Human-AI Interaction in Product Recommendation

Jing Dong, Prakirt Jhunjhunwala, Yash Kanoria
Columbia Business School, Columbia University
jing.dong@gsb.columbia.edu, prakirt2203@gmail.com, ykanoria@gmail.com

Abstract

We study the strategic interaction between a customer (or user) and an AI agent in product recommendation. The user conveys preferences through a costly, noisy message, incurring cognitive communication cost. The agent interprets this signal to form a posterior belief and provides a set of recommendations that balances the diversity in the recommendation with the search cost incurred by the user to evaluate the recommendations. The objective is to optimally trade off the utility the user derives from the best recommendation against communication and search costs, modeled through a rational inattention framework. Our main contribution is a formal characterization of the optimal interaction strategy, derived through a high-dimensional approximation that yields the near-optimal trade-off between utility and costs. We show how the optimal strategy depends on feature-space dimension and cost parameters, and we identify distinct regimes where the balance between communication and search shifts sharply. A key insight from our result is that, even in high dimensional setting, the optimal strategy cannot rely only on communication or search in isolation. Instead, both mechanism must be jointly optimized to achieve high utility.

1 Introduction

On e-commerce platforms such as Amazon, customers are often confronted with an overwhelming number of options when searching for a category product. A single query, for instance headphones, can yield thousands of results, and evaluating product specifications and reviews quickly becomes a tedious task [Oulasvirta et al., 2009]. Modern AI agents can process vast amounts of product information (e.g., descriptions, reviews, specifications) at a scale and speed far beyond human capability, enabling them to identify and recommend items with exceptional efficiency. Despite their capabilities, AI agents typically lack comprehensive knowledge of user preferences. To compensate, they rely on short interactions with the user to infer likely needs. Yet customers themselves may struggle to clearly articulate their preference, or may be unaware of features worth considering.

Despite these strengths, the effectiveness of user–agent interactions in recommendation systems is fundamentally limited by two user-side bottlenecks: the *ability to communicate* preferences and the *capability to evaluate* recommended items. The customer seeks to maximize utility from the chosen item but incurs costs both for signaling preferences and for evaluating the recommendations. This paper develops a theoretical framework to study and offer insights into optimizing Human-AI interaction based decision-making in e-commerce setting under cognitive constraints. The core mechanism is twofold: *user communication sharpens the recommendation quality, while the agent's provision of multiple recommendations diversifies the set to account for preference uncertainty*.

Our model is grounded in the framework of Rational Inattention (RI) [Sims, 2003], which models the cost of information acquisition. Our primary contribution is to analyze the user-agent interaction in a high-dimensional setting, where the number of product features is large. We characterize the optimal information exchange from the user and the optimal recommendation set size from the agent as a

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: MLxOR: Mathematical Foundations and Operational Integration of Machine Learning for Uncertainty-Aware Decision-Making.

function of the dimension of the feature space and the system's cost parameters. Our results reveal that for many practical settings, an optimal strategy does not rely on preference communication or recommendation diversification in isolation, but rather requires the joint optimization of both.

Related Work: The literature on the implication of human-AI interaction in decision making is growing [Agarwal et al., 2023, Angelova et al., 2023]. [Kleinberg et al., 2018] show how predictions from machine learning can augment human judgments in socially sensitive contexts, while [Boyacı et al., 2024] analyze the interplay between human decision makers and machine-provided input under bounded rationality. Closest to our work is [Castro et al., 2023], who study the trade-off between output fidelity and communication cost under single-output recommendations, showing that preference heterogeneity can introduce substantial bias. [Liang, 2025] considers delegation to an AI clone, finding that noise in high dimensions degrades performance and humans may outperform. In contrast, we show providing multiple outputs lets users select the best fit, yielding significant gains, and that in high dimensions users provide richer information to ensure credible recommendations.

