

# Agentic Artificial Intelligence for Inclusive Customer Engagement: Bridging Personalization, Accessibility, and Ethical Interaction

Jessy Christadoss<sup>1</sup>, Dr. Manas Ranjan Panda<sup>2\*</sup>

<sup>1</sup>Independent Researcher, Senior Quality Engineer, Integral Ad Science, USA, LinkedIn: [www.linkedin.com/in/jessy-christadoss](https://www.linkedin.com/in/jessy-christadoss)

<sup>2</sup>Independent Researcher, Partner, Wipro Limited, India, LinkedIn: [linkedin.com/in/manaspanda01](https://www.linkedin.com/in/manaspanda01)

## Abstract

An agentic artificial intelligence (AI) system that can pursue complex goals, transforms the way customer engagement functions. This paper presents a mathematical model for agentic AI that balances deep personalization with universal accessibility. We propose a framework in which an AI agent, governed by utility functions, dynamically adjusts its strategy of interaction (such as communication channels and content complexity) depending on user states in real time (e.g., preferences, context, and ability). The model integrates ethics to prevent discrimination and protect user privacy. One major contribution is formalizing the "inclusivity/personalization trade-off," showing how systems can optimize for broad access without losing individual relevance. Simulations using the model showed significant improvements in engagement metrics across different users compared to traditional reactive systems. The paper concludes by discussing the ethical implications and future directions for the deployment of autonomous agents in customer-facing roles.

**Keyword:** Agentic AI, Ethical-AI, Modelling Mathematics, Personalized Content, Accessible Design, User Engagement, Inclusiveness.

## 1. Introduction

The paradigms of customer engagement are shifting from reactive support to proactive, contextual interaction. Traditional rule-based chatbots and recommender systems operate on a narrow set of triggers and historical data, often failing to accommodate the diverse needs of dynamic users. This is especially true for users with disabilities, those using assistive technology, or individuals in nonstandard contextual situations. The emergence of agentic AI, defined as systems capable of setting goals, making plans, and executing autonomous actions, offers an unprecedented opportunity to bridge this gap.

This paper argues that agentic AI can be systematically designed to harmonize the often-competing demands of hyper-personalization and universal accessibility. We move beyond conceptual discussion by presenting formal mathematical models that capture the core mechanics of inclusive agentic engagement. The model provides a blueprint for developing AI agents that not only understand what users want but also recognize how those needs should be met, all while operating within strict ethical boundaries. This work aims to lay a foundational framework for the next generation of customer technology interactions that are both equitable and intelligent.

## 2. Literature Review

Research has laid the groundwork crucial for understanding the components of this model. The concept of software agency, rooted in the Belief-Desire-Intention (BDI) model formalized by Rao and Georgeff (1995) [1], provides robust architectures for autonomous agents. In the realm of customer engagement, the journey begins with simple algorithmic personalization, such as collaborative filtering, which evolved into sophisticated deep learning models for sequence prediction, as highlighted by Zhang et al. (2019) [2]. Simultaneously, the field of accessibility in HCI has long advocated the principles of universal design, articulated by Shneiderman et al. (2016) [3], emphasizing the need to build systems usable by all people.

The ethical dimension of AI has become increasingly prominent, with seminal works such as Mittelstadt et al. (2016) [4] outlining principles of beneficence, non-maleficence, autonomy, and justice as they apply to algorithms. However, significant gaps exist in the literature, particularly in the integration of these disparate fields. While some studies have explored adaptive interfaces for users [5], others have focused on ethical constraints in machine learning [6]; there has been a lack of a unified, formalized model that combines proactive, goal-seeking agency with principled inclusivity and ethics in customer contexts. This paper seeks to fill that gap by synthesizing pre-2024 research threads into a cohesive agentic AI framework.

## 3. A Mathematical Model for Agentic Interaction

The proposed agentic AI system is modelled as a utility-maximizing agent operating within a constrained Markov Decision Process (MDP) framework. In this formulation, the agentic system aims to select actions that maximize the expected value of its utility function over time, conditioned on the dynamic state of the user and the surrounding environment. This approach enables the system to balance immediate and long-term objectives while adhering to ethical and contextual constraints. The fundamental components of the proposed model are defined as follows:

Let  $S$  be the state space, representing the user context, including demographic data  $D$ , real-time accessibility needs  $A_n$  (e.g., "requires screen reader," "high ambient noise"), and psychological state  $P$  (e.g., "frustrated", "time-pressured"). Let  $A$  be the action space, encompassing all possible engagement actions, such as sending a message, modifying an interface, or escalating to a human agent. Each action  $a \in A$  has an associated personalization score  $Per(a, s)$  and an accessibility score  $Acc(a, s)$ .

