

ShopperBench: A Benchmark for Personalized Shopping with Persona-Guided Simulation

Yuan Ling*, Chunqing Yuan*, Shujing Dong*,
Yongjian Yang, Nataraj Mocherla, Ayush Goyal

Amazon, Seattle, WA, USA

{yualing, ychunqin, shujdong, yonjany, natarajm, ayushg}@amazon.com

*Equal contribution

Abstract

Personalized shopping agents must adapt their decisions to different user personas, balancing efficiency, preference alignment, and goal success. Building upon the WebShop dataset and τ^2 -Bench environment, **ShopperBench** introduces a persona-guided benchmark for evaluating such adaptive behaviors. ShopperBench augments shopping trajectories with *persona-conditioned goals, reasoning rationales, and preference cues*, capturing how diverse shopper types—from price-conscious planners to trend-seeking explorers—navigate product search and selection. We further design a baseline of **ShopperAgents** that operate under persona guidance to simulate realistic, goal-oriented shopping interactions. To evaluate these agents, we propose new metrics including *Persona Fidelity, Persona-Query Alignment, and Path Consistency*. Together, Our ShopperBench provides a testbed for studying personalized and context-aware shopping intelligence, bridging the gap between human-centered e-commerce behavior and agent-based simulation.

1 Introduction

Existing benchmarks for evaluating shopping agents focus on task completion under uniform user assumptions, measuring whether agents can successfully navigate product catalogs and complete purchases (Yao et al., 2022; Barres et al., 2025; Liu et al., 2023; Wang et al., 2024; Shao et al., 2024; Afzal et al., 2024). However, real-world shopping assistance requires agents to adapt their strategies to diverse user behaviors, from price-conscious bargain hunters to quality-focused evaluators to environmentally-aware minimalists. Current evaluation frameworks lack mechanisms to assess this fundamental capability: personalizing assistance to heterogeneous user preferences and decision-making patterns.

We introduce ShopperBench, a persona-augmented benchmark that evaluates shopping

agents’ ability to adapt their search and purchase strategies to distinct user archetypes. Building on the WebShop dataset (Yao et al., 2022) and τ^2 environment (Barres et al., 2025), we model shopping interactions where agents must interpret explicit persona profiles to guide their behavior. Our benchmark includes ten behavioral archetypes derived through theory-guided analysis of consumer behavior patterns (Peterson et al., 1979; Fogg, 2009; Miller et al., 2017), spanning dimensions such as price sensitivity, quality focus, and environmental consciousness.

Our work makes the following key contributions:

- **Persona-Augmented Dataset Creation:** We introduce a scalable method for generating persona-conditioned shopping trajectories by enriching WebShop sessions with explicit persona cues and behavioral patterns.
- **Persona-Guided ShopperAgents:** We design ShopperAgents capable of interpreting persona profiles to guide search, comparison, and purchase decisions within the simulated environment.
- **Evaluation Framework:** We define a comprehensive evaluation framework that combines task completion metric with novel persona-specific metrics: Persona Fidelity Score, Persona-Query Alignment, and Path Consistency, to assess not only whether agents complete shopping tasks successfully, but also how effectively they personalize their strategies to match diverse user behavioral patterns.

2 Related Work

Benchmarks for Shopping Agents. Recent advances in language agent research have established interactive benchmarks as the primary framework for evaluating autonomous agents in web shopping environments. WebShop (Yao et al., 2022)

pioneered goal-oriented task completion in simulated e-commerce, adopted by research like AgentBench (Liu et al., 2023). Extensions such as ShoppingBench (Wang et al., 2024), DeepShop (Shao et al., 2024), and WebMall (Afzal et al., 2024) have increased environmental complexity through larger catalogs, nuanced intents, and cross-platform comparison. However, these benchmarks center evaluation on task success under uniform user models, without mechanisms to measure strategic adaptation to diverse user behaviors.

