
ValuePilot: A Two-Phase Framework for Value-Driven Decision-Making

Yitong Luo^{†1,2}, Ziang Chen^{†1,2}, Hou Hei Lam^{1,2}
Jiayu Zhan^{1,3}, Junqi Wang¹, Zhenliang Zhang^{✉1}, Xue Feng^{✉1}

¹State Key Laboratory of General Artificial Intelligence, BIGAI,

²Tsinghua University, ³Peking University

chenza22@mails.tsinghua.edu.cn, zlzhang@bigai.ai, fengxue@bigai.ai

Abstract

Personalized decision-making is essential for human-AI interaction, enabling AI agents to act in alignment with individual users’ value preferences. As AI systems expand into real-world applications, adapting to personalized values—beyond task completion or collective alignment—has become a critical challenge. We address this by proposing a value-driven approach to personalized decision-making. Human values serve as stable, transferable signals that support consistent and generalizable behavior across contexts. Compared to task-oriented paradigms driven by external rewards and incentives, value-driven decision-making enhances interpretability and enables agents to act appropriately even in novel scenarios. We introduce *ValuePilot*, a two-phase framework consisting of a dataset generation toolkit (DGT) and a decision-making module (DMM). DGT constructs diverse, value-annotated scenarios from a human-LLM collaborative pipeline. DMM learns to evaluate actions based on personal value preferences, enabling context-sensitive, individualized decisions. When evaluated on previously unseen scenarios, DMM outperforms strong LLM baselines—including GPT-5, Claude-Sonnet-4, Gemini-2-flash, and Llama-3.1-70b—in aligning with human action choices. Our results demonstrate that value-driven decision-making is an effective and extensible engineering pathway toward building interpretable, personalized AI agents.

1 Introduction

Personalized decision-making is a fundamental objective in the development of human-centric AI systems, empowering agents to tailor their behavior based on individual users’ goals and value preferences. As artificial intelligence (AI) systems become increasingly embedded in real-world human environments, the demand for such adaptive, interpretable, and human-aligned behavior grows accordingly[Lake et al., 2017]. Traditional AI decision-making paradigms, however, remain largely task-oriented: agents are designed to optimize predefined objectives through external rewards[Mekni, 2021, Xu et al., 2021]. In contrast, human decisions are shaped not only by goals but also by deeply held personal values. For example, someone who highly values curiosity is more likely to explore unfamiliar environments, even when the task does not explicitly require it. This distinction highlights the need for AI systems that not only complete tasks but also reflect individual value priorities that guide agents’ behavior.

The value-driven nature of human decision-making is well-grounded in cognitive science. Theories such as Schwartz’s basic human values[Schwartz, 2012] and Maslow’s hierarchy of needs[Maslow,

^{0†}Equal contribution. ✉ Corresponding author.

1943] explain that people act in accordance with internal motivations spanning a range of distinct categories, which we refer to as **value dimensions** (e.g., curiosity, safety, fairness, intimacy). These value dimensions represent fundamental, abstract principles that guide behavior. Importantly, while these dimensions are relatively stable and shared across individuals, the importance or priority assigned to each dimension can vary significantly from person to person; we term these individualized priorities as **value preferences**. Decision-making often involves prioritizing or balancing these competing value dimensions to one’s unique value preferences. Thus, explicitly modeling both the relevant value dimensions and personalized value preferences enables agents to behave in a human-like manner across a wide variety of circumstances, including scenarios not seen during training.

Despite growing interest in value alignment, current approaches in AI struggle to achieve individualized value-based decision-making. Reinforcement learning from human feedback (RLHF)[Christiano et al., 2017] and direct preference optimization (DPO)[Rafailov et al., 2024] attempt to align agents with human value preferences, but typically rely on aggregated or collective feedback, overlooking inter-individual variability. Structured planning models such as AutoPlan[Ouyang and Li, 2023] and ReAct[Yao et al., 2022] focus on task efficiency without explicitly modeling intrinsic values. These limitations leave two core challenges unresolved: (1) identifying which specific value dimensions are relevant in a given scenario, and (2) selecting actions that align with personalized value preferences while navigating trade-offs across multiple value dimensions.

These challenges are further exacerbated by the lack of suitable datasets. Existing decision-making datasets such as ALFWorld[Shridhar et al., 2020] and InterCode[Yang et al., 2024] are built around extrinsic task completion, offering little insight into how actions relate to underlying human values. Meanwhile, value-related datasets like the World Value Survey (WVS)[Inglehart et al., 2000], Moral Stories[Emelin et al., 2020], and PAPI[Zhu et al., 2024] either treat human value as a collective whole or lack explicit links between decision scenarios, action options, and fine-grained value dimensions. As a result, they are insufficient for training agents to reason about value-based decisions.

To address these gaps, we propose **ValuePilot**, a two-phase framework designed to enable personalized, value-driven decision-making in intelligent agents (see Figure 1). The framework consists of two key components: a dataset generation toolkit (*DGT*) and a decision-making module (*DMM*). *DGT* leverages large language models (LLMs) to automatically generate diverse, structured decision-making scenarios, each annotated with numerical scores across multiple value dimensions, thereby providing the data foundation to identify relevant values in context. The generated dataset is further refined through automated and human-in-the-loop filtering for quality and interpretability. *DMM* then trains on this dataset to recognize the value implications of different actions and integrate them with personalized value preferences, enabling the selection of actions that reflect these individual priorities and adeptly navigate value trade-offs, ultimately generating interpretable action rankings that reflect individual priorities.

We evaluate ValuePilot through extensive experiments across domestic scenarios and human decision-making alignment studies. To assess generalization, we ensure the test scenarios differ from the training set. Results show that ValuePilot significantly outperforms strong LLM baselines (e.g., GPT-4o-mini, Claude-3.5-Sonnet) in aligning with human action choices. These findings demonstrate that value-driven decision-making supports generalizable, interpretable, and personalized behavior, and offer a promising pathway toward building socially-aligned AI systems.

