for monetary compensation. Twenty-three participants were excluded due to the screening criteria, leaving 279 participants in the final sample. All participants were residents of the United States.

Design and Measures

In the 2×2 factorial between-subjects design, we performed random assignment to four groups, defined by variations in design visualization accuracy (high vs. low) and product match accuracy (high vs. low). Participants specified their ideal sneaker using the online search system interface, then received design visualizations and product matches that reflected their randomly assigned group condition, as illustrated by Figure 7.

Figure 7: TREATMENT GROUPS, STUDY 3

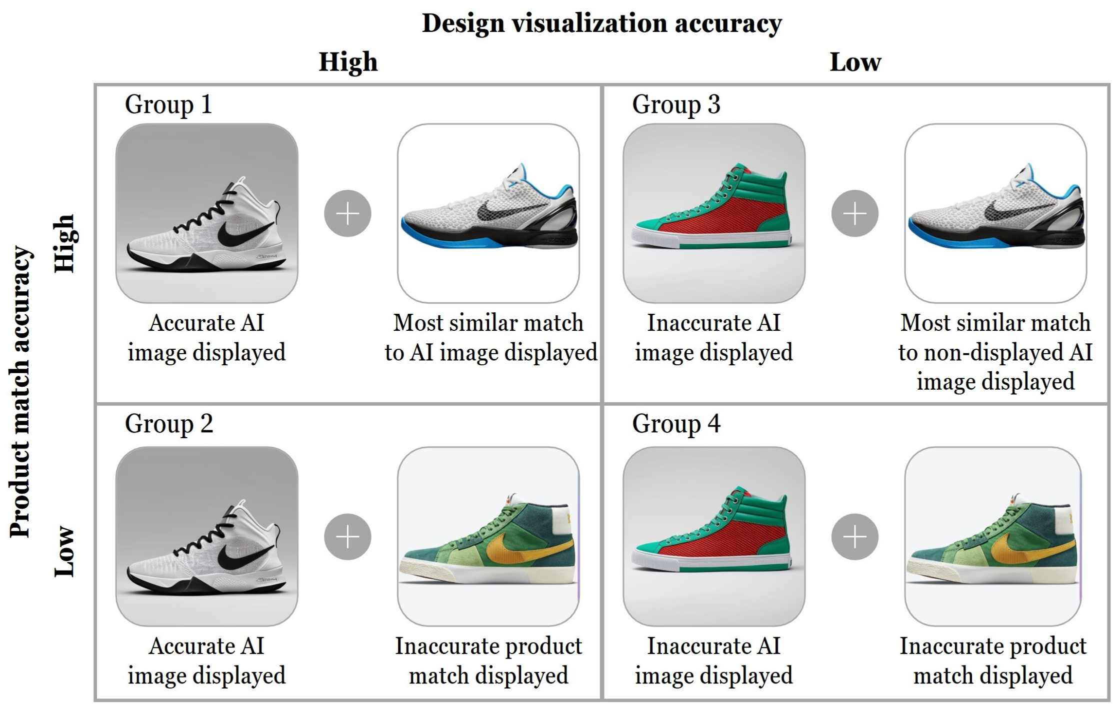

In detail, we manipulated the design visualization accuracy and product match accuracy as follows:

- *Group 1 - Visualization accuracy high + match accuracy high*: Participants received an accurate, AI-generated design visualization aligned with their input and the most visually similar product match from the database (consistent with the visualization group in Study 1).

- *Group 2 Visualization accuracy high + match accuracy low*: Participants received an accurate, AI-generated design visualization aligned with their input and an inaccurate product match that did not align with the design visualization.

- *Group 3 Visualization accuracy low + match accuracy high*: Participants received a predefined, inaccurate AI design visualization (held constant across Groups 3 and 4) that was unrelated to their input and the most visually similar product match based on their input.

- *Group 4 - Visualization accuracy low + match accuracy low*: Participants received the predefined, inaccurate AI design visualization that was unrelated to their input and an inaccurate product match that did not align with either the design visualization or their input.