### \*Corresponding author:

Dr. Manas Ranjan Panda

**Citation:** Christadoss J and Ranjan Panda M (2026) Agentic Artificial Intelligence for Inclusive Customer Engagement: Bridging Personalization, Accessibility, and Ethical Interaction. American J Sci Edu Re: AJSER-289.

**Received Date:** 08 December, 2025; **Accepted Date:** 27 December, 2025; **Published Date:** 05 January, 2026

The agent's objective is to find a policy  $\pi(s)$  that maps states to actions and maximizes the expected cumulative utility:

$$U(s, a) = \alpha \cdot \text{Per}(a, s) + \beta \cdot \text{Acc}(a, s) - \gamma \cdot \text{EthCost}(a, s)$$

where  $\alpha, \beta, \gamma$  are weighting coefficients that calibrate the trade-off between personalization, accessibility, and ethical cost. The function  $\text{EthCost}(a, s)$  quantifies the potential ethical violation of an action, such as privacy intrusion or algorithmic bias.

### 3.1. State and Action Spaces Formulation

The state space  $S$  is a complex construct; it not just one thing. It is composed of many parts like the user demographic  $D$  which is mostly static and the real-time accessibility needs  $A_n$  which can change very quickly. For example, a user might be in a quiet library then walk into a loud construction area, so the  $A_n$  for audio output needs to be updated. Then there is the psychological state  $P$ , such as whether the user is frustrated or in a hurry. This is inferred from behaviour, such as their typing speed or if they are repeating the same question. Thus, the state  $S$  is a combination of all these elements  $S = f(D, A_n, P)$ .

The action space ( $A$ ) includes all the things an AI agent can do. This can be something simple, like sending a text message or showing an image, or something more advanced, like automatically changing the website's contrast or font size, or even deciding to connect the user to a human helper. Each action  $a$  in the set  $A$  is evaluated using two main scores. The personalization score  $\text{Per}(a, s)$  shows how well the action matches the user's preferences and past behavior, while the accessibility score  $\text{Acc}(a, s)$  shows how easy the action is for the user to use in their current situation. Sometimes, an action might be very personalized but not very accessible or the other way around.

### 3.2. The Utility Function and Policy Optimization

The core of the agent's decision-making is the utility function  $U(s, a)$ . This function is a weighted sum: it combines the good factors and subtracts the bad ones. The term  $\alpha \cdot \text{Per}(a, s)$  is for personalization, where alpha ( $\alpha$ ) is a tuning parameter; if the

company values personalization highly, they set alpha high. Similarly, the term  $\beta \cdot \text{Acc}(a, s)$  is for accessibility, with its own weight ( $\beta$ ). The critical part is the subtraction of  $\gamma \cdot \text{EthCost}(a, s)$ . This acts as a penalty or a brake on the system to stop it from choosing an action that is unethical, even if it is very personalized and accessible.

The policy  $\pi(s)$  is the brain of the agent. It is a function that looks at the current state  $s$  and then picks the action  $a$  that promises the highest expected utility  $U(s, a)$  over time. The agent is not just thinking about the immediate reward but also the future consequences of its actions on user engagement. Thus, finding the optimal policy  $\pi^*$  involves solving the constrained MDP, which means maximizing the cumulative utility while always respecting the rules embedded in the  $\text{EthCost}$  function. The weights  $\alpha, \beta$ , and  $\gamma$  are not fixed forever; they can be adjusted by the system administrator to reflect changing business goals or ethical standards.

### 3.3. Quantifying Ethical Cost in Decision-Making

The  $\text{EthCost}(a, s)$  function is how we make the agent behave morally; it translates ethical principles into a numerical penalty. For example, if an action  $a$  would require accessing the user's medical history without clear consent for the task, the  $\text{EthCost}$  for that action would be very high, preventing the agent from choosing it. Similarly, if an action is likely to produce a biased outcome against a certain demographic group  $D_i$ , the cost function would be triggered to make that action less attractive to the agent.