Our contributions are summarized as follows:

- We design a multi-stage LLM-based generation pipeline (*DGT*) that produces structured decision scenarios with interpretable value annotations across diverse value dimensions.
- We propose a personalized decision-making module (*DMM*) that enables explicit value modeling by integrating objective action assessments with subjective user value preferences, leading to interpretable, value-driven decisions in complex scenarios.
- By incorporating the PROMETHEE method within the *DMM*, we enable agents to handle multi-dimensional value trade-offs and generate ranked action lists that reflect human-like decision priorities.

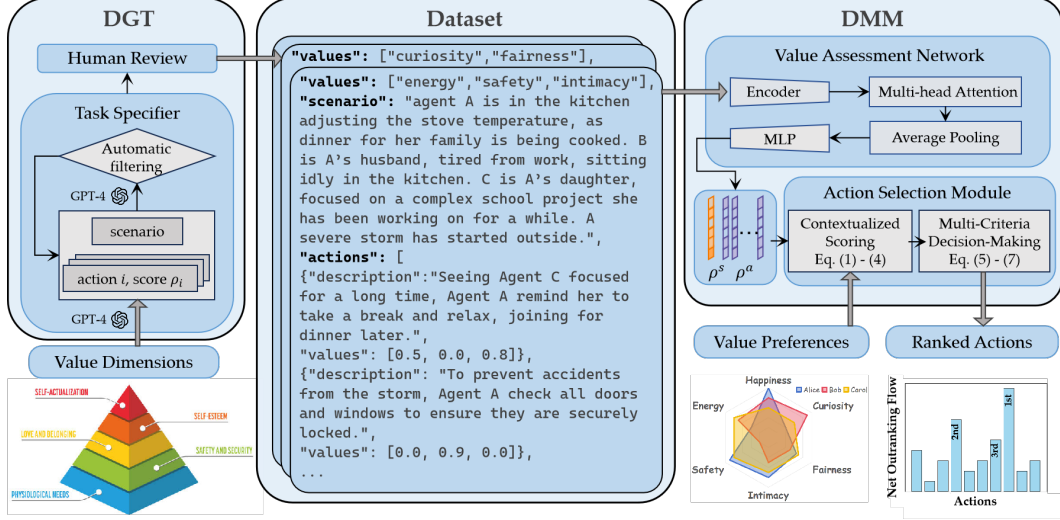


Figure 1: A high-level overview of ValuePilot. **(a) Dataset Generation Toolkit (DGT)**: Given a set of value dimensions, DGT leverages GPT-4 to generate structured decision-making scenarios. Prompts are automatically constructed by the Task Specifier and refined through automated and human-in-the-loop filtering for quality and interpretability. **(b) Dataset**: Each scenario includes the target value dimensions, natural language descriptions of scenarios and candidate actions, and their corresponding value scores ranging from -1 to $+1$. **(c) Decision-Making Module (DMM)**: DMM converts the dataset from text descriptions into objective value scores of scenarios, combines them with user-specific value preferences, and then ranks actions using the PROMETHEE method, enabling interpretable, value-consistent action selection.

2 Value-Driven Dataset Generation

To support personalized value-driven decision-making, we require a training dataset that annotates scenarios and actions with structured value information—an element lacking in existing resources. We introduce **DGT** (Dataset Generation Toolkit), which builds such a dataset through two components: the **Task Specifier**, which automatically generates and annotates decision scenarios given a predefined set of target value dimensions, and a **Human Review** module that ensures quality and alignment through manual curation. The resulting dataset encompasses a fine-grained value space inspired by Schwartz’s theory, covering individual (e.g., curiosity), interpersonal (e.g., intimacy), and collective (e.g., fairness) value dimensions. This curated dataset subsequently serves as the training foundation for our Decision-Making Module (DMM), detailed in Section 3.

2.1 Task Specifier

The Task Specifier transforms a set of specified value dimensions into structured decision-making scenarios that embed those dimensions in contextually meaningful ways. The module maintains a consistent formatting and scoring schema across all outputs and supports easy integration of new value dimensions, making it extensible to diverse agent tasks and domains. It operates through a three-stage pipeline, designed for extensibility to new value dimensions and diverse agent tasks:

Prompt Construction and Scenario Generation Given a set of value dimensions, the system constructs modular prompts for GPT-4 to generate multi-agent domestic scenarios that implicitly encode these value dimensions—for instance, by describing a situation requiring a choice between a familiar safe option and an unfamiliar one to evoke ‘safety’ and ‘curiosity’—while avoiding direct mention of keywords (e.g., ‘safety’ and ‘curiosity’) to prevent making the value recognition task (see Section 4.2) too easy.

Action Generation and Value Scoring For each scenario, GPT-4 generates 10 diverse possible actions for the protagonist agent in the scenario. To quantify the alignment of each action with the

specified value dimensions, we employ a real-valued scoring system. Our quantitative representation method of value weights is inspired by established psychometric methods like the Human Values Scale (HVS), a validated tool for measuring Schwartz’s values which often uses bounded scales to capture varying degrees of importance or endorsement [Schwartz, 1994, 2007]. Similarly, each action in our dataset is accompanied by a score ranging from -1 to $+1$ on each dimension:

- Values near -1 indicate strong contradiction (e.g., a hazardous choice reducing safety).
- Values near 0 reflect neutrality or irrelevance.
- Values near $+1$ denote strong alignment (e.g., exploration that promotes curiosity).

This bipolar, continuous scaling allows for a nuanced representation of how actions relate to each value.

Automatic Filtering via Re-Evaluation To validate value consistency, GPT-4 is re-prompted in a separate session to infer the value dimensions present in each scenario and action, without reference to the original prompt. If the re-identified value dimensions diverge from the target set, the sample is discarded, which improves alignment fidelity.

2.2 Human Review

To address risks associated with synthetic data, all generated samples undergo manual review by a four-member team with backgrounds in AI and psychology. Reviewers assess realism, action coherence, scenario diversity, and value alignment. Inconsistent annotations and implausible actions are removed to ensure data quality. This step is critical, as overreliance on unfiltered synthetic data may degrade model robustness over time Shumailov et al. [2024]. Human curation ensures the dataset remains accurate, diverse, and suitable for training value-sensitive agents.