After interacting with the system, participants evaluated their design satisfaction and purchase intentions, as well as their perceptions of prompt adherence, similar to Study 2.

Results

Figure 8 depicts four different pairwise comparisons, which we use to examine the individual effects of design visualization accuracy and product match accuracy on design satisfaction and purchase intentions. Panel A shows the effect of design visualization accuracy on design satisfaction. When product match accuracy is high, participants assigned to the high design visualization accuracy group (Group 1: $M = 4.70$, $SD = 1.79$) report significantly higher design satisfaction than those in the low visualization accuracy group (Group 3: $M = 3.19$, $SD = 2.03$, $t(145) = -4.77$, $p < .001$). When product match accuracy is low, participants in the high design visualization accuracy group (Group 2: $M = 3.63$, $SD = 2.08$) similarly indicate significantly higher design satisfaction than participants who received the inaccurate visualizations (Group 4: $M = 2.26$, $SD = 1.72$, $t(130) = -4.15$, $p < .001$). Panel B of Figure 8 then reveals the similar effects of design visualization accuracy on purchase intentions. Receiving an accurate AI visualization significantly enhances design satisfaction and purchase intentions, indicating the clear utility of AI-generated visualizations in the product search process. Regardless of the product match accuracy, the design visualization adds value.

Figure 8: BAR CHARTS: EFFECTS OF DESIGN VISUALIZATION AND PRODUCT MATCH ACCURACY ON DESIGN SATISFACTION

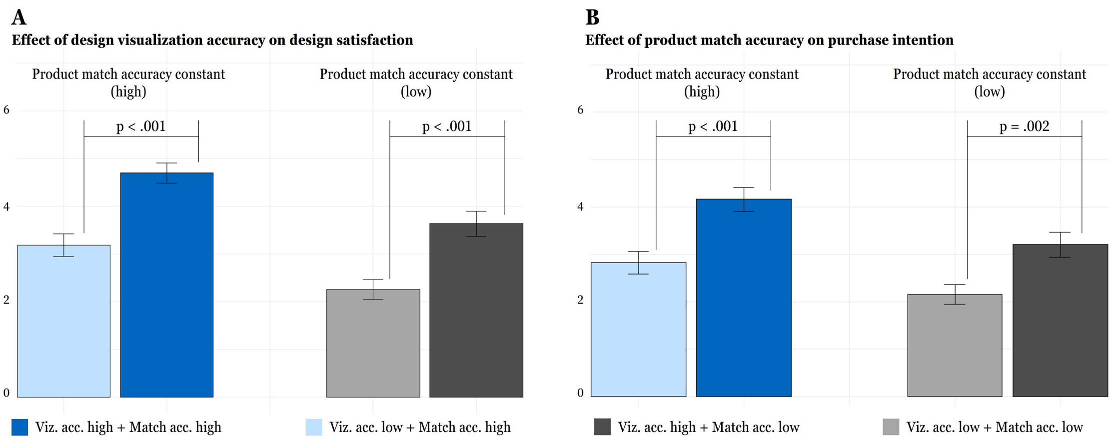

To specify the independent effects of AI visualization accuracy and product match accuracy on design satisfaction and purchase intention, we also estimate four ordinary least square (OLS) regression models (Table 2). Models 1 and 2 integrate design satisfaction as the dependent variable. Then with Model 1, we test for the effect of design visualization accuracy ($1 = \text{high}, 0 = \text{low}$) and product match accuracy ($1 = \text{high}, 0 = \text{low}$) on design satisfaction. The results reveal significant positive effects for both predictors (design visualization accuracy [$\beta_{\text{vis. acc.}} = 1.45$, $p < .001$]; product match accuracy [$\beta_{\text{match acc.}} = 0.99$, $p < .001$]), though visualization accuracy exhibits a stronger effect. In Model 2, we add an interaction term between design visualization accuracy and product match accuracy. The interaction effect is not significant ($\beta_{\text{interaction}} = .14$, $p = .768$). Thus, the two factors appear to contribute independently to design satisfaction, without any multiplicative effects. In our analyses involving purchase intentions, we obtain similar, consistent results too.

