

Designing for Decisions

Overcoming Choice Paralysis in Human-Centered AI — With Lessons from Fashion E-Commerce

By Mao Huilin

m757621807@gmail.com

Prologue: The Girl Who Was Late

It's 7:40 AM as the girl stands in front of her wardrobe, in a good mood today. The new skirt she bought online a few days ago finally arrived, and she has been looking forward to wearing it out for the first time. Checking her weather app, she notes it is 26°C, meaning no jacket is needed.

She takes the skirt out and lays it on the bed, and next, she starts looking for a top to match it. However, white feels too ordinary, black seems a bit heavy, and for blue, the color and cut don't seem to go quite right. She suddenly remembers a styling video she swiped through a few days ago claiming that "a tight top with a loose bottom balances your proportions better," but another blogger argued that "monochrome looks are much more sophisticated."

Opening her bookmarks folder, she reviews one picture, two pictures, three pictures, and begins trying on clothes based on those images. The first outfit doesn't look the way it did in the picture; with the second outfit, something feels off; and on the third outfit, still, it's hard to put her finger on what's wrong. She doesn't stop as she keeps looking, thinking maybe the next picture will have the answer, maybe it'll be fine if she just changes one piece, or maybe she should try just one more time.

Twenty minutes pass, leaving the bed piled high with clothes. She glances at the time again and realizes she is almost late. In the end, she silently picks up the jacket she wore yesterday—not because it looks the best, and not because she likes it the most, but simply because, at the very least, it won't be a mistake. She closes the wardrobe and heads out the door.

Throughout the entire morning, she did so many things: she checked the weather, recalled styling tips, searched for inspiration, constantly tried things on, constantly compared, and constantly rejected. And yet, she never actually made the single most important move: deciding. While AI has learned how to recommend, it still hasn't learned how to help people decide.

We Solved the Wrong Problem

For decades, technology has become remarkably good at helping us find. Google helped us find information, Amazon helped us find products, Netflix helped us find something worth watching, Spotify helped us find music we might love, and today, Fashion AI is helping us find clothes that match our style. Across different industries, different products, and different algorithms, they all pursued the same unified goal: to help people find.

However, this focus uncovers a fundamental limitation in our current technological landscape. While recommendation systems can influence decisions, they ultimately cannot complete the decision process for us.

Deconstructing the Decision Process: Decision Friction and The Decision Loop

When we think about decisions, we often imagine a single moment: a person standing in front of two options, a final click on the "Buy Now" button, or a simple "yes" or "no." But real decisions rarely happen that way. Long before a decision becomes visible, something else has already begun. People observe, compare, imagine, hesitate, and search for reassurance, sometimes changing their minds several times before taking any action. The visible choice is only the end of a much longer journey. Yet most digital products only interact with the final moment. They recommend another product, another outfit, another movie, or another option, as if the difficulty lies in choosing between options. But perhaps the harder part is everything that happens before a person is ready to choose. If decision is a process, what exactly happens inside that process, and why do people hesitate, even after they have found what they were looking for? To answer these questions, we first need to understand what happens before a decision is made.

Before trying to improve decisions, it is worth observing one, not as a theory, but as it naturally unfolds. One morning, a girl stood in front of her wardrobe. At first, it seemed like she was simply choosing between clothes, but as the process unfolded, something else became visible. She was not moving from one outfit to another; she was moving from one thought to another. She checked the weather, then remembered today's schedule. She picked up a shirt, then imagined how other people might see her. She changed into another outfit, then wondered whether it matched the new skirt she had been looking forward to wearing. She opened Pinterest, closed it, and returned to the wardrobe. Nothing happened only once; almost everything repeated. She searched, evaluated, imagined, doubted, and then she searched again. By the time she finally left home, one thing had become clear: she had spent far less time changing clothes than changing her mind.