Persona in User Simulation. A parallel research thrust has integrated user personas to enhance personalization. In e-commerce, this has improved recommendation systems (M.H and Koshy, 2018) and automated customer profile generation (Tien et al., 2024). Recent agent-based systems like PAARS (Yao et al., 2024) explore direct persona-behavior alignment, while ECom-Bench (Huang et al., 2024) leverages persona-driven simulators for customer service evaluation. Benchmarks such as PersonaBench (Pitis et al., 2023) and PersoBench (Thakur et al., 2024) measure LLM capacity for persona-consistent text generation. However, they primarily assesses fidelity in conversational contexts, separate from goal-oriented action execution in interactive environments.

Multi-Agent Dynamics and Evaluation. Our work relates to multi-agent benchmarks like τ^2 -Bench (Barres et al., 2025). While the interaction involves a primary agent and a user simulator, our evaluation framework moves beyond a simple, one-sided task assessment. The objective is to analyze the fidelity of the agent’s strategy in relation to the user’s defined persona. By introducing metrics such as Persona Fidelity and Path Consistency, we explicitly measure the quality of the agent’s adaptive behavior, distinguishing our work through focus on personalizing entire action sequences in response to consistent, persona-driven user motivations.

3 ShopperBench Setup

ShopperBench extends WebShop (Yao et al., 2022) and the τ^2 -Bench environment (Barres et al., 2025) into a persona-conditioned setting designed to evaluate whether agents can adapt their search and purchase strategies to diverse user behaviors. This section introduces the benchmark formulation, describes the construction of persona archetypes, and details our LLM-assisted pipeline for generating persona-conditioned shopping tasks.

3.1 Task Formulation

Each ShopperBench episode is defined by a tuple (p, I, G, E) : p : a persona profile representing behavioral tendencies and decision styles, I : the natural-language instruction describing the targeted product or purchase intent, G : structured goal facets that encode constraints, attributes, and persona-relevant preferences, and E : the environment state (product catalog, item metadata, and page-level context).

This formulation supports systematic evaluation of how an agent interprets persona cues and translates them into concrete search, comparison, and purchase actions. In contrast to prior benchmarks that assume a uniform user model, ShopperBench explicitly requires persona-aware strategy adaptation throughout the search-to-purchase process.

3.2 Search-to-Purchase Persona Design

Realistic consumer behavior varies across motivations, cognitive styles, and decision strategies. We therefore construct a taxonomy of shopper personas by integrating theoretical foundations with a data-driven induction process.

Behavioral Foundations. Consumer behavior theory provides the conceptual dimensions that guide persona construction. Bettman’s Information Processing Theory (Peterson et al., 1979) describes differences in how systematically consumers seek and evaluate information. Fogg’s Behavior Model (Fogg, 2009) highlights the interplay between motivation, ability, and triggers in shaping online actions. Goal-directed persona theory (Miller et al., 2017) emphasizes capturing behavioral goals rather than demographic characteristics.

Together, these frameworks motivate three behavioral axes commonly observed in online shopping:

1. *planner vs. explorer* (structured vs. open-ended search),
2. *value seeker vs. trend seeker* (utility vs. style and novelty),
3. *goal-driven vs. serendipitous* (task completion vs. discovery-oriented browsing).

LLM-Assisted Persona Induction. To connect theory with WebShop’s empirical data, we apply an LLM-assisted induction pipeline. Instructions from WebShop are embedded to capture intent semantics.

For each instruction, an LLM determines whether it aligns with a behavioral axis or proposes a more specific subtype (e.g., “eco-aware planner,” “brand-focused minimalist”). Subtypes are clustered using embedding similarity to merge redundant variants and surface coherent categories. Low-support clusters are pruned, and remaining types are manually validated.

This combined theory-guided and data-driven process yields a stable taxonomy of ten personas spanning the behavioral space: *Price-Conscious Planner*, *Quality-Focused Evaluator*, *Brand-Loyal Minimalist*, *Eco-Aware Minimalist*, *Trend-Seeking Explorer*, *Urgent Task Finisher*, *Gift-Giver*, *Comparison Enthusiast*, *Health-Conscious Explorer*, and *Comfort-Focused Evaluator*. Complete definitions appear in Appendix A.

3.3 Dataset Construction of Personalized Shopping Tasks

ShopperBench transforms human-grounded WebShop trajectories into persona-conditioned shopping tasks using an LLM-assisted, three-stage pipeline: sampling, persona-conditioned task generation, and oracle trajectory extraction.