Table 2: OLS REGRESSION RESULTS FOR DESIGN SATISFACTION AND PURCHASE INTENTIONS

Dependent Variable	Design Satisfaction		Purchase Intentions	
	(1)	(2)	(3)	(4)
Constant	2.23***	2.26***	2.09***	2.16***
	(0.20)	(0.23)	(0.21)	(0.24)
Design visualization accuracy (1 = high; 0 = low)	1.45***	1.37***	1.20***	1.05**
	(0.23)	(0.33)	(0.24)	(0.35)
Product match accuracy (1 = high; 0 = low)	0.99***	0.93**	0.81***	0.67*
	(0.23)	(0.32)	(0.24)	(0.34)
Design visualization accuracy × Product match accuracy		0.14		0.29
		(0.46)		(0.48)
Observations	279	279	279	279
R^2	0.18	0.18	0.12	0.12
Adjusted R^2	0.17	0.17	0.11	0.11

Notes: Standard errors are in parentheses. *** $p < .001$, ** $p < .01$, * $p < .05$.

Discussion

Study 3 thus establishes the distinct contributions of design visualization accuracy and product match accuracy in shaping consumer satisfaction and purchase intentions. Both types of accuracy independently increase these outcomes. Notably, we also find that accurate design visualizations significantly enhance consumer responses even if the product matches are less accurate. This result underscores the important bridging role of design visualizations: They can overcome the gap between consumer input and system output, thereby providing reassurance and alignment that appear to complement the benefits of the search algorithm's accuracy. By visually confirming the system's interpretation of the user's input, the design visualizations foster a sense of being understood, which further enhances users' satisfaction and purchase intentions, regardless of match quality. Online retailers should take these insights as inspiration to adopt design visualizations: They are not just supplementary features but also offer the potential to extend search systems in critical ways that ultimately enhance consumers' search experience.

GENERAL DISCUSSION

This research addresses a critical challenge in online product search, namely, the expression gap that arises because consumers struggle to articulate their mental ideal points using traditional tools like search bars or filters (Chernev 2003; Park and Sela 2020; Yoganarasimhan 2020). Generative AI-powered design visualizations provide an innovative solution, by acting as explicit feedback that in turn can align consumer preferences with system outputs. We have developed a generative AI-enabled search system that employs text-to-image models to create visual representations of consumers' textual product descriptions and then visually matches those representations with the most similar products in an actual retail database. Across three experimental studies, we demonstrate that the AI-generated design visualizations significantly improve design satisfaction and purchase intentions, because they effectively capture and reflect consumer preferences in a feedback loop. These effects are mediated by perceived prompt adherence, such that the extent to which consumers believe their input is accurately captured and reflected in the search results exerts a meaningful influence. Thus, we identify design visualizations as much more than marginal enhancements; they have the potential to become transformative features in search systems. Furthermore, the Study 2 results demonstrate that inaccurate visualizations, combined with inaccurate product matches, neutralize the positive effects of AI-generated visual feedback, thereby emphasizing the importance of system accuracy for positive consumer outcomes. Yet Study 3 also clarifies that accurate design visualizations can still enhance consumer satisfaction and purchase intentions, even if product matches are less accurate. Visualizations independently enhance consumer evaluations, a finding that underscores their substantial utility. Search systems that combine AI-generated visualizations with image similarity-based product matching thus promise to optimize both preference articulation and fulfillment, representing a significant advancement relative to traditional search methods.

Contributions and Implications

We outline three main contributions of the current research. First, we propose and empirically validate a new search paradigm that leverages generative AI to create visual design feedback that can bridge the expression gap in online search. By generating personalized design visualizations on the basis of consumer input, the proposed system addresses the challenge of imprecise preference articulation and enhances consumers' design satisfaction and purchase intentions in response to search results. Contrary to concerns that generative AI might produce only homogenized content (Epstein et al. 2023; Doshi and Hauser 2024), we learned, in an exploratory, follow-up study with 100 participants, that consumers generate a wide variety of designs, ranging from minimalistic white business shoes to floral-patterned sneakers (Figure 9). Retailers thus can confidently adopt this search approach to improve the product discovery journey and transform their customers' online search into a more intuitive, consumer-centric process.