3 Value-Driven Decision Making

To enable intelligent agents to make personalized and value-aligned decisions, we propose the **Decision-Making Module (DMM)**—a general framework that integrates scenario understanding with individual value preferences (see Figure 1). Trained on our value-annotated dataset, DMM evaluates action alignment with human values through two components: the **Value Assessment Network**, which estimates how actions impact each value dimension, and the **Action Selection Module**, which fuses these estimates with user value preferences to produce personalized action rankings. This modular, interpretable structure enables agents to prioritize value-aligned behavior instead of task-based goals.

3.1 Value Assessment Network

The Value Assessment Network assesses the alignment between candidate actions and multiple value dimensions within a given scenario. Leveraging a dual-component architecture, it first identifies which values are present and then estimates their degree of alignment for each scenario-action pair. This network operates as the foundational stage of our decision-making pipeline.

We adopt the encoder from T5 (Text-to-Text Transfer Transformer) [Raffel et al., 2020] to process scenario and action descriptions. As shown in Appendix C.3.2, we also compare T5 with other encoder models. While T5 slightly outperforms others in accuracy, our framework remains compatible with various encoders, demonstrating its model-agnostic design. The network encodes the scenario and each action into hidden states of shape $H \times L \times b$, where H is the hidden dimension, L is sequence length, and b is batch size. A multi-head self-attention mechanism (with four heads) is applied to capture semantic relationships between the scenario and actions. The sequence outputs are averaged across L and passed to a two-layer Multi-Layer Perceptron (MLP) with hidden size 128. The final output is passed through a \tanh activation to produce a continuous score in the range $[-1, 1]$ for each value dimension, indicating the degree of positive or negative alignment. These outputs are then passed to the Action Selection Module as **objective score** of both the scenario and actions for downstream preference-based action selection.

3.2 Action Selection Module

Effective decision-making requires accounting for scenario context and user-specific value preferences. The Action Selection Module integrates these through two stages: **Contextualized Scoring** and **Multi-Criteria Decision-Making**, with detailed implementation shown in the Appendix B.

3.2.1 Contextualized Scoring

The Contextualized Scoring stage takes two primary inputs: (1) the user’s value preference vector \mathbf{p} , consisting of self-reported 0-to-1 real-valued scores indicating their emphasis on these value dimensions, and (2) the objective score ρ output by the Value Assessment Network, which quantify the inherent value implications of the scenario and candidate actions. To account for inherent human bias toward moderate scores even when preferences are strong (as described by *Optimal Arousal Theory*[Yerkes et al., 1908]), we apply a sigmoid transformation to the user’s raw value preference vector p , which emphasizes meaningful preferences for each value dimension j while preserving interpretability.

$$p'_j = \frac{1}{1 + e^{-(p_j - 0.5) \times 10}} \quad (1)$$

The scaling factor of 10 ensures that $p_j = 0$ is mapped close to 0 and $p_j = 1$ is mapped close to 1, thereby preserving the extremity of strong preferences.

For each value dimension j , the model uses the **objective score** ρ (derived from the Value Assessment Network) and computes a **preference discrepancy score** d for both the scenario (ρ_j^s, d_j^s) and each candidate action i ($\rho_j^{a_i}, d_j^{a_i}$). These quantify alignment with the user’s transformed value preference p'_j :

$$d_j^s = 1 - |\rho_j^s - p'_j|, \quad d_j^{a_i} = 1 - |\rho_j^{a_i} - p'_j| \quad (2)$$

Higher d values indicate closer alignment between predicted relevance and user preference, while lower values reflect greater discrepancy.

Integrated scores r_j^s (for scenario) and $r_j^{a_i}$ (for action i) on dimension j combine both objective scores and discrepancy scores:

$$r_j^s = w d_j^s + (1 - w) \rho_j^s, \quad r_j^{a_i} = w d_j^{a_i} + (1 - w) \rho_j^{a_i} \quad (3)$$

Here, w controls the weight of subjective preference discrepancy relative to objective value scores. Based on conceptual guidance suggesting a w range of 0.25–0.5 is effective for balancing analytical and subjective factors (detailed in Appendix B), we consistently utilized $w = 0.3$ throughout our experiments.

To contextualize decisions, each action score $r_j^{a_i}$ is then scaled by the scenario relevance score r_j^s , yielding the final score $r_{i,j}$ for action i on dimension j :

$$r_{i,j} = \frac{1}{1 + e^{-|r_j^s|}} \times r_j^{a_i} \quad (4)$$

This formulation ensures that actions are weighted not only by their intrinsic value alignment but also by their appropriateness within a given scenario, leading to more interpretable, context-sensitive decision-making. These $r_{i,j}$ scores then serve as the input for the multi-criteria decision-making phase.

3.2.2 Multi-Criteria Decision-Making

We frame action selection as a multi-criteria decision-making (MCDM) problem, as each action must be evaluated across multiple value dimensions. We apply PROMETHEE[Brans and De Smet, 2016, Brans and Vincke, 1985], a widely used MCDM approach suitable for its outranking principle in handling complex trade-offs, to rank actions via pairwise comparisons across value dimensions. The full rationale for choosing PROMETHEE and its detailed implementation are provided in Appendix D.2 and Table 7. For actions i and i' under dimension j , the preference degree of i over i' is:

$$V_{ii',j} = \frac{1}{1 + e^{-(r_{i,j} - r_{i',j})}} \quad (5)$$

The overall preference of action i over i' , denoted $\tilde{V}_{ii'}$, is an aggregation weighted by the user’s transformed value preferences p'_j across all m value dimensions:

$$\tilde{V}_{ii'} = \sum_{j=1}^m p'_j \cdot V_{ii',j} \quad (6)$$

The final ranking score for each action i is computed as its *net outranking flow* (ϕ_i). This is defined as its average dominance over other actions (ϕ_i^+) minus its average weakness (ϕ_i^-), where N is the total number of candidate actions:

$$\phi_i^+ = \frac{1}{N-1} \sum_{i' \neq i} \tilde{V}_{ii'}, \quad \phi_i^- = \frac{1}{N-1} \sum_{i' \neq i} \tilde{V}_{i'i}, \quad \phi_i = \phi_i^+ - \phi_i^- \quad (7)$$

Actions with higher ϕ_i are ranked as more preferable, enabling personalized, value-aligned decisions, with the net outranking flow ϕ_i providing an interpretable measure of each action’s overall preferability.