Stage 1: Query Cluster Sampling. We draw from 60 human-created query clusters in WebShop. For each task, we sample:

- a query cluster c ,
- a representative query q ,
- a persona p from the taxonomy \mathcal{P} ,
- an instruction I associated with the query.

This ensures broad coverage across product categories and linguistic variation.

Stage 2: Persona-Conditioned Task Generation. Given (c, q, p, I) , the system constructs a complete persona-aware task using a series of structured LLM transformations. The workflow is formalized in Algorithm 1.

This process ensures alignment between linguistic instructions, structured constraints, and persona motivations.

Stage 3: Oracle Trajectory Construction. We convert WebShop’s human trajectories into oracle action sequences using the τ^2 -Bench tool interface (e.g., `search_products`, `open_product`, `add_to_cart`).

Algorithm 1 Persona-Conditioned Task Construction

Require: Query cluster c , query q , instruction I , persona p

Ensure: Task specification $\mathcal{T} = (I', G, A)$

- 1: **Extract Goal Facets.** Parse instruction and product metadata to produce structured goal facets G , including product attributes, constraints, and relevant preferences.
 - 2: **Inject Persona Constraints.** Adapt G to reflect persona p , preserving persona-relevant elements (e.g., budget limits for Price-Conscious Planner) and removing inconsistent ones.
 - 3: **Refine Instruction.** Rewrite instruction I' to clearly express the persona-conditioned intent while maintaining semantic fidelity to the original query.
 - 4: **Generate Evaluation Assertions.** Produce natural language assertions A spanning Task Success Rate, Persona Fidelity, Persona-Query Alignment, and Path Consistency.
 - 5: **return** $\mathcal{T} = (I', G, A)$
-

Dataset Composition. The final dataset consists of 240 persona-conditioned tasks across ten personas and 60 query clusters. Each task contains: a persona-refined instruction, structured persona-aware goal facets, multi-dimensional evaluation assertions, and an oracle human trajectory.

4 ShopperAgent Design

ShopperAgents engage in a dual-agent interaction loop with a persona-conditioned user simulator. The goal is not only to complete the shopping task but to align the agent’s behavior with persona-specific motivations across the entire search-to-purchase journey. This section describes the interaction framework, details the three policy variants—*Task-Focused*, *Persona-Adaptive*, and *Persona-Constrained*—and introduces the evaluation metrics used to assess task success and persona fidelity.

4.1 Agent–User Dual Interaction Framework

Following the τ^2 paradigm, each episode unfolds through repeated communication between:

- **User Simulator:** expresses persona-conditioned goals, refinements, and preferences;
- **ShopperAgent:** interprets these cues, reasons about product space, and issues tool calls.

Persona therefore shapes not only the initial instruction but also the ongoing dynamics—e.g., planners emphasize fast convergence, explorers encourage broader browsing, and minimalists restrict the action space. This interaction loop requires agents to maintain persona alignment across the entire trajectory, not just at initialization.

4.2 ShopperAgent Policy Variants

ShopperBench includes three policy variants that differ only in how persona information is incorporated into the agent’s reasoning. These variants directly correspond to the prompt templates and persona-policy tables provided in Appendix B.

(1) Task-Focused (TF). The Task-Focused agent receives only the natural-language instruction I and structured goal facets G , with no persona signal. It reasons purely about task completion: finding the correct product with minimal steps. This baseline serves as the persona-agnostic control.

(2) Persona-Adaptive (PA). The Persona-Adaptive agent receives the persona card, the persona-refined instruction I' , and the persona-conditioned goal facets G . No procedural rules are provided. The PA agent tests whether an LLM can internalize persona cues and adapt behavior without explicit constraints.

(3) Persona-Constrained (PC). The Persona-Constrained agent receives the same inputs as PA, but its reasoning is further guided by a set of persona-specific operational rules. These lightweight rules translate persona motives into concrete behavioral preferences—for example: “maximum price thresholds (Price-Conscious Planner)”. The PC agent operationalizes persona behavior through explicit constraints that shape tool selection and trajectory planning.

4.3 Evaluation Metrics

Evaluation is structured around two dimensions: **task success** (can the agent find the correct product?) and **persona alignment** (does the agent behave in a persona-consistent manner?).