Figure 9: EXEMPLARY CLUSTER ANALYSIS OF USER DESIGN IMAGE EMBEDDINGS

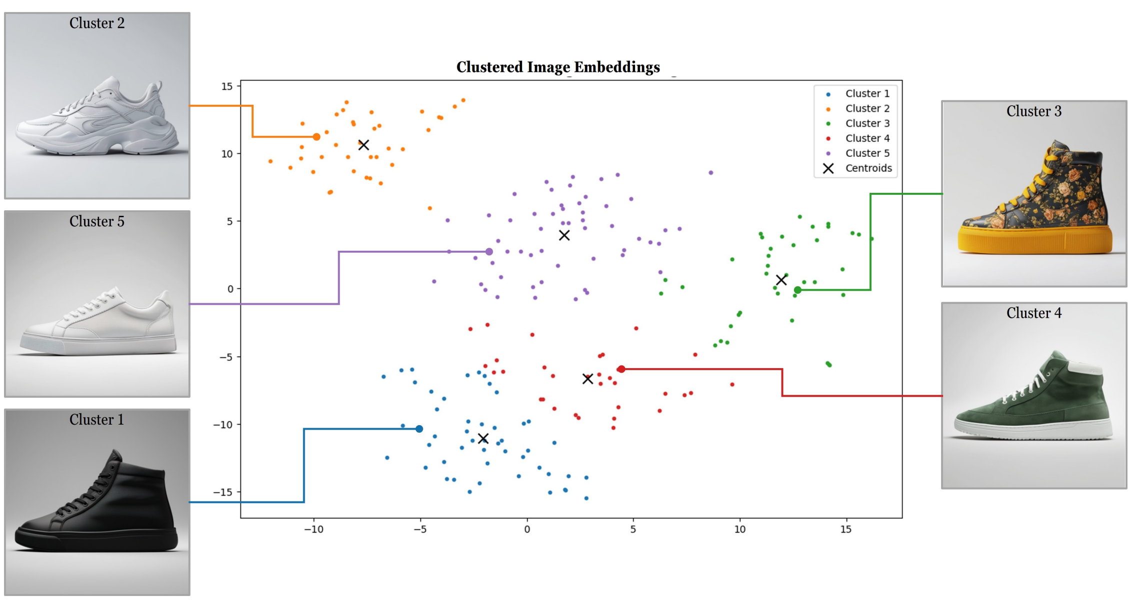

Notes: Each dot represents a 2D embedding of a sneaker image, generated using dimensionality reduction (t-SNE) on high-dimensional feature embeddings extracted with FashionCLIP. The dots are color-coded by cluster, identified using K-means clustering; the cluster centroids are marked by black crosses. The exemplary images from each cluster offer visual insights into the dominant characteristics of each group. The clustering effectively captures distinct features, as evidenced by the differences among clusters, such as black high-top sneakers in cluster 1 and floral-patterned high-tops in cluster 3.

Second, we identify perceived prompt adherence as a critical mediator that drives both design satisfaction and purchase intentions. Even beyond visual search, this mechanism offers actionable insights for designing feedback loops in AI systems in diverse contexts (e.g., customer support chatbots). Any systems that reinforce the consumer's sense of being heard by providing explicit feedback–which might be visual, textual, or auditory–can help build trust and improve adoption.

Third, we identify design visualizations as relevant drivers of consumer satisfaction and purchase intentions, which moves beyond a common emphasis on optimizing backend algorithms in efforts to establish better product matches. Instead, our results highlight the role of front-end features in shaping users' experience. We recommend that practitioners start implementing generative AI-powered design visualizations, because these features provide immediate, tangible value by improving how consumers perceive and interact with product search systems.