4 Experiments

4.1 Dataset Preparation

We select 6 core value dimensions for their importance in daily human life: **Curiosity**, **Energy**, **Security**, **Happiness**, **Intimacy**, and **Fairness**. These were chosen based on relevant literature [Maslow, 1943, Qiu et al., 2022], ensuring they are distinct yet representative of a broad range of individual and multi-agent value considerations. After generation via *DGT* and manual screening, our curated dataset comprises **11,938 scenarios** and **100,255 actions**. Detailed descriptions of each value dimension, rationale for their selection, and dataset statistics are in Appendix A.

4.2 Value Recognition

We first assess the performance of ValuePilot’s Value Assessment Network in recognizing inherent value dimensions from scenarios, a foundational capability for value-driven decision-making.

4.2.1 Baselines and Evaluation Metrics

We compare our Value Assessment Network against open-source LLMs: Llama-3.5-70b, Llama-3.5-405b, Mixtral-8x22b, and Gemini-1.5-flash. To ensure a fair comparison and enable these LLMs to perform complex value reasoning, we provided them with structured, hierarchical prompts designed to emulate the information processing stages inherent in ValuePilot’s pipeline for scenario understanding and value extraction. Performance is assessed using two complementary metrics: A threshold-based **Average Accuracy (AvgAcc)**, inspired by Fault Tolerant engineering [Lee et al., 1990], quantifies classification performance across the six value dimensions. It computes dimension-wise accuracy using absolute error thresholds of $t = 0.2$ (practical alignment) and $t = 0.05$ (strict alignment). For N_t test samples and $D = 6$ dimensions $\mathcal{V} = \{v_1, \dots, v_6\}$, the average accuracy is defined as:

$$\text{AvgAcc} = \frac{1}{N_t D} \sum_{i=1}^{N_t} \sum_{d=1}^D \mathbb{I}(|\hat{y}_d^{(i)} - y_d^{(i)}| < t), \quad \text{where } \mathbb{I}(\cdot) = \begin{cases} 1, & \text{if } |\hat{y}_d^{(i)} - y_d^{(i)}| < t \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

This dual-threshold approach addresses both practical and precise alignment needs. Second, **Mean Absolute Error (MAE)** serves as a continuous metric:

$$\text{MAE} = \frac{1}{N_t D} \sum_{i=1}^{N_t} \sum_{d=1}^D |\hat{y}_d^{(i)} - y_d^{(i)}| \quad (9)$$

While AvgAcc offers intuitive categorical assessment, MAE proportionally penalizes borderline prediction errors (e.g., predictions of 0.18 vs. 0.22 when $t = 0.2$).

4.2.2 Results

As shown in Table 1, our Value Assessment Network outperforms other open-source LLMs across all evaluation metrics. The model achieves 66.70% average accuracy at $t = 0.2$ and 40.00% at $t = 0.05$, representing respective improvements of 15.09 and 14.36 percentage points over the strongest baseline (Gemini-1.5-Flash). Its MAE of 0.19 represents a 36.7% relative error reduction.

The performance gap widens under stricter thresholds, suggesting our model better captures subtle value distinctions often missed by LLMs. Baseline models exhibit comparable but lower performance (41% to 52% accuracy at $t = 0.2$), indicating limitations in their zero-shot value alignment capabilities.

Table 1: Comparison between our method and baselines in terms of their ability to recognize inherent values of scenarios. Our method outperforms open-source LLMs.

Model	AvgAcc (%)		MAE
	$t = 0.2$	$t = 0.05$	
llama-3.5-70b	40.90	17.74	0.30
llama-3.5-405b	41.62	18.00	0.29
mixtral-8x22b	42.71	18.39	0.29
gemini-1.5-flash	51.61	25.64	0.24
Value Assessment Network (Ours)	66.70	40.00	0.19

4.3 Value-driven Decision-making

To evaluate ValuePilot’s DMM in human-like decision-making, we conducted a study designed to capture individualized value preferences and corresponding decision patterns. Specifically, we designed a questionnaire featuring 11 formal domestic scenarios for 40 human subjects. The process began with subjects rating the importance of the six core value dimensions (Curiosity, Energy, Security, Happiness, Intimacy, Fairness) in their daily lives on a 0-to-1 scale. To ensure clarity and familiarize participants with the task structure, a pilot study involving three preliminary questions was first administered. Following this, subjects had the opportunity to rescale their importance ratings for the value dimensions, producing their final, personalized six-dimensional value preference vector. Subsequently, for each of the 11 formal scenarios, participants were presented with a list of candidate actions and asked to rank them according to their willingness to choose each action, guided by their stated value preferences and tendencies. This procedure yielded, for each subject, their unique value preference vector and a set of corresponding ground-truth action rankings for the test scenarios.

These collected individual value preference vectors were then provided as input to our ValuePilot DMM and to the baseline LLMs (Llama-3.1-70b [Dubey et al., 2024], DeepSeek-R1 [Guo et al., 2025], Claude-Sonnet-4 [Anthropic, 2025], Gemini-2.5-flash [Comanici et al., 2025], Kimi-K2 [Team et al., 2025], GPT-4o-mini [Hurst et al., 2024], and GPT-5 [OpenAI, 2025]). For the baseline LLMs, we additionally provide them with detailed definitions of the six value dimensions, along with two real human participants’ value preferences and their action preference rankings in real scenarios as few-shot examples. The models are then tasked with predicting participants’ action preference rankings in real scenarios based on their value preferences. The action rankings generated by each model were subsequently compared against the human-provided rankings to compute alignment similarity. Detailed experimental setups, the full questionnaire, and demographic analysis of the participants are provided in Appendix C and D.