This real-world observation reveals that decision-making is a complex, multi-value conflict rather than a simple selection task. The girl had enough clothes, knew today's weather, had saved dozens of outfit inspirations, and even knew the styling principles

she wanted to follow, yet she still couldn't decide. At first glance, it seemed like hesitation, but the more closely the process was observed, the more another pattern began to appear: every outfit solved one problem, and created another. The white shirt looked professional but perhaps too formal; the T-shirt felt comfortable but perhaps too casual; the new skirt was beautiful but the blue top didn't quite match its silhouette. Changing clothes was never the difficult part; changing the standard by which each outfit was judged was. Every new outfit introduced a new question, and every answer gave rise to another concern. She was not searching for more options; she was trying to satisfy several values at the same time: professional, comfortable, beautiful, natural, weather-appropriate, and suitable for today's schedule. None of these values was wrong, but the difficulty was that they could not all be maximized at once. The resistance created by these competing values is what this framework defines as **Decision Friction**. **Decision Friction** does not come from having too few options, nor simply from having too many. It emerges when several reasonable values compete, and no recommendation can determine which one should matter most today. Recommendation can generate another option, but it cannot decide which value should come first—that decision still belongs to the person.

Because of this friction, people don't make decisions in a straight line; instead, they move through a cyclical loop. After observing the same pattern again and again, one thing became difficult to ignore: people rarely move directly from a situation to a decision. Instead, they return to the same questions, the same concerns, and the same possibilities, over and over again. They search, not because there are no options, but because they hope another option will resolve the conflict. They imagine, not only how an outfit looks, but how the day might unfold, how others might respond, and how they themselves might feel. They evaluate, not against a single standard, but against several values at once—professional, comfortable, confident, appropriate, and beautiful. Then friction appears: one outfit satisfies one value, while another satisfies something else. No option feels completely right, so they adjust, turning to another outfit, another standard, and another search, causing the process to begin again. This

repeated process is what this framework terms **The Decision Loop**. This loop is not hesitation or indecisiveness; it is a natural response to unresolved Decision Friction. The loop does not continue because people enjoy comparing; it continues because each comparison creates new information, new questions, and sometimes, new conflicts. A decision is not another step inside the loop; it is the moment a person finally decides that continuing the loop is no longer more valuable than leaving it. Ultimately, recommendation changes options, but decision changes commitment.

From Recommendation to Decision Intelligence: A Paradigmatic Shift in Design

Every generation of AI has become better at answering the same question: "What should I recommend next?" The answers have become increasingly accurate, more personalized, more relevant, and more efficient; yet, accuracy alone has not eliminated hesitation because recommendation and decision are fundamentally not the same problem. Recommendation asks, "What is the next best option?" whereas decision asks, "What is preventing this person from choosing?" These are entirely different questions: one focuses on options while the other focuses on people; one predicts preference while the other seeks to understand uncertainty; one optimizes information while the other reduces Decision Friction. This boundary is precisely where Recommendation Intelligence ends and where Decision Intelligence begins.

Decision Intelligence is not an AI that makes decisions for people; rather, it is an AI that understands where a decision becomes difficult. It recognizes the values competing inside a decision, identifies the source of Decision Friction, and helps people clarify what matters now—not forever, not universally, but only for this specific decision. Imagine the same girl standing in front of her wardrobe: a traditional recommendation system says, "Here are ten outfits you may like," whereas

a Decision Intelligence system asks, "What matters most today? Professional? Comfort? Confidence? Weather? Meeting? Walking?" Once the most important value becomes clear, many options quietly disappear, not because the AI removed them, but because the decision itself became clearer. The goal of Decision Intelligence is not to increase recommendations, but to reduce unnecessary uncertainty—not by replacing human judgment, but by helping people understand it. The future of AI is not simply generating better answers; it is learning to ask better questions.

For decades, digital products have been designed to optimize information: search engines help people find information, e-commerce platforms help people discover products, and recommendation systems help people explore options. Each generation became better at delivering answers, yet people still hesitate—not because information is missing, but because understanding is. Consequently, the challenge is no longer information design; it is **decision design**. Designing for decisions begins with a simple realization: people do not need another recommendation at every moment; they need different kinds of support at different moments of the decision process. Sometimes they need information, sometimes they need reassurance, sometimes they need to understand why they are hesitating, and sometimes they simply need clarity. If decisions unfold as a process, then AI must be designed to support that process rather than interrupt it, changing the role of AI completely. The goal is no longer to maximize recommendation accuracy, but to reduce unnecessary Decision Friction while preserving human judgment, turning recommendation into one component of the experience, not the entire experience.

To successfully implement decision design, AI must support three distinct architectural layers:

- **Layer One: Understand the Situation.** Every decision begins with context. This layer addresses the foundational reality of the user: *Who is the person? Where are they going? What are they trying to accomplish today?* Without

context, recommendation is only prediction; with context, recommendation finally becomes meaningful.