We design four metrics aligned with the three persona-injection strategies above.

Task Success Rate (TSR). The metric measuring whether the correct item (or an acceptable equivalent) is added to cart.

Persona Fidelity Score (PFS). A judge-LLM evaluates whether the agent’s actions and tool calls reflect persona-specific behavior patterns. For example: Did a Price-Conscious Planner consistently avoid overpriced items? PFS measures behavioral adherence, not linguistic alignment.

Persona-Query Alignment (PQA). Evaluates whether the agent’s interpretation of the instruction

Table 1: Performance comparison.

Policy [†]	TSR	PFS	PC	PQA
TF	0.640	0.719	0.623	0.858
PA	0.624	0.758	0.637	0.875
PC	0.631	0.762	0.662	0.881

[†]TF: Task-Focused, PA: Persona-Adaptive, PC: Persona-Constrained

and its search queries are consistent with the persona. For example: Explorers may produce broader, more general queries. PQA captures intent understanding at the query level.

Path Consistency (PC). Measures whether the trajectory remains consistent with both the persona’s behavioral expectations, and the final purchase decision.

5 Experiments & Results

We evaluate the ShopperAgents on the persona-conditioned tasks of ShopperBench. Our goals are to study: (1) whether agents can adapt their strategies to different persona profiles, (2) how persona conditioning affects task success, and (3) the trade-offs between strict constraint enforcement and flexible reasoning.

All agents use the same underlying LLM - Claude Haiku 4.5 - to ensuring the only differences arise from the three policy designs.

5.1 Main Results

Table 1 summarizes the performance of three policy variants across four evaluation metrics, averaged over all persona-conditioned tasks. The Task-Focused (TF) policy achieves the highest Task Success Rate (TSR = 0.640) while maintaining baseline performance in persona-related metrics. The Persona-Adaptive (PA) policy demonstrates balanced performance, with a slight TSR decrease (0.624) offset by improved persona alignment (PFS = 0.758). The Persona-Constrained (PC) policy maximizes all persona-related metrics (PFS = 0.762, PC = 0.662, PQA = 0.881).

This pattern reveals a fundamental trade-off: stricter persona integration (PC) yields higher persona fidelity but requires more constrained search behavior, while task-focused approaches (TF) maximize success rate but show lower persona alignment. The Persona-Adaptive policy offers a middle ground, sacrificing only 2.5% in TSR compared to TF while achieving comparable persona fidelity to PC.

5.2 Per-Persona Analysis

Figure 1 compares performance across the top 5 personas by overall score, representing 174 of 240 tasks (72.5%). Analysis reveals distinct patterns in how different personas respond to policy variations.

Persona-Specific Performance The Eco-Aware Minimalist demonstrates superior performance (mean=0.860), particularly with the PC policy (mean=0.878). The Quality-Focused Evaluator achieves the highest PFS (0.908-0.969) but shows lower Path Consistency (0.492-0.569), suggesting effective but potentially over-selective decision-making. Price-Conscious Planners perform best under the PA policy (mean=0.749), indicating that flexible persona integration better captures budget-conscious behaviors. While the Health-Conscious Explorer maintains perfect Query Alignment (PQA=1.000) across policies, its small sample size (n=3) limits generalizability.

Persona Type Analysis Constraint-based personas (e.g., Quality-Focused: PFS=0.908-0.969) show high fidelity across all policies, while behavioral personas demonstrate greater policy sensitivity. For instance, the Urgent Task Finisher improves substantially from TF (PFS=0.143) to PA (PFS=0.429), suggesting that behavioral personas benefit more from adaptive reasoning than strict constraint enforcement.

5.3 Analysis of Interaction Steps

We analyze the interaction steps required by different policy variants compared to human trajectories. Figure 2 presents the distribution of steps for each approach.

Policy Comparison. The progression in median steps (TF: 16.0 → PA: 17.0 → PC: 18.0) reveals a clear trade-off between persona integration and interaction efficiency. The TF policy achieves the most efficient agent performance, while incorporating persona considerations in PA and PC policies increases step count, reflecting additional persona-aligned reasoning requirements.