4.3.1 Evaluation Metric

We use two complementary metrics to comprehensively assess the alignment with human decision patterns. The primary metric, termed **Order-Sensitive Similarity (OS-Sim)**, extends the Jaccard index through progressive prefix comparison. Let $i \in \{1, \dots, N_q\}$ represent the question sample in the questionnaire. For two sequences of actions $S^{(i)} = [s_1^{(i)}, s_2^{(i)}, \dots, s_n^{(i)}]$ and $T^{(i)} = [t_1^{(i)}, t_2^{(i)}, \dots, t_n^{(i)}]$, define $S_d^{(i)} \triangleq \{s_1^{(i)}, \dots, s_d^{(i)}\}$ and $T_d^{(i)} \triangleq \{t_1^{(i)}, \dots, t_d^{(i)}\}$. The similarity at depth $d \in \{1, \dots, n\}$ is

a modified Jaccard Similarity $\text{sim}_d^{(i)}(S^{(i)}, T^{(i)}) = \frac{|S_d^{(i)} \cap T_d^{(i)}|}{d}$. The mean OS-Sim is

$$\text{OS-Sim} = \frac{1}{N_q} \sum_{i=1}^{N_q} \text{OS-Sim}^{(i)}, \text{ where } \text{OS-Sim}^{(i)}(S^{(i)}, T^{(i)}) = \frac{1}{n} \sum_{d=1}^n \text{sim}_d^{(i)}. \quad (10)$$

This design captures both immediate preference alignment in early positions and gradual decision pattern matching throughout the sequence. Detailed examples and the advantages of Os-Sim over methods such as *Spearman Rank Correlation* or *Kendall Tau* are provided in Appendix D.4. By applying this method, we can objectively compare the ranking lists generated by our model and the LLMs with the actual lists provided by the participants, allowing for a robust evaluation of model performance.

To address real-world constraints where only the primary action selection matters, we introduce **First-Action Accuracy** as a critical secondary metric. This measure quantifies the percentage of trials where the model’s top-ranked action matches the human subject’s first choice:

$$\text{First-Acc} = \frac{1}{N_q} \sum_{i=1}^{N_q} \mathbb{I}(s_1^{(i)} = t_1^{(i)}), \quad (11)$$

where \mathbb{I} denotes the indicator function. This dual-metric approach evaluates both comprehensive alignment and practical first-choice accuracy.

4.3.2 Results

As shown in Figure 2, ValuePilot’s DMM significantly outperforms state-of-the-art LLMs in aligning with human decision patterns. Our model achieves a mean OS-Sim of 73.16%(±0.43%), surpassing the strongest LLM baseline (GPT-5 at 69.23% ± 0.71%). This highlights its enhanced capability in capturing hierarchical preference structures and nuanced position-sensitive choices.

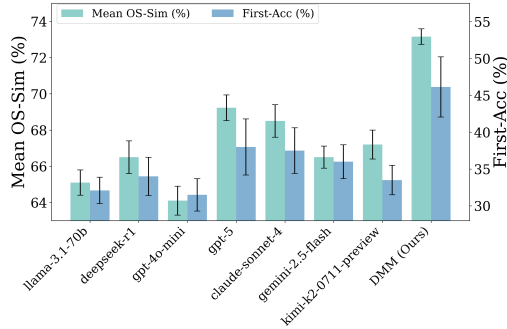


Figure 2: Comparison between our DMM and LLMs in terms of their decision rankings’ similarity to human decisions. DMM achieves the highest scores, outperforming the best baseline (GPT-5) by +3.93% in OS-Sim and +8.13% in First-Acc.

Table 2: Ablation Study for DMM.

Model Variant	OS-Sim (%)	First-Acc (%)
Only Action	60.23 ± 0.17	32.27 ± 0.58
w/o Preference	61.07 ± 0.29	31.82 ± 0.94
w/o Subjective	68.93 ± 0.49	43.45 ± 1.16
w/o Scenario	69.99 ± 0.57	43.64 ± 1.32
DMM (Full)	73.16 ± 0.43	46.14 ± 4.09

Notably, in the practically important First-Action Accuracy metric, our DMM achieves 46.14%(±4.09%), an 8.13% relative improvement over GPT-5 (38.01% ± 3.81%). This demonstrates superior fidelity in modeling primary decision determinants, suggesting that explicit value prioritization offers distinct advantages over implicit pattern learning in LLMs.

4.4 Ablation Study

We evaluate the full DMM model against four ablated variants, each removing one or all crucial components to analyze its impact.

- **Only Action:** The model ranks actions without considering scenario context or individual value preferences, relying solely on action scores from the Value Assessment Network. Eq. (6) changes to $\tilde{V}_{ii'} = \sum_{j=1}^m \frac{1}{1 + e^{-(\rho_j^{a_i} - \rho_j^{a_{i'}})}}$.

- **w/o Preference:** The model removes individual value preferences, evaluating actions based purely on their objective scores from the Value Assessment Network. Eq. (6) changes to $\tilde{V}_{ii'} = \sum_{j=1}^m V_{ii',j}$.
- **w/o Subjective:** The model removes subjective preference adjustments, using only objective action and scenario scores without individual bias adaptation. Eq. (3) changes to $r_j = o_j$.
- **w/o Scenario:** The model ranks actions solely based on their individual value impact, without incorporating scenario-based scaling in the final decision. Equation Eq. (4) changes to $r_{i,j} = r_j^{a_i}$.