From an economic and information-theoretic perspective, our model builds on the rational inattention (RI) framework [Sims, 2003, 2005]. RI has been applied extensively to understand decision-making under costly information acquisition Zhong [2022], Miao et al. [2022], Sims [2006]. We also refer the reader to the review [Maćkowiak et al., 2023] for more details. Closely related are approaches in information theory, such as rate–distortion methods [Tishby and Polani, 2010], which also study trade-offs between informativeness and cost. While most prior work applies RI to single-agent decision problems, we extend this line of inquiry to human–AI interaction, characterizing the joint optimization of user communication and search effort.

2 Modeling the interaction and obtaining optimal strategy

We consider a market with a continuous space of products, each represented by a feature vector $\boldsymbol{\theta} \in \mathcal{S}^{d-1} := \{\boldsymbol{\theta} \in \mathbb{R}^d : \|\boldsymbol{\theta}\|_2 = 1\}$. Customers (or users) are similarly characterized by a preference vector $\mathbf{h} \in \mathcal{S}^{d-1}$, which encodes their most preferred product profile or *true preferences*. We assume customer preferences are uniformly distributed over the sphere, $\mathbf{h} \sim p(\mathbf{h}) = \text{Uniform}(\mathcal{S}^{d-1})$.

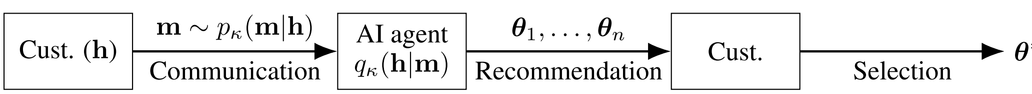

Communication: The interaction starts with the customer transmitting a message (or context) $\mathbf{m}$ to the agent, where the message $\mathbf{m}$ is a noisy representation of the true preferences $\mathbf{h}$. We model $\mathbf{m}$ as drawn from the von-Mises Fisher (vMF) distribution $p_\kappa(\mathbf{m}|\mathbf{h}) \propto \exp\big(\kappa\langle \mathbf{h}, \mathbf{m}\rangle\big)$, where κ denotes the message precision. Higher values of κ imply that $\mathbf{m}$ is more closely aligned with $\mathbf{h}$, which leads to more accurate recommendations. We assume that customers can choose the precision of their message, but they also incur a cost that increases with precision. We model this cost using the RI framework, which posits that the cost of information is proportional to the information gain measured by the KL divergence between the prior belief and posterior distribution induced by the information exchange. In our case, the agent's prior over customer preferences is $\mathbf{h} \sim p(\mathbf{h})$, which is updated to the posterior $q_\kappa(\mathbf{h}|\mathbf{m})$ upon receiving message $\mathbf{m}$. The communication cost is therefore defined as $\lambda_c D_{\text{KL}}\big(q_\kappa(\mathbf{h}|\mathbf{m}) \| p(\mathbf{h})\big)$, where λ_c is the communication cost parameter.

Recommendation: To generate recommendations, the agent forms a posterior $q_\kappa(\mathbf{h}|\mathbf{m})$ using the prior $p(\mathbf{h})$ and the customer's message $\mathbf{m}$. A key feature of our model is the practical constraint that the agent cannot analytically compute this posterior. Instead, it can only generate samples from it. As such, the recommendation menu $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n\}$ is constructed by generating n independent samples from the posterior $\boldsymbol{\theta}_i \sim q_\kappa(\mathbf{h}|\mathbf{m})$ being an independent draw.

Selection (or Search): For any product $\boldsymbol{\theta}$, the utility is determined by its match with the customer's true preferences, measured by the dot product $\langle \mathbf{h}, \boldsymbol{\theta}\rangle$. After receiving the recommendations $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n\}$, the user evaluates each option and chooses the best item $\boldsymbol{\theta}^* = \arg\max_i \langle \mathbf{h}, \boldsymbol{\theta}_i\rangle$, receiving the utility $\max_i \langle \mathbf{h}, \boldsymbol{\theta}_i\rangle$. We make the assumption that this is a noise-free selection process, meaning the user perfectly identifies and chooses the recommendation that maximizes the utility.