Implementing  $\text{EthCost}$  requires a framework for detecting potential harm. This could involve checking actions against a set of rules defined by ethicists or using a separate machine learning model to predict the fairness and privacy impact of an action. The weight  $\gamma$  controls how strong these ethical constraints are. A low  $\gamma$  means the agent might ignore ethical concerns for a small gain in personalization, while a high  $\gamma$  makes the agent very cautious and conservative in its actions to avoid any possible ethical violation.

**Table 1:** Key Variables in the Agentic AI Model.

Variable	Notation	Description
State	$S$	The combined state of the user and environment.
Action	$A$	The set of all possible engagement actions.
Personalization Score	$\text{Per}(a, s)$	A measure of action relevance to user preferences (0-1).
Accessibility Score	$\text{Acc}(a, s)$	A measure of action usability given user needs (0-1).
Ethical Cost	$\text{EthCost}(a, s)$	A penalty for actions that violate ethical guidelines.
Utility Function	$U(s, a)$	The overall objective function to be maximized.

## 4. The Personalization-Accessibility Trade-Off

The development of agentic artificial intelligence systems is facing a fundamental tension. This is the conflict between making things personal for a user and ensuring they are accessible to everyone. For example, a system might want to send a user a highly personalized data visualization, but that complex graphic is completely unusable for a person who relies on a screen reader. On the other hand, if the system always provides information in the most basic and accessible format, such as simple text and audio, this can be inefficient for an expert user. That expert user might prefer a dense, visual dashboard for a quick overview. This conflict means the AI cannot simply maximize one goal without compromising the

other; it must make a strategic choice. Therefore, the agentic AI must be designed to navigate this compromise intelligently. It does this by not viewing personalization and accessibility as a single fixed setting. Instead, it treats them as dynamic priorities that can change from one interaction to the next, based on who the user is and what situation they are currently in. The core innovation is to make the method of personalization itself adaptable, ensuring that the tailored experience is, by its very nature, also an accessible one for that specific individual at that specific moment.

Agentic AI faces a core conflict: deep personalization often creates accessibility barriers, and universal accessibility can reduce efficiency for advanced users. The system navigates this by dynamically balancing two priorities. It adjusts its strategy in real time, ensuring that the way it personalizes an interaction is always the most accessible method for that particular user's current context and needs.

**Mathematical Formalization of the Trade-Off**


The trade-off is quantitatively managed within the agent's utility function  $U(s, a)$ . The terms for personalization and accessibility are weighted against each other, and their combined value is reduced by any ethical costs.

$$U(s, a) = \alpha(s) \cdot P(s, a) + \beta(s) \cdot \text{Acc}(s, a) - \gamma \cdot C(s, a)$$


Where:

- $U(s, a)$ : The total utility of taking action  $a$  in state  $s$ .
- $P(s, a)$ : The personalization score of action  $a$  for a user in state  $s$  (range 0 to 1).
- $\text{Acc}(s, a)$ : The accessibility score of action  $a$  for a user in state  $s$  (range 0 to 1).
- $\alpha(s), \beta(s)$ : Dynamic weighing coefficients for personalization and accessibility. Their values are derived from the user state  $s$ .
- $C(s, a)$ : A cost function representing ethical or resource penalties.
- $\gamma$ : A fixed weight for the cost penalty.

The key is that  $\alpha$  and  $\beta$  are functions of the state  $s$ , making the trade-off fluid.



**Figure-1:** Dynamic Weight Adjustment Process.



**Figure-2:** The Trade-Off Frontier.

The graph below illustrates the relationship between personalization and accessibility scores for different user contexts. The "feasible region" shows all possible combinations that the system can achieve. The optimal point is chosen dynamically on this frontier based on the user state.

**5. Dynamic User State Modeling**

The effectiveness of an agent depends largely on its capacity to perceive and accurately model the user's state  $S$ . This state is not static; rather, it is a dynamic construct that evolves continuously in real time. At any given moment, the state  $S_t$  is updated through user interactions, sensor data, and feedback, and may even reflect changes in the user's mood. For instance, if a user repeatedly ignores complex visual messages, the agent may infer a need for simpler visual layouts or higher contrast, leading to an update in the accessibility component  $A_n$ .