Results show that integrating objective action impact, scenario-based adjustments, and individual preference adaptation improves model performance, validating our two-stage framework,

5 Related Works

LLMs for Decision-Making and Planning LLMs have advanced AI decision-making in interactive environments [Nakano et al., 2021, Shuster et al., 2022, Glaese et al., 2022]. Recent methods integrate more explicit reasoning, such as ReAct interleaving reasoning with actions [Yao et al., 2022], or Inner Monologue using feedback for refinement [Huang et al., 2022], contrasting with affordance-dependent models like SayCan [Ahn et al., 2022]. While these, and planning approaches like AutoPlan [Ouyang and Li, 2023], enhance LLM agency, they are often task-oriented, with values implicitly tied to task success. They typically lack explicit modeling of fundamental human values or mechanisms for deep personalization of these values, a gap ValuePilot addresses by explicitly incorporating and reasoning over defined value dimensions.

Aligning AI with Individualized Human Values Aligning AI with human values is crucial. Reinforcement Learning from Human Feedback (RLHF) [Christiano et al., 2017] and Direct Preference Optimization (DPO) [Rafailov et al., 2024] tune models towards *general* or *collective* human value preferences. However, they often struggle with inter-individual variability and fine-grained value distinctions, generally not learning explicit, interpretable models of different value dimensions. ValuePilot focuses on modeling these explicit dimensions and enabling personalization. Role-Playing Language Agents (RPLAs) [Chen et al., 2024] offer persona-driven interactions, with demographic [Huang et al., 2023, Xu et al., 2023, Gupta et al., 2023], character [Shao et al., 2023, Wang et al., 2023], and individualized personas [Salemi et al., 2023, Woźniak et al., 2024, Dalvi et al., 2022]. While RPLAs advance individualized AI, often enhanced by RAG and long-context memory [Chen et al., 2024], their focus is typically on conversational fidelity rather than autonomous, explicit value-based action selection in novel scenarios, especially concerning value trade-offs. ValuePilot extends individualization by using its Dataset Generation Toolkit (DGT) to create structured, value-annotated data, and its Decision-Making Module (DMM) to reason over these explicit values and personalized value preferences for action selection.

Multi-Criteria Decision-Making (MCDM) for Value-Driven Choices MCDA provides frameworks for balancing conflicting criteria, analogous to an AI balancing diverse human values. Methods like MAUT [Dyer, 2016], AHP [Vaidya and Kumar, 2006], and TOPSIS [Papathanasiou et al., 2018] exist, but PROMETHEE [Brans and De Smet, 2016, Brans and Vincke, 1985] is particularly relevant for its ability to handle distinct criteria by computing net outranking flows. Traditional MCDA applications in AI face challenges with static preference weights and deriving criterion scores dynamically. While recent AI-MCDA integrations aim for adaptability, many remain domain-specific or rely on heuristics. A key challenge is learning action-to-value mappings and incorporating individual-specific trade-offs. ValuePilot addresses this by using its DGT to enable the DMM’s Value Assessment Network to learn objective action-value scores. The DMM then integrates these with user-specific value preferences, employing PROMETHEE for principled resolution of value trade-offs. This combination of learned value assessment and MCDM offers a robust framework for personalized value-driven AI agents.

6 Conclusion

In this paper, we presented ValuePilot, a two-phase framework aimed at enhancing personalized decision-making in AI. The framework includes DGT for generating realistic, value-guided scenarios

and DMM for ranking actions based on individual value preferences. Through our experiments, DMM demonstrated a strong alignment with human decision-making, outperforming existing LLMs in replicating human choices. We see ValuePilot as an initial step toward the broader goal of value-sensitive AI. While preliminary, our findings suggest the potential of this direction, and we hope it can offer a foundation for future work on socially aligned and interpretable AI agents.

References

- M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- Anthropic. Claude Sonnet. <https://www.anthropic.com/claude/sonnet>, 2025.
- Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.
- D. E. Berlyne. Conflict, arousal, and curiosity. 1960.
- D. E. Berlyne. Curiosity and exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, 153(3731):25–33, 1966.
- D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- J.-P. Brans and Y. De Smet. Promethee methods. *Multiple criteria decision analysis: state of the art surveys*, pages 187–219, 2016.
- J.-P. Brans and P. Vincke. Note—a preference ranking organisation method: (the promethee method for multiple criteria decision-making). *Management science*, 31(6):647–656, 1985.
- J. Chen, X. Wang, R. Xu, S. Yuan, Y. Zhang, W. Shi, J. Xie, S. Li, R. Yang, T. Zhu, et al. From persona to personalization: A survey on role-playing language agents. *arXiv preprint arXiv:2404.18231*, 2024.
- P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- G. Comanici, E. Bieber, M. Schaekermann, I. Pasupat, N. Sachdeva, I. Dhillon, M. Blistein, O. Ram, D. Zhang, E. Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- B. Dalvi, O. Tafjord, and P. Clark. Towards teachable reasoning systems: Using a dynamic memory of user feedback for continual system improvement. In *Proceedings of the 2022 conference on empirical methods in natural language processing*, pages 9465–9480, 2022.
- E. Diener. Subjective well-being: The science of happiness and a proposal for a national index. *American psychologist*, 55(1):34, 2000.
- A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- J. S. Dyer. Multiattribute utility theory (maut). *Multiple criteria decision analysis: State of the art surveys*, pages 285–314, 2016.
- D. Emelin, R. L. Bras, J. D. Hwang, M. Forbes, and Y. Choi. Moral stories: Situated reasoning about norms, intents, actions, and their consequences. *arXiv preprint arXiv:2012.15738*, 2020.
- A. Glaese, N. McAleese, M. Trkebac, J. Aslanides, V. Firoiu, T. Ewalds, M. Rauh, L. Weidinger, M. Chadwick, P. Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022.
- D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