- **Layer Two: Clarify What Matters.** People rarely hesitate because they lack options; they hesitate because several values compete, such as being professional, comfortable, confident, practical, or elegant. AI should not immediately generate more options, but should instead first help people understand which value matters most for this specific decision, as clarity effectively reduces friction before recommendation even begins.
- **Layer Three: Support Reflection.** Every decision teaches something, a moment that most digital products completely forget, but Decision Intelligence should not. After the decision is finalized, AI can help people recognize long-term patterns, identifying when they hesitate most, which values consistently guide them, and which situations create the strongest Decision Friction. Here, the goal is not to judge, but to foster self-understanding, turning every individual decision into a small piece of future wisdom.

These three layers do not replace recommendation; rather, they completely redefine its place. Recommendation is no longer the final destination, but becomes merely one step inside a much larger, cohesive decision experience. Ultimately, the true value of future AI products will not be measured by how accurately they predict our preferences, but by how well they help us understand ourselves before we decide.

The Blueprint: Implementing Decision Design

Recommendation systems have become remarkably good at answering questions, but perhaps the next generation of AI should begin by asking better ones. Decision Design is not defined by a new algorithm; rather, it is defined by a different sequence of questions, where every question changes what the AI tries to understand, and every answer changes what kind of recommendation becomes possible. Designing for

decisions begins by redesigning the conversation, starting with **Shift One: From "Who are you?" to "What is happening today?"** Traditional AI begins with identity, asking questions like *"Who are you? What do you usually wear? What brands do you like?"* to build a user profile. But decisions are rarely made from identity alone; they are made inside a situation, meaning the same person may dress completely differently for an interview, a weekend trip, or a rainy afternoon. Therefore, the goal of Decision AI is not simply to know the user, but to understand today. Instead of asking, *"What do you usually wear?"* it may first ask, *"What kind of day is today?"* offering context-driven choices such as interview, office, travel, rain, date, or meeting. This is not personalization; it is context, establishing the foundational **Design Principle: Context is not user information. Context is today's reality.**

The next evolution requires **Shift Two: From "What do you like?" to "What matters today?"** While recommendation systems learn preferences, Decision Intelligence learns priorities. People rarely hesitate because they cannot find enough clothes; they hesitate because several values compete at the same time, including comfort, professionalism, confidence, elegance, and convenience. Instead of immediately generating more outfits, AI should first help people answer a simpler question: *"What matters most today?"* Under this framework, the recommendation changes because the priority changes, not because the person changed, demonstrating the second **Design Principle: Recommendation compares clothes. Decisions compare values.**

This shift in priority further demands **Shift Three: From "What should we recommend?" to "How can we reduce uncertainty?"** Traditional recommendation systems often assume that more options lead to better decisions, but Decision Design makes the opposite assumption: the goal is not to maximize choice, but to reduce unnecessary uncertainty. Once today's situation and today's priorities become clear, many possible outfits no longer need to be considered, allowing the recommendation

to become smaller, clearer, and more meaningful. It is no longer the beginning of the experience, but rather the consequence of understanding, which aligns with the third **Design Principle: Recommendation should express a decision, not replace one.**

Finally, product architecture must realize **Shift Four: From "Is this recommendation successful?" to "What did the user learn?"** Most recommendation systems end when a decision is made, but Decision Intelligence begins learning after the decision. Every completed decision reveals something, not only about clothes, but about the person, such as which situations create hesitation, which values appear again and again, and which choices are rarely regretted. Reflection transforms one successful recommendation into knowledge that improves the next decision, proving that while recommendation helps today, reflection helps tomorrow. This underscores the final **Design Principle: A recommendation disappears after today. A reflection improves tomorrow.**

Bringing it all together, Decision Design is not a new recommendation algorithm; it is a fundamentally different way of designing the conversation between people and AI. The shift is surprisingly small—the AI simply asks different questions—but the resulting experience becomes entirely distinct. Instead of asking, *"What should we recommend?"* Decision Intelligence begins with a better question: *"What does this person need before they can decide?"* Perhaps the future of Fashion AI will not be defined by better recommendations, but by better questions. Consequently, before building an AI product, creators must run through a rigorous framework:

- Does the AI understand today's situation?
- Does it understand what matters today?
- Does it reduce uncertainty instead of increasing choices?
- Does the user learn something after the decision?

If the answer is "no" to most of these questions, you may have built a recommendation system, not a Decision Intelligence system.