The search cost represents the effort a customer expends to find the best product among a set of n recommendations. Initially, the customer views all n recommendations as equally likely to be the best (uniform prior), and after search, the belief collapses to certainty. The total information gain is same as the entropy of uniform distribution over the set $[n]$, i.e., $\log n$. Consequently, using the RI framework, we model the search cost as $\lambda_s \log n$, where λ_s is the search cost parameter.

Objective: The overall objective of an average customer is given by

$$\mathcal{P}(\kappa, n) := \mathbb{E}_{\mathbf{h}\sim p(\mathbf{h})}\mathbb{E}_{\mathbf{m}\sim p_\kappa(\mathbf{m}|\mathbf{h})}\mathbb{E}_{\boldsymbol{\theta}_i\sim q_\kappa(\mathbf{h}|\mathbf{m})}\Big[\max_i \langle \mathbf{h}, \boldsymbol{\theta}_i\rangle - \lambda_s \log n - \lambda_c D_{\mathrm{KL}}\big(q_\kappa(\mathbf{h}|\mathbf{m})\|p(\mathbf{h})\big)\Big]$$

Our central objective is to maximize the objective $\mathcal{P}(\kappa, n)$ by jointly optimizing the two key decision variables: the message precision κ and the recommendation set size n.

2.1 Main Result

In general, the problem of finding the optimal values of the control variables κ and n that solve the optimization problem $\max_{\kappa,n} \mathcal{P}(\kappa, n)$ is analytically intractable. To gain insights, we use a high-dimensional approximation of the problem and solve it to obtain near-optimal solution. Our main result is as follows.

Theorem 2.1. *Suppose $c_s := d\lambda_s$ and $c_c := d\lambda_c$. Consider the following optimization problem.*

$$\mathit{OPT} := \max_{\rho,\alpha,w,y} \rho w + \sqrt{1-\rho^2}\sqrt{1-w^2}x - c_s\alpha + \frac{c_c}{2}\log(1-\rho^2) \;\; s.t. \; I_\rho(w,x) = -\alpha, \quad (1)$$

where $\rho \in [0,1)$, $\alpha \geq 0$ and $I_\rho(w,x) := \frac{\rho(w-\rho)}{1-\rho^2} + \frac{1}{2}\log(1-w^2) + \frac{1}{2}\log(1-x^2) - \frac{1}{2}\log(1-\rho^2)$. Suppose $\tilde\rho$ and $\tilde\alpha$ are the solution of the optimization problem OPT, and $\tilde\kappa = \frac{\tilde\rho}{1-\tilde\rho^2}d$, and $\tilde n = e^{\tilde\alpha d}$, then, we have the following result

$$\big|\mathcal{P}(\kappa^*, n^*) - \mathcal{P}(\tilde\kappa, \tilde n)\big| \leq O\Big(\frac{1}{\sqrt{d}}\Big),$$

where κ^* and n^* are the optimal solution of $\max_{\kappa,n} \mathcal{P}(\kappa, n)$.

Theorem 2.1 provides a high-dimensional approximation OPT (Eq. (1)) of the optimization problem $\max_{\kappa,n} \mathcal{P}(\kappa, n)$, and implies that $\max_{\kappa,n} \mathcal{P}(\kappa, n) = \text{OPT} + O\big(1/\sqrt{d}\big)$. A crucial part of the result in Theorem 2.1 is that, in high-dimensional setting, one can approximate the utility that a customer receives (i.e., $\mathbb{E}[\max_i \langle \mathbf{h}, \boldsymbol{\theta}_i]\rangle$), by establishing a Large Deviation Principle on the set of observed utilities $\{\langle \mathbf{h}, \boldsymbol{\theta}_i\rangle\}_{i\in[n]}$. Essentially, the LDP allows us to replace the complex stochastic optimization over the randomly generated recommendation set with a simpler, deterministic optimization problem that can be solved analytically. Mathematically, for given κ and n such that $\alpha = \frac{\log n}{d}$, and ρ satisfies $\kappa = \frac{\rho}{1-\rho^2}d$, the achieved utility satisfies $\mathbb{E}\big[\max_i \langle \mathbf{h}, \boldsymbol{\theta}_i\rangle\big] = f_{\rho,\alpha}(\rho) + O(1/\sqrt{d})$, where