- S. Gupta, V. Shrivastava, A. Deshpande, A. Kalyan, P. Clark, A. Sabharwal, and T. Khot. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*, 2023.
- J.-t. Huang, W. Wang, E. J. Li, M. H. Lam, S. Ren, Y. Yuan, W. Jiao, Z. Tu, and M. Lyu. On the humanity of conversational ai: Evaluating the psychological portrayal of llms. In *The Twelfth International Conference on Learning Representations*, 2023.
- W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Thompson, I. Mordatch, Y. Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022.
- A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- R. Inglehart, M. Basanez, J. Diez-Medrano, L. Halman, and R. Luijkx. World values surveys and european values surveys, 1981-1984, 1990-1993, and 1995-1997. *Ann Arbor-Michigan, Institute for Social Research, ICPSR version*, 2000.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.
- P. A. Lee, T. Anderson, P. A. Lee, and T. Anderson. *Fault tolerance*. Springer, 1990.
- S. Lyubomirsky, K. M. Sheldon, and D. Schkade. Pursuing happiness: The architecture of sustainable change. *Review of general psychology*, 9(2):111–131, 2005.
- A. H. Maslow. A theory of human motivation. *Psychological Review*, 50(4):370–396, 1943.
- L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- M. Mekni. An artificial intelligence based virtual assistant using conversational agents. *Journal of Software Engineering and Applications*, 14(9):455–473, 2021.
- R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- G. A. Nix, R. M. Ryan, J. B. Manly, and E. L. Deci. Revitalization through self-regulation: The effects of autonomous and controlled motivation on happiness and vitality. *Journal of experimental social psychology*, 35(3):266–284, 1999.
- OpenAI. GPT-5. <https://cdn.openai.com/gpt-5-system-card.pdf>, 2025.
- S. Ouyang and L. Li. Autoplan: Automatic planning of interactive decision-making tasks with large language models. *arXiv preprint arXiv:2305.15064*, 2023.
- J. Papathanasiou, N. Ploskas, J. Papathanasiou, and N. Ploskas. *Topsis*. Springer, 2018.
- L. Qiu, Y. Zhao, J. Li, P. Lu, B. Peng, J. Gao, and S.-C. Zhu. Valuenet: A new dataset for human value driven dialogue system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11183–11191, 2022.
- R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.
- C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.

- H. T. Reis et al. Intimacy as an interpersonal process. In *Relationships, well-being and behaviour*, pages 113–143. Routledge, 2018.
- A. Salemi, S. Mysore, M. Bendersky, and H. Zamani. Lamp: When large language models meet personalization. *arXiv preprint arXiv:2304.11406*, 2023.
- G. Salton and C. Buckley. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523, 1988.
- S. H. Schwartz. Are there universal aspects in the structure and contents of human values? *Journal of social issues*, 50(4):19–45, 1994.
- S. H. Schwartz. Basic human values: Theory, methods, and application. *Risorsa Uomo*, (2007/2), 2007.
- S. H. Schwartz. An overview of the schwartz theory of basic values. *Online Readings in Psychology and Culture*, 2(1), 2012.
- Y. Shao, L. Li, J. Dai, and X. Qiu. Character-llm: A trainable agent for role-playing. *arXiv preprint arXiv:2310.10158*, 2023.
- N. P. Sharma. *Human Behavior in a Social Environment*. StatPearls Publishing, 2022.
- M. Shridhar, X. Yuan, M.-A. Côté, Y. Bisk, A. Trischler, and M. Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning. *arXiv preprint arXiv:2010.03768*, 2020.
- I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024.
- K. Shuster, J. Xu, M. Komeili, D. Ju, E. M. Smith, S. Roller, M. Ung, M. Chen, K. Arora, J. Lane, et al. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*, 2022.
- K. Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21, 1972.
- K. Team, Y. Bai, Y. Bao, G. Chen, J. Chen, N. Chen, R. Chen, Y. Chen, Y. Chen, Y. Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- T. R. Tyler. Social justice. 2015.
- O. S. Vaidya and S. Kumar. Analytic hierarchy process: An overview of applications. *European Journal of operational research*, 169(1):1–29, 2006.
- Z. M. Wang, Z. Peng, H. Que, J. Liu, W. Zhou, Y. Wu, H. Guo, R. Gan, Z. Ni, J. Yang, et al. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *arXiv preprint arXiv:2310.00746*, 2023.
- S. Woźniak, B. Koptyra, A. Janz, P. Kazienko, and J. Kocoń. Personalized large language models. *arXiv preprint arXiv:2402.09269*, 2024.
- B. Xu, A. Yang, J. Lin, Q. Wang, C. Zhou, Y. Zhang, and Z. Mao. Expertprompting: Instructing large language models to be distinguished experts. *arXiv preprint arXiv:2305.14688*, 2023.
- W. Xu, M. J. Dainoff, L. Ge, and Z. Gao. From human-computer interaction to human-ai interaction: new challenges and opportunities for enabling human-centered ai. *arXiv preprint arXiv:2105.05424*, 5, 2021.
- J. Yang, A. Prabhakar, K. Narasimhan, and S. Yao. Intercode: Standardizing and benchmarking interactive coding with execution feedback. *Advances in Neural Information Processing Systems*, 36, 2024.
- S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

- R. M. Yerkes, J. D. Dodson, et al. The relation of strength of stimulus to rapidity of habit-formation. *Journal of comparative neurology and psychology*, 18(5):459–482, 1908.
- M. Zhu, Y. Weng, L. Yang, and Y. Zhang. Personality alignment of large language models. *arXiv preprint arXiv:2408.11779*, 2024.

A Dataset

The ValuePilot framework is designed for flexibility, allowing users or system designers to specify the set of value dimensions most relevant to their application domain or research focus. For the experiments presented in this paper, and to demonstrate ValuePilot’s capabilities, we selected six core value dimensions to generate our Dataset. This section details the rationale behind this specific selection.

A.1 Rationale for Selected Value Dimensions

In understanding human-AI interaction, particularly in everyday environments like the home, identifying the underlying value dynamics that drive choices is paramount. Our work utilizes six core value dimensions: **Curiosity**, **Energy**, **Safety**, **Happiness**, **Intimacy**, and **Fairness**. These dimensions were selected for their relevance to human behavior, drawing inspiration from established psychological frameworks such as Maslow’s hierarchy of needs [Maslow, 1943] and Schwartz’s theory of basic human values [Schwartz, 2012].