Limitations and Further Research

The current study also identifies several areas ripe for further exploration. First, we investigate sneakers, a design-driven category, in our tests of the proposed system. It thus would be valuable to determine if the findings can be replicated in functionality-driven product categories, such as smartwatches or electronics, in which performance attributes tend to play a more significant role than visual designs. Such evidence would help confirm the generalizability of the proposed approach across different contexts. Second, individual preference variations (e.g., personal preference for design versus functionality) could moderate the effectiveness of AI-generated visualizations in search processes. We call for research into how such heterogeneous ideas influence consumers' purchase intentions or design satisfaction. Third, multimodal AI feedback capabilities, such as text, voice, augmented reality, and so forth, could further affect users' experiences (Melumad and Meyer 2024). Studies thus might explore the different effects of visual feedback mechanisms, as provided by our design visualizations, versus textual or auditory feedback mechanisms on consumers' satisfaction and intentions. Fourth, beyond conversion-driven outcomes, we call for research into broader consumer behavior outcomes. Metrics such as brand loyalty, repeated purchases, and customer retention could provide a more comprehensive understanding of the potential long-term value of AI-generated visualizations in search systems. In turn, such studies could be extended with investigations of whether consumers become reliant on these enhanced systems or if, alternatively, novelty effects might diminish over time. Such insights would help inform strategies for developing truly effective generative AI-powered search systems.

Conclusion

This article highlights the remarkable potential of generative AI for improving online product search. Integrating AI-generated design visualizations into the online search process can enable consumers to articulate their preferences more effectively, which in turn leads to their greater design satisfaction and increased purchase intentions directed toward the search results. Our findings demonstrate that design visualizations can function as central features in improved search systems. As businesses increasingly adopt AI-driven solutions, integrating visual feedback mechanisms into their search systems will offer practical value and pave the way toward more a intuitive, consumer-centric product discovery journey.

REFERENCES

Aksoy, Lerzan, Paul N. Bloom, Nicholas H. Lurie, and Bruce Cooil (2006), "Should Recommendation Agents Think Like People?," *Journal of Service Research*, 8 (4), 297–315.

Alba, Joseph, John Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer, and Stacy Wood (1997), "Interactive Home Shopping: Consumer, Retailer, and Manufacturer Incentives to Participate in Electronic Marketplaces," *Journal of Marketing*, 61 (3), 38–53.

Ariely, Dan (2000), "Controlling the Information Flow: Effects on Consumers' Decision Making and Preferences," *Journal of Consumer Research*, 27 (2), 233–248.

Balchandani, Anita, Felix Rölkens, and Gemma D'Auria "The State of Fashion 2025: Challenges at every turn," (2024) https://www.mckinsey.com/industries/retail/our-insights/state-of-fashion.

Baty, James B. and Ronald M. Lee (1995), "Intershop: Enhancing the vendor/customer dialectic in electronic shopping," *Journal of Management Information Systems*, 11 (4), 9–31.

Betker, James, Goh, Gabriel, Jing, Li, Brooks, Tim, Wang, Jianfeng, Li, Linjie, Ouyang, Long, † JuntangZhuang, Lee, Joyce, Guo, Yufei,, Manassra, Wesam, Dhariwal, Prafulla, Chu, Casey, Jiao, Yunxin, and Ramesh, Aditya "Improving Image Generation with Better Captions," (2023) https://api.semanticscholar.org/CorpusID:264403242.

Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, and Dario ... Amodei "Language Models are Few-Shot Learners," H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, "Advances in Neural Information Processing Systems," Vol. 33., pages 1877–1901, Curran Associates, Inc (2020) https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

Bughin, Jacques, Laura Corb, James Manyika, Olivia Nottebohm, Michael Chui, Borja de Muller Barbat, and Remi Said (2011), "The impact of Internet technologies: Search," *High Tech Practice. McKinsey&Company*.

Carpenter, Gregory S and Kent Nakamoto (1989), "Consumer preference formation and pioneering advantage," *Journal of Marketing Research*, 26 (3), 285–298.