Dynamic modeling enables the agent to behave in a more human-like and context-aware manner, adapting its responses to changing circumstances. For example, when a user is in a bright outdoor environment, the agent may increase contrast or switch to audio output. Although such inferences may not always be perfect, they allow the system to respond intelligently to situational cues. Unlike traditional systems, where accessibility parameters were stored as fixed values in a database, this approach treats the model as a time-varying function that evolves with both the user and the environment.


**6. Ethical Constraint as Utility Penalty**

Ethical considerations are not applied as external rules but are embedded directly within the mathematical framework of the agent. The function  $\text{EthCost}(a, s)$  serves as a penalty term that activates when an action violates ethical principles. For example, if an action involves the use of private health data for marketing purposes, it would incur a high ethical cost, thereby reducing its likelihood of being selected by the agent. This integration ensures that ethical reasoning becomes an inherent part of the decision-making process, rather than an afterthought.

The total utility then looks something like:

$$U'(a, s) = U(a, s) - \text{EthCost}(a, s)$$

In some cases, the system also identifies when a particular action negatively impacts a specific user group  $D_i$ , and it applies an additional penalty. This formal design ensures that the agent continues to uphold principles of fairness and human autonomy, even when the optimization process tends to prioritize engagement or performance metrics too aggressively.



**Figure 3:** Flowchart of the Agentic AI System for Customer Engagement.

### 7. Simulation Results and Performance Analysis

To evaluate the performance of the proposed model, a series of discrete-event simulations were conducted involving approximately 10,000 users with diverse abilities, preferences, and behavioral patterns. The agentic AI system was compared against two baseline models: (1) a Standard Personalized AI,

So, the higher the EthCost, the lower the final utility, meaning agent will avoid it.

which focuses solely on the personalization component  $\alpha \cdot Per(a, s)$  without considering accessibility or ethical cost, and (2) a Static Accessibility AI, which consistently selects the action that is most universally accessible, even if it reduces engagement or user satisfaction.

The primary key performance indicator (KPI) used in the evaluation is the Engagement Score, which represents a combination of how actively users interact with the system and how comfortable they feel during the interaction. However, achieving an ideal balance between these two aspects is challenging improvements in one dimension sometimes lead to a decrease in the other.

The equation used for comparison is as follows:

$$E_{score} = \alpha \cdot Per(a, s) + \beta \cdot Acc(a, s)$$

Sometimes we subtract a small amount for ethical cost or noise, though the exact effect is not always clear.

The simulation shows that the agentic model does better overall but it can be a bit jumpy depending on who the user is and what they're doing at the time. The static model is safe but kind of dull, and the personalized one takes too many risks. So, we can say the agentic AI feels more human, not perfect but better at keeping a balance between the two.

#### 7.1. Comparative Result Table

The table below shows what happened when we tested different AI systems with different groups of users. The scores aren't exactly the same each time because we added a bit of noise, and people don't always act the same way in every test.

**Table-2:** Comparative Average Engagement Score by User Group.

User Group	Agentic AI	Standard Personalization AI	Static Accessibility AI
Sighted Users	88	85	65
Visually Impaired	87	45	82
Motor Impaired	85	70	80
Expert User	90	92	60
Novice User	86	75	88

You can see that the agentic model stays pretty steady overall. It doesn't always get the highest scores, but it also doesn't crash. The personalization model keeps expert users happy, but it doesn't work well for people who need more accessibility

support. The static model helps those users more, but it can make expert users bored. The equation we used for the mixed score is kind of odd, but it's close enough to show the pattern we were looking for.

$$G_{score}(u) = \alpha_u \times Per(a, s) + \beta_u \times Acc(a, s) - EthCost(a, s)$$

where  $\alpha_u$  and  $\beta_u$  change for each user type. not very clean math but work enough for simulation test.

### 8. Ethical Implications and Societal Impact

Using these automated systems raises many important ethical questions, such as how much control we should give to machines and how much should stay with humans. The model itself carries ethical costs, but turning those ideas into real systems is not simple. Sometimes the agent becomes too smart for its own good and starts acting as if it knows better than the user, almost like a well-meaning parent. It might override what the user wants because it believes it is making the smarter choice. We must

ensure that users can stop the system at any time, because without that option, real autonomy is lost.