$$f_{\rho,\alpha}(z) := \max_{w,x} \;\; zw + \sqrt{1-z^2}\sqrt{1-w^2}x \;\text{ such that } I_\rho(w,x) = -\alpha,$$

where $I_\rho(w,x)$ is the rate function of suitably constructed random variables. In order to establish the result, we use a change of variable by decomposing the posterior on customer preferences as $\mathbf{h} = W\mathbf{m} + \sqrt{1-W^2}\mathbf{Y}$, where $\mathbf{Y}$ represents the uncommunicated information and satisfies the orthogonality condition $\langle \mathbf{Y}, \mathbf{m}\rangle = 0$. Similarly, we write $\boldsymbol{\theta}_i = W_i\mathbf{m} + \sqrt{1-W_i^2}\mathbf{Y}_i$. Afterwards, we have $\langle \mathbf{h}, \boldsymbol{\theta}_i\rangle = WW_i + \sqrt{1-W^2}\sqrt{1-W_i^2}X_i$ where $X_i = \langle \mathbf{Y}, \mathbf{Y}_i\rangle$. The LDP establishes that as the dimension grows, the set of points $\{(W_i, X_i)\}_{i\in[n]}$, with high probability, fall within an *observable region*, whose boundary is defined by the level sets of the rate function $I_\rho(w,x)$.

3 Main Insights

We derive insights on the relationship between the optimal interaction strategy and the model parameters using the analytic solution of OPT. The insights we derive are as follows.

Jointly optimizing message precision and recommendation set size: A key insight from our analysis is that the optimal interaction requires balancing message precision κ and recommendation set size n, with the trade-off governed by the cost parameters λ_c and λ_s. Even in high dimensions, communication and search cannot be treated as independent levers; they must be jointly optimized.

Even in scenarios where the users under-communicate, the AI agent can improve outcomes by eliciting more precise information, thereby achieving a better utility–cost trade-off. Figure 1 provides the incremental gain of using a joint optimization compared to best of the two benchmark strategies (pure search or pure communication). The results reveal that the gains from joint optimization (over κ and ρ) are concentrated in a narrow band of intermediate cost regimes.

Balanced Regime: A central finding of this work is the existence of a balanced regime, where both message precision and recommendation set size play critical roles. This regime arises when the scaled cost parameters c_s and c_c in OPT (Eq. (1)) are comparable, which occurs when $\lambda_s \sim 1/d$ and $\lambda_c \sim 1/d$. In this case, the optimal interaction strategy is a hybrid strategy: the user provides partial preference information, with the information gain from the message scaling linearly in d, while the agent supplies a recommendation set whose size grows exponentially in d.

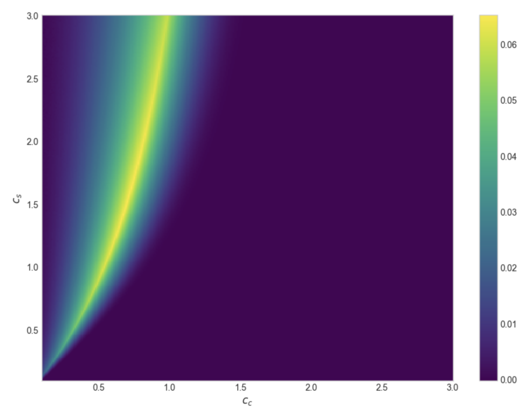

Figure 1: Incremental payoff of the joint strategy relative to the best of the two benchmark strategies: pure search ($\kappa = 0$, and optimal n) or pure communication ($n = 1$ and optimal κ).

The intuition is as follows. Since the prior $p(\mathbf{h})$ implies that total uncertainty about preferences (entropy) scales linearly with d, any meaningful communication must reduce uncertainty at a proportional rate. Our results confirm that in the balanced regime, the user's communicated information scales linearly with d. This is similar to a scenario where the user effectively specifies a subset of features while leaving others unspecified. Further, the unspecified features define a large subspace of feasible products. Within this subspace, the most preferred product differs along the uncommunicated dimensions. To account for this residual uncertainty, the agent must provide multiple recommendations. In high dimensions, adequate coverage requires a recommendation set whose size grows exponentially with the number of unspecified features, akin to providing a constant number of options for each feature left unspecified.