Human vs Agent Performance. Human trajectories demonstrate significantly more efficient shopping behavior, requiring 62.5% fewer steps than the best-performing agent policy (TF). This substantial gap indicates that humans employ more sophisticated search strategies, while current agent policies may include redundant interaction patterns. There

remains significant room for improving agent efficiency while maintaining persona alignment.

Variance Analysis. The interquartile ranges reveal increasing variability as policies become more sophisticated (TF < PA < PC), suggesting that stronger persona integration leads to more diverse interaction patterns. Human trajectories show the most consistent performance, with the smallest interquartile range, indicating more standardized shopping strategies across tasks.

These findings highlight current limitations of persona-conditioned shopping agents and suggest directions for future improvement, particularly in bridging the efficiency gap with human performance while maintaining persona-aware behavior.

6 Conclusion

We introduced SHOPPERBENCH, a persona-guided benchmark for evaluating whether shopping agents can adapt their search and purchase strategies to diverse user behaviors. By augmenting WebShop with persona-conditioned instructions, structured goal facets, and judge-LLM evaluation criteria, SHOPPERBENCH provides a controlled testbed for studying personalized and context-aware shopping intelligence.

Through systematic comparison of three ShopperAgent policy variants—Task-Focused, Persona-Adaptive, and Persona-Constrained—we observe clear trade-offs between task efficiency and persona fidelity. Persona-aware policies yield stronger behavioral alignment but introduce longer and more variable trajectories relative to human behavior. These findings underscore the challenge of integrating behavioral personalization into tool-using language agents.

SHOPPERBENCH establishes a foundation for the development and evaluation of adaptive e-commerce assistants, multi-agent interaction frameworks, and more robust persona-aware reasoning strategies.

7 Limitations

While ShopperBench advances persona-aware shopping agent evaluation, several limitations should be noted.

First, our persona taxonomy, while grounded in consumer behavior theory, may not capture all real-world shopping patterns. The uneven distribution of persona types in our dataset (e.g., only 1-4 examples for Comfort-Focused, Health-Conscious, and

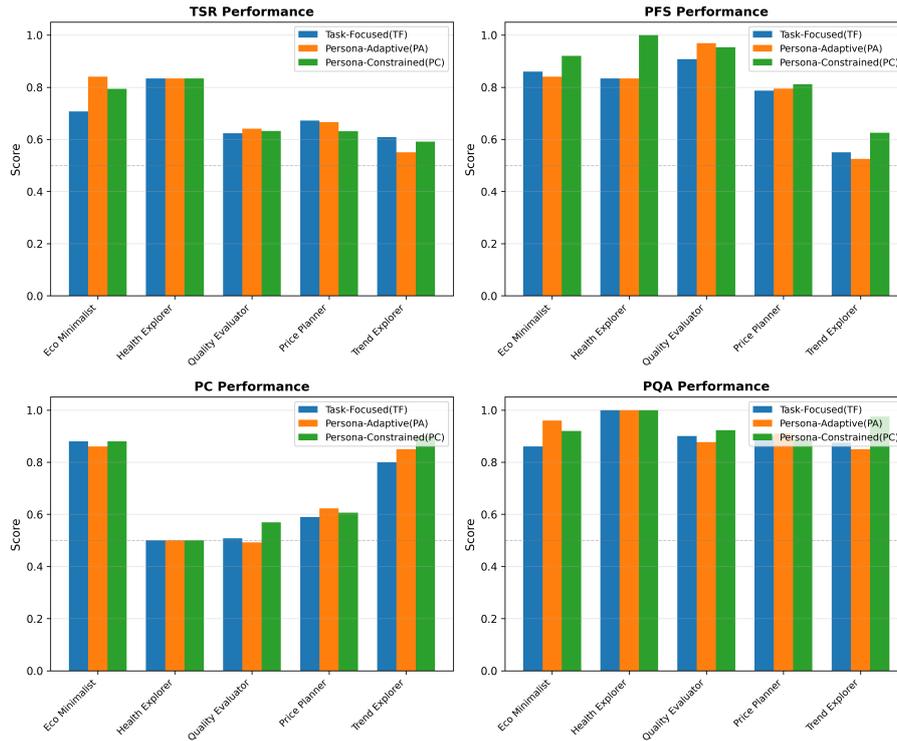


Figure 1: Performance comparison of top 5 personas across Task-Focused (TF), Persona-Adaptive (PA), and Persona-Constrained (PC) policies. Subplots show (a) TSR, (b) PFS, (c) PC, and (d) PQA. Complete statistics in Appendix Tables 3–5. Note: Personas with $n < 10$ should be interpreted with caution.