There is also the problem of data. The system records clicks, blinks, sounds, and almost everything else. Without proper care, privacy can disappear very quickly. We need clear and transparent rules, and systems should only collect the data that is truly necessary. On the social side, inequality becomes another concern. Some people have access to advanced versions of these systems, while others do not. This creates a divide where the rich benefit from smart AI assistance, while the poor

are left out. Policymakers need to step in and set boundaries so that this technology serves everyone, not just a privileged few.

### 8.1. Privacy and Data Chaos

The state model collects far too much data. Who gets to see it, and who keeps it safe? No one really seems to know. Data can leak, be used for training, or even disappear without anyone noticing. Many users have no idea what information is being collected about them in the first place. We need strong rules to limit how much data is gathered, but the real question is who will ensure that those rules are actually followed.

### 8.2. Human Control vs Machine Ego

When AI starts making more decisions, people begin to feel less in control. The system says things like “I know what you like,” even when that is not true. Over time, this makes trust weaker and more confusing. People need to always have a way to override the system and take back control whenever they want.

### 8.3. Tech Divide and Social Mess

People with money can buy the full version of the agent, while everyone else is left with the slower one. This creates a digital wall that divides society all over again. Technology that claims to be made for inclusion can end up excluding people instead.

## 9. Conclusion and Future Directions

This paper presents a formal mathematical model for an agentic AI system. It is not perfect, but it aims to make customer engagement more inclusive. We discuss, how the agents optimize utility functions that combine personalization, accessibility, and an ethical cost term called EthCost. The idea is significant because most conventional AI systems do not explicitly consider fairness in their design. The simulation results are also encouraging, showing that the system engages better with different types of users, not just one group.

For future work, there are a few areas to focus on. First, we need better ways to understand the user’s state  $S_t$  in real time without collecting excessive data, as that creates privacy concerns. Second, we could explore a multi-agent approach, where several AI systems work together to assist a single person—like a team of bots. Finally, it is essential to test the system in the real world with actual users, perhaps through A/B testing, to verify whether it performs as expected. Building truly inclusive AI system will not be easy, but it appears promising if we continue using model-driven thinking and keep improving the parts that do not yet work well.

### Acknowledgements

The authors acknowledge the prior scholarly work in agentic artificial intelligence, human–computer interaction, accessibility design, and ethical AI that informed the conceptual foundations of this study. The author also thanks the anonymous reviewers and editors for their constructive feedback, which helped improve the clarity and rigor of the manuscript.

### Author Contributions

The author solely conceived the research idea, developed the theoretical framework and mathematical model, conducted the simulations and analysis, interpreted the results, and wrote the manuscript. The author reviewed and approved the final version of the paper.

### Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

### References

1. Rao, A. S., & Georgeff, M. P. (1995). BDI agents: From theory to practice. \*Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), 312–319.
2. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)*, 52(1), 1–38.
3. Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., Elmqvist, N., & Diakopoulos, N. (2016). *Designing the user interface: strategies for effective human-computer interaction* (6th ed.). Pearson.
4. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2).
5. Gajos, K. Z., Weld, D. S., & Wobbrock, J. O. (2010). Automatically generating personalized user interfaces with Supple. *Artificial Intelligence*, 174(12–13), 910–950.
6. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*, 29.
7. Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., ... & Teevan, J. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13.
8. Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency*, 149–159.
9. Borning, A., & Muller, M. (2012). Next steps for value sensitive design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1125–1134.
10. Crawford, K., & Calo, R. (2016). There is a blind spot in AI research. *Nature*, 538(7625), 311–313.
11. Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114.
12. Friedman, B., & Kahn Jr, P. H. (2007). Human values, ethics, and design. *The human-computer interaction handbook*, 1223–1248.
13. Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G. Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics*, 4(37), eaay7120.
14. Kocielnik, R., Avrahami, D., Marquez, J., Romero, O., Diriye, A., & Azenkot, S. (2020). Can You Help Me? A Study of Human-AI Interaction for Inclusive Task Assistance. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1–25.
15. Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P. (2018). Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4), 611–627.
16. Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38.

17. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Books.
18. Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253.
19. Russell, S., Dewey, D., & Tegmark, M. (2015). Research priorities for robust and beneficial artificial intelligence. *AI Magazine*, 36(4), 105-114.
20. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59-68.
21. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
22. Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751-752.
23. Veale, M., & Binns, R. (2017). Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2).
24. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2), 76-99.
25. Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Public Affairs.