Search-only regime: When communication is highly expensive ($\lambda_c \gg 1/d$ and $\lambda_s \sim 1/d$), it results in $c_c \to \infty$ in OPT (Eq. (1)). In this case, the optimal solution is achieved by setting $\rho = 0$, yielding an optimal message precision of $\kappa^* \approx \tilde{\kappa} = 0$. An example of this regime arises in the context of *window shopping*, where users find it difficult, or even impossible, to articulate their preferences, as they themselves are unsure of what they might like. Consequently, the cognitive burden shifts entirely to exhaustive search, and the agent must provide an overwhelmingly large recommendation set. This mirrors a brute-force search experience on a generic e-commerce platform with no AI-assistance, where finding a suitable product depends solely on the user's search effort.

Communication-only regime: When search is very expensive ($\lambda_c \sim 1/d$ and $\lambda_s \gg 1/d$), we obtain $c_s \to \infty$ in OPT (Eq. (1)). The optimal solution in this regime requires $\alpha = 0$, which implies that the recommendation set size n remains $O(1)$. A natural example is a *deep research* setting, where the desired output is a long, detailed report. Here, it is prohibitively costly for the user to evaluate multiple lengthy reports, so the strategic burden shifts entirely to communication. The optimal approach is for the user to provide high-fidelity preference information upfront, enabling the agent to tailor a single precise output. In practice, this corresponds to the agent asking targeted clarifying questions to elicit detailed user preferences before undertaking the deep research task.

In the above three regimes (balanced, search-only and communication-only), the user's payoff is a non-zero fraction of the highest achievable payoff (max payoff under no costs). In a *non-operational regime*, where both costs are very high ($\lambda_s \gg 1/d$ and $\lambda_c \gg 1/d$), any meaningful interaction is too costly to be worthwhile. The optimal strategy degenerates to random selection, and the user's expected payoff consequently approaches zero. Alternatively, in a *frictionless regime*, where either $\lambda_s \ll 1/d$ or $\lambda_c \ll 1/d$, the user's payoff is highest achievable as this is similar to having no costs. Here, $\lambda_s \ll 1/d$ represents a scenario where the user can effortlessly evaluate all options to find the perfect match, and $\lambda_c \ll 1/d$ represents a futuristic case where the user relegates the purchase task to a perfect AI clone that already knows the user's true preferences.

References

Nikhil Agarwal, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. Combining human expertise with artificial intelligence: Experimental evidence from radiology. Technical report, National Bureau of Economic Research, 2023.

Victoria Angelova, Will S Dobbie, and Crystal Yang. Algorithmic recommendations and human discretion. Technical report, National Bureau of Economic Research, 2023.

Tamer Boyacı, Caner Canyakmaz, and Francis De Véricourt. Human and machine: The impact of machine input on decision making under cognitive limitations. *Management Science*, 70(2): 1258–1275, 2024.

Francisco Castro, Jian Gao, and Sébastien Martin. Human-ai interactions and societal pitfalls. *arXiv preprint arXiv:2309.10448*, 2023.

Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. Human decisions and machine predictions. *The quarterly journal of economics*, 133(1):237–293, 2018.

Annie Liang. Artificial intelligence clones. *arXiv preprint arXiv:2501.16996*, 2025.

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273, 2023.

Jianjun Miao, Jieran Wu, and Eric R Young. Multivariate rational inattention. *Econometrica*, 90(2): 907–945, 2022.

Antti Oulasvirta, Janne P Hukkinen, and Barry Schwartz. When more is less: the paradox of choice in search engine use. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 516–523, 2009.

Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3): 665–690, 2003.

Christopher A Sims. Rational inattention: a research agenda. 2005.

Christopher A Sims. Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163, 2006.

Naftali Tishby and Daniel Polani. Information theory of decisions and actions. In *Perception-action cycle: Models, architectures, and hardware*, pages 601–636. Springer, 2010.

Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582, 2022.