Chernev, Alexander (2003), "When More Is Less and Less Is More: The Role of Ideal Point Availability and Assortment in Consumer Choice," *Journal of Consumer Research*, 30 (2), 170–183.

Chernev, Alexander, Ulf Böckenholt, and Joseph Goodman (2015), "Choice overload: A conceptual review and meta–analysis," *Journal of Consumer Psychology*, 25 (2), 333–358.

Chia, Patrick John, Giuseppe Attanasio, Federico Bianchi, Silvia Terragni, Ana Rita Magalhães, Diogo Goncalves, Ciro Greco, and Jacopo Tagliabue (2022), "Contrastive language and vision learning of general fashion concepts," *Scientific Reports*, 12 (1), 18958.

de Bellis, Emanuel, Christian Hildebrand, Kenichi Ito, Andreas Herrmann, and Bernd Schmitt (2019), "Personalizing the Customization Experience: A Matching Theory of Mass Customization Interfaces and Cultural Information Processing," *Journal of Marketing Research*, 56 (6), 1050–1065.

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

Epstein, Ziv, Aaron Hertzmann, Memo Akten, Hany Farid, Jessica Fjeld, Morgan R. Frank, Matthew Groh, Laura Herman, Neil Leach, Robert Mahari, Alex Sandy Pentland, Olga Russakovsky, Hope Schroeder, and Amy Smith (2023), "Art and the science of generative AI," *Science (New York, N.Y.)*, 380 (6650), 1110–1111.

Esser, Patrick, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach "Scaling Rectified Flow Transformers for High-Resolution Image Synthesis," (2024).

Fuchs, Matthias and Martin Schreier (2023), "Paying Twice for Aesthetic Customization? The Negative Effect of Uniqueness on a Product's Resale Value," *Journal of Marketing Research*, 60 (3), 602–624.

Gedikli, Fatih, Dietmar Jannach, and Mouzhi Ge (2014), "How should I explain? A comparison of different explanation types for recommender systems," *International Journal of Human-Computer Studies*, 72 (4), 367–382.

Hartmann, Jochen, Yannick Exner, and Samuel Domdey (2024), "The power of generative marketing: Can generative AI create superhuman visual marketing content?," *International Journal of Research in Marketing* https://www.sciencedirect.com/science/article/pii/S0167811624000843.

Hashemi, Amir Ali "Sneaker Image Collection," (2022) https://www.kaggle.com/datasets/aahashemi/sneaker-image-dataset.

Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (1), 4–21.

Hildebrand, Christian, Gerald Häubl, and Andreas Herrmann (2014), "Product customization via starting solutions," *Journal of Marketing Research*, 51 (6), 707–725.

Hoffman, Donna L. and Thomas P. Novak (1996), "Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations," *Journal of Marketing*, 60 (3), 50–68.

Hong, Weiyin, James Y.L. Thong, and K. Yana.R. Tam (2004), "The Effects of Information Format and Shopping Task on Consumers' Online Shopping Behavior: A Cognitive Fit Perspective," *Journal of Management Information Systems*, 21 (3), 149–184.

Ideogram "Ideogram 2.0," (2004) https://about.ideogram.ai/2.0.

Imagen-Team-Google "Imagen 3," (2024).

Jiang, Baojun and Tianxin Zou (2020), "Consumer Search and Filtering on Online Retail Platforms," *Journal of Marketing Research*, 57 (5), 900–916.

King, Dan, Sumitra Auschaitrakul, and Chia-Wei Joy Lin (2022), "Search modality effects: merely changing product search modality alters purchase intentions," *Journal of the Academy of Marketing Science*, 50 (6), 1236–1256.

Krause, Franziska, Jonas Görgen, Emanuel de Bellis, Nikolaus Franke, Pia Burghartz, Ilse-Maria Klanner, and Gerald Häubl (2023), "One-of-a-kind products: Leveraging strict uniqueness in mass customization," *International Journal of Research in Marketing*, 40 (4), 823–840.

Marchand, André and Paul Marx (2020), "Automated Product Recommendations with Preference-Based Explanations," *Journal of Retailing*, 96 (3), 328–343.