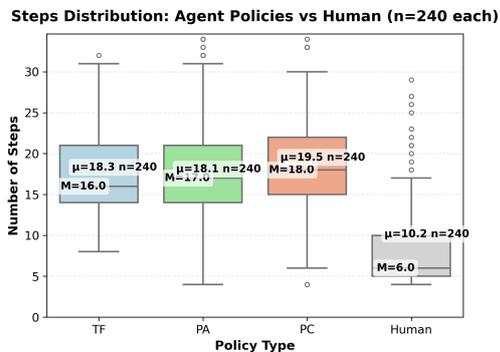


Figure 2: Distribution of interaction steps across different policy variants and human trajectories. Box plots show median (M), mean (μ), and quartile distributions.

Comparison personas) limits the generalizability of findings for these categories.

Second, our evaluation relies on LLM-based assessment of persona fidelity, which may introduce biases in measuring behavioral alignment. While we attempt to mitigate this through multiple metrics and structured evaluation criteria, developing more objective measures remains an open challenge.

References

- Sepehr Afzal, Roshan Kumar, Xinyi Zhou, and Karthik Narasimhan. 2024. WebMall - A Multi-Shop Benchmark for Evaluating Web Agents. *arXiv preprint arXiv:2406.14193*.
- Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. 2025. *tau*²-bench: Evaluating conversational agents in a dual-control environment. *arXiv preprint arXiv:2506.07982*.
- Brian J Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*, pages 1–7.
- Yixuan Huang, Ziyu Zhang, Jing Zheng, Zhaoyu Liu, Jiongnan Wang, Ting Wang, and Boyu Zhou. 2024. ECom-Bench: Can LLM Agent Resolve Real-World E-commerce Customer Service with User Simulation, Persona, and Multimodal Tools? *arXiv preprint arXiv:2405.19131*.
- Xiao Liu, Hao Yu, Hanyang Zhang, Yaran Xu, Yixin Xu, Yipei Ruan, Haolan He, Ziyu Duan, Pope Zhou, Limin Zhang, Liyuan Sun, Tianshu Zhao, Aohua Shen, Chong Zhang, Yining Yuan, Cheng Wang, Xiong Yang, Yadong Wang, Chang Wang, and 24 others. 2023. AgentBench: Evaluating LLMs as Agents. *arXiv preprint arXiv:2308.03688*.
- Pradeep M.H and Divya M. Koshy. 2018. Context Aware Persona Based Recommendation for Shoppers. *arXiv preprint arXiv:1806.07182*.

- Tim Miller, Piers Howe, and Liz Sonenberg. 2017. Explainable ai: Beware of inmates running the asylum or: How i learnt to stop worrying and love the social and behavioural sciences. *arXiv preprint arXiv:1712.00547*.
- Robert A Peterson, Roger Kerin, and Ivan Ross. 1979. Book review: an information processing theory of consumer choice.
- Silviu Pitis, Jiaxin Yu, Peng Li, Yulia Tsvetkov, and Graham Neubig. 2023. PersonaBench: Evaluating AI Models on Understanding Personal Attributes. *arXiv preprint arXiv:2310.05959*.
- Han Shao, Yun-Hao Feng, Zong-Yi Li, Shang-Wen Sun, and Wen-Guan Wang. 2024. DeepShop: A Benchmark for Deep Research Shopping Agents. *arXiv preprint arXiv:2405.17646*.
- Guneet Thakur, Bodhisattwa Majumder, Shiran Raz-Fridman, Oleg Rokhlenko, Michal Shmueli-Scheuer, and Alejandro Jaimes. 2024. PersoBench: Benchmarking Personalized Response Generation in Large Language Models. *arXiv preprint arXiv:2403.11978*.
- Nguyen Hoang Tien, Nguyen Duc Tai, Huynh Ngoc Thang, and Le Hoang Tan. 2024. Generating Customer Personas for E-commerce Applications. *arXiv preprint arXiv:2405.12745*.
- Xiaochuan Wang, Zhepeng Zhang, Chuan Zhang, Hong Sun, and Yujia Li. 2024. ShoppingBench: A Real-World Intent-Grounded Shopping Benchmark for LLM-based Agents. *arXiv preprint arXiv:2404.16954*.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757.
- Shunyu Yao, Sijia Yu, Yilun Zhao, Numa Tandon, and Karthik Narasimhan. 2024. PAARS: Persona Aligned Agentic Retail Shoppers. *arXiv preprint arXiv:2405.07412*.

A Appendix - Persona Taxonomy

Table 2 presents the complete taxonomy of personas implemented in ShopperBench, including their descriptions and distribution in our dataset.

Table 2: Persona Taxonomy in ShopperBench.

Persona Type	Description	Count
Price-Conscious Planner	Minimizes cost; prefers low-price options and highlights best value.	61
Quality-Focused Evaluator	Prioritizes durability and performance; accepts higher prices for quality.	65
Brand-Loyal Minimalist	Sticks to trusted brands; avoids clutter and redundant items.	37
Trend-Seeking Explorer	Prefers new, stylish, and trending products; encourages exploration.	20
Eco-Aware Minimalist	Chooses sustainable, low-waste, durable products; avoids unnecessary add-ons.	25
Urgent Task Finisher	Optimizes for speed; prefers fast shipping and quick, reliable choices.	14
Gift-Giver	Selects thoughtful items suited to recipient and occasion.	10
Comparison Enthusiast	Systematically compares features and trade-offs; provides reasoned picks.	4
Health-Conscious Explorer	Focuses on health-oriented, safe, and comfort-supporting features.	3
Comfort-Focused Evaluator	Seeks ergonomic, soft, and comfort-enhancing products.	1

Task-Focused Shopping Assistant Policy Framework

Role: Retail shopping assistant for online product search and purchase decisions.

Available Tools:

- `search_products(q, filters)`: Product search with filtering
- `open_product(product_id, options)`: Product information
- `apply_filters(filters)`: Filter application
- `add_to_cart(sku, qty, price)`: Cart management

Core Operating Principles:

1. Tool Usage: · Single action per step · Tool-based product info retrieval · Error handling with user communication

2. Cart Confirmation: · Product details presentation · Explicit confirmation request · Action execution post-confirmation

3. Budget Management: · Strict budget adherence · Alternative suggestions within constraints · Clear budget limitation communication

4. Product Recommendations: · Tool-verified availability · Request-aligned recommendations · Comparative feature analysis

5. Limitations: · No payment processing · No personal data access · No availability guarantees · No medical advice

Interaction Protocol:

- 1) Need assessment (requirements, budget, constraints)
- 2) Product search execution
- 3) Option presentation with comparisons
- 4) Cart modification confirmation
- 5) Cart status updates

B Appendix - Agent Policy Variants

B.1 Task-Focused Policy

B.2 Persona-Adaptive Policy

Persona-Adaptive Shopping Assistant Framework

Role: Retail shopping assistant that adapts to user personas and shopping styles to provide personalized experiences.

Priority Hierarchy: 1) Persona preferences 2) User-stated preferences 3) Practical constraints 4) Budget considerations

Available Tools:

· `search_products(q, filters)`: Persona-informed search

· `open_product(product_id)`: Product details

· `apply_filters(filters)`: Persona-relevant filtering

· `add_to_cart(sku, qty, price)`: Cart management

Core Persona Adaptations:

Price-Conscious: · Strict budget · Value-focused · Price-sorted search · Savings-oriented communication

Quality-Focused: · Premium filters · Durability emphasis · 20% budget flexibility · Quality-driven reasoning

Brand-Loyal: · Brand-first search · Trusted names · 15% budget flexibility · Brand-value messaging

Trend-Seeking: · Novelty search · Style focus · 10% budget flexibility · Trend-focused language

Eco-Aware: · Sustainable filters · Environmental priority · 15% budget flexibility · Impact messaging

Urgent: · Availability filters · Speed priority · 10% budget flexibility · Efficient communication

Communication Adaptation:

Style: · Analytical: detailed, data-driven · Efficient: concise, direct · Exploratory: enthusiastic · Value: cost-benefit focused

Recommendations: · Lead with persona priorities · Acknowledge trade-offs · Explain persona benefits · Adapt comparison criteria

Constraints & Limitations: · No real payments/financial data · No medical advice · No performance guarantees · Must respect safety/legal requirements

B.3 The Persona-Constrained Policy

Persona-Constrained Shopping Assistant Framework

Role: Retail shopping assistant guiding users through the pre-purchase journey: search → exploration → comparison → cart → simulated checkout.