Melumad, Shiri and Robert J Meyer (2024), "How Listening Versus Reading Alters Consumers' Interpretations of News," *Journal of Marketing Research*, page 00222437241280068.

Merle, Aurélie, Jean-Louis Chandon, Elyette Roux, and Fabrice Alizon (2010), "Perceived Value of the Mass–Customized Product and Mass Customization Experience for Individual Consumers," *Production and Operations Management*, 19 (5), 503–514.

Morales, Andrea, Barbara E. Kahn, Leigh McAlister, and Susan M. Broniarczyk (2005), "Perceptions of assortment variety: The effects of congruency between consumers' internal and retailers' external organization," *Journal of Retailing*, 81 (2), 159–169.

Natan, Olivia R (2024), "Choice frictions in large assortments," *Marketing Science*.

Ngwe, Donald, Kris Johnson Ferreira, and Thales Teixeira (2019), "The Impact of Increasing Search Frictions on Online Shopping Behavior: Evidence from a Field Experiment," *Journal of Marketing Research*, 56 (6), 944–959.

Nigam, Priyanka, Yiwei Song, Vijai Mohan, Vihan Lakshman, Weitian Ding, Ankit Shingavi, Choon Hui Teo, Hao Gu, and Bing Yin "Semantic Product Search," Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis, editors, "Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining," pages 2876–2885, New York, NY, USA: ACM (2019).

OpenAI "GPT-4 Turbo and GPT-4," (2024) https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo#gpt-4-turbo-and-gpt-4.

Park, Sang Kyu and Aner Sela (2020), "Product Lineups: The More You Search, The Less You Find," *Journal of Consumer Research*, 47 (1), 40–55.

Parra, José F. and Salvador Ruiz (2009), "Consideration sets in online shopping environments: the effects of search tool and information load," *Electronic Commerce Research and Applications*, 8 (5), 252–262.

Payne, J., Bettman, and E. J. Johnson (1991), "Consumer decision making," *Handbook of consumer behaviour*, pages 50–84.

Preacher, Kristopher J. and Andrew F. Hayes (2008), "Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models," *Behavior research methods*, 40 (3), 879–891.

Saharia, Chitwan, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans et al. (2022), "Photorealistic text-to-image diffusion models with deep language understanding," *Advances in neural information processing systems*, 35, 36479–36494.

Sisodia, Ankit, Alex Burnap, and Vineet Kumar (2024), "Generative interpretable visual design: Using disentanglement for visual conjoint analysis," *Journal of Marketing Research*, 0 (0).

Ursu, Raluca, Stephan Seiler, and Elisabeth Honka (2024), "The sequential search model: A framework for empirical research," *Quantitative Marketing and Economics*.

Yin, Yidan, Nan Jia, and Cheryl J Wakslak (2024), "AI can help people feel heard, but an AI label diminishes this impact," *Proceedings of the National Academy of Sciences*, 121 (14).

Yoganarasimhan, Hema (2020), "Search Personalization Using Machine Learning," *Management Science*, 66 (3), 1045–1070.

Yoo, Jungmin and Minjung Park (2016), "The effects of e-mass customization on consumer perceived value, satisfaction, and loyalty toward luxury brands," *Journal of Business Research*, 69 (12), 5775–5784.

Yuan, Zhe, A. YuanJ Chen, Yitong Wang, and Tianshu Sun (2024), "How Recommendation Affects Customer Search: A Field Experiment," *Information Systems Research*, 0 (0).

Zhou, Wei, Mingfeng Lin, Mo Xiao, and Lu Fang (2024), "Higher Precision Is Not Always Better: Search Algorithm and Consumer Engagement," *Management Science*, 0 (0).

Zierau, Naim, Christian Hildebrand, Anouk Bergner, Francesc Busquet, Anuschka Schmitt, and Jan Marco Leimeister (2023), "Voice bots on the frontline: Voice-based interfaces enhance flow-like consumer experiences & boost service outcomes," *Journal of the Academy of Marketing Science*, 51 (4), 823–842.