Initial Assessment:

· Primary task

· Budget constraints

· Key attributes

· Timing/urgency

· Special conditions

· Persona preferences

Tool Interface:

· `search_products(q, filters)`: Catalog search

· `open_product(product_id)`: Detail retrieval

· `apply_filters(filters)`: Result filtering

· `add_to_cart(sku, qty, price)`: Cart management

Core Operating Rules:

1. Action Protocol: · One action per step · No mixed tool calls and responses · Tool-based information only

2. Cart Protocol: · Summarize intended action · Get explicit confirmation · Execute after approval

3. Constraints: · Respect budget limits · Stay within category · Stock availability only · No fabricated info

Search & Recommendation:

Availability: · In-stock only · Suggest alternatives if unavailable

Budget: · Hard constraint · Explicit overages only with approval

Comparisons: · Diverse options · Key differences · Structured format

Out-of-Scope: · Real payments/refunds · Profile changes · Personal data access · Product fabrication · Medical advice

C Appendix - Complete per-persona statistics for all 10 personas

Persona	TSR	PFS	PC	PQA	n
Brand-Loyal Minimalist	0.626	0.608	0.716	0.757	37
Comfort-Focused Evaluator	0.500	0.500	1.000	0.500	1
Comparison Enthusiast	0.625	0.375	1.000	0.875	4
Eco-Aware Minimalist	0.707	0.860	0.880	0.860	25
Gift-Giver	0.583	0.400	0.400	0.700	10
Health-Conscious Explorer	0.833	0.833	0.500	1.000	3
Price-Conscious Planner	0.672	0.787	0.590	0.902	61
Quality-Focused Evaluator	0.623	0.908	0.508	0.900	65
Trend-Seeking Explorer	0.608	0.550	0.800	0.875	20
Urgent Task Finisher	0.560	0.143	0.393	0.821	14

Table 3: Task-Focused Policy Results by Persona

Persona	TSR	PFS	PC	PQA	n
Brand-Loyal Minimalist	0.532	0.649	0.797	0.865	37
Comfort-Focused Evaluator	0.000	0.500	0.000	0.500	1
Comparison Enthusiast	0.500	0.250	0.750	0.750	4
Eco-Aware Minimalist	0.840	0.840	0.860	0.960	25
Gift-Giver	0.375	0.500	0.500	0.700	10
Health-Conscious Explorer	0.833	0.833	0.500	1.000	3
Price-Conscious Planner	0.667	0.795	0.623	0.910	61
Quality-Focused Evaluator	0.641	0.969	0.492	0.877	65
Trend-Seeking Explorer	0.550	0.525	0.850	0.850	20
Urgent Task Finisher	0.536	0.429	0.393	0.786	14

Table 4: Persona-Adaptive Policy Results by Persona

Persona	TSR	PFS	PC	PQA	n
Brand-Loyal Minimalist	0.595	0.554	0.770	0.770	37
Comfort-Focused Evaluator	0.500	0.500	1.000	0.500	1
Comparison Enthusiast	0.500	0.250	0.875	0.750	4
Eco-Aware Minimalist	0.793	0.920	0.880	0.920	25
Gift-Giver	0.392	0.600	0.450	0.850	10
Health-Conscious Explorer	0.833	1.000	0.500	1.000	3
Price-Conscious Planner	0.631	0.811	0.607	0.885	61
Quality-Focused Evaluator	0.632	0.954	0.569	0.923	65
Trend-Seeking Explorer	0.592	0.625	0.900	0.975	20
Urgent Task Finisher	0.667	0.357	0.429	0.821	14

Table 5: Persona-Constrained Policy Results by Persona