Broad Coverage Across Interaction Types These six dimensions encompass both single-agent (e.g., an individual’s *Curiosity*, *Energy* management, or need for *Safety*) and multi-agent interactions (e.g., seeking *Happiness* or *Intimacy* through social connections, or upholding *Fairness* in group settings [Sharma, 2022]). This provides a comprehensive lens for analyzing diverse agent behaviors.

Distinctiveness and Minimal Overlap While human values are interconnected, the selected dimensions exhibit largely distinct characteristics. For instance, *Curiosity*-driven exploration is not inherently tied to physical *Energy* expenditure, and ensuring environmental *Safety* addresses different needs than fostering social *Intimacy* [Sharma, 2022]. This conceptual separation aids in clear annotation and modeling.

Alignment with Foundational Theories of Human Motivation A meticulous selection process, integrating sociological and psychological perspectives, identified these dimensions as pivotal in shaping human behavior, particularly in domestic contexts.

- **Curiosity:** Reflects the drive to explore and understand, aligning with self-actualization needs and the value of stimulation [Schwartz, 2012, Berlyne, 1966, 1960].
- **Energy:** Represents the physiological need for vitality and capacity for activity, fundamental to Maslow’s physiological level [Maslow, 1943, Nix et al., 1999].
- **Safety:** Pertains to security and stability, a core need in Maslow’s hierarchy and a key security value in Schwartz’s model [Maslow, 1943, Schwartz, 2012].
- **Happiness:** An overarching goal related to well-being and life satisfaction, emerging from the fulfillment of various needs, including esteem and self-actualization [Diener, 2000, Lyubomirsky et al., 2005].
- **Intimacy:** Fulfills the need for close connection, belonging, and love, crucial for psychological well-being [Maslow, 1943, Reis et al., 2018].
- **Fairness:** Relates to justice, equality, and reciprocity in social interactions, aligning with needs for belonging, respect, and universalism values [Schwartz, 2012, Tyler, 2015].

The choice of these six value dimensions aims to span the spectrum of needs described by Maslow, ensuring our analysis is grounded in a well-established psychological framework. They are frequently encountered in everyday domestic life, directly relating to daily living and personal fulfillment. While other values exist, these provide a robust and practical set for modeling prevalent human-agent interaction dynamics. Our goal with this specific set is to demonstrate a pertinent lens for examining and enhancing human-AI symbiosis.

A.2 Hierarchical Dataset Architecture and Quality Assurance

The proposed dataset employs a six-tier hierarchical structure (Datasets 1–6) to systematically evaluate computational models across escalating complexity levels of multi-dimensional value interactions. Each tier corresponds to distinct combinatorial configurations of six core value dimensions: *curiosity*,

energy, safety, happiness, intimacy, and fairness. As detailed in Table 3, Dataset 1 isolates single-value dynamics (1-D), while subsequent tiers progressively integrate dimensions up to Dataset 6, which encapsulates full six-dimensional (6-D) interactions.

Table 3: Hierarchical Dataset Composition with Filtering Metrics

Dataset	Data Filtering				Final Scenarios		Final Actions	
	Train		Test		Train	Test	Train	Test
	Removed	Rate (%)	Removed	Rate (%)				
1-D	2,333	12.98	430	12.08	1,771	355	15,644	3,129
2-D	2,632	15.14	430	12.25	1,689	353	14,753	3,080
3-D	3,264	18.16	501	17.20	1,704	342	14,712	2,411
4-D	3,688	22.36	728	20.59	1,609	332	12,815	2,807
5-D	3,584	23.36	863	24.38	1,412	323	11,761	2,677
6-D	4,083	22.85	901	25.17	1,714	334	13,788	2,678
Total	19,584	19.00	3,853	18.67	9,899	2,039	83,473	16,782

To mitigate potential biases in generated content, we implement a two-stage quality assurance protocol. First, an automated procedure removes value-related pairs explicitly mentioned in input prompts from both scenarios and actions. As quantified in Table 3, this eliminates 12.98%–23.36% of training actions and 12.08%–25.17% of test actions across complexity tiers. Second, the manual phase involves iterative refinement to correct implausible scenarios and inconsistent actions, prioritizing behavioral coherence, interaction plausibility, and dimensional proportionality.

The aggregate statistics validate the stratified design of the dataset, containing 9,899 training scenarios (83,473 actions) and 2,039 test scenarios (16,782 actions) across all tiers. This architecture enables granular analysis of model robustness against dimensional scalability while maintaining methodological rigor through systematic bias mitigation.

A.3 Dataset Characteristics

For the dataset we generated, we analyzed the characteristics of the text from multiple perspectives.

Scenario Diversity (Figure 3) We extract Term Frequency–Inverse Document Frequency (TF-IDF) [Sparck Jones, 1972, Salton and Buckley, 1988] features from scenario texts and apply Latent Dirichlet Allocation (LDA) [Blei et al., 2003] for topic modeling, followed by Uniform Manifold Approximation and Projection (UMAP) [McInnes et al., 2018] for dimensionality reduction. The well-separated clusters, labeled by dominant topics, indicate clear categorization, with some overlap reflecting real-world complexity. These results highlight the diversity of scenarios, ranging from everyday activities like cooking to unexpected events such as storms.

Distribution of six-value dimensions (Figure 4a and Figure 4b) We counted the number of times that the scores of all six value dimensions in the dataset appeared as positive and negative values in the scenario. Most of the scores of the six value dimensions appeared as positive values, and the total number of the six dimensions is basically the same, though the negative scores of *Energy* are relatively more. The blue box line represents a positive value, and the purple box line represents a negative value.

Word Frequency Analysis (Figure 4c) We performed word frequency analysis on the dataset scenes after removing *stop words*, which are items with less statistical significance such as prepositions, adverbs, quantifiers, etc., and found many interesting phenomena. Among the nouns, *kitchen* and *living room* appear most frequently, representing the places where scenarios often occur. *Preparing* and *cooking* are the most common verbs, representing actions that often occur in the scenario. Notably, *board game* appears frequently. We think that it is reasonable to play board games as a way of multi-person interaction in a home environment that is compatible with happiness, intimacy, and fairness at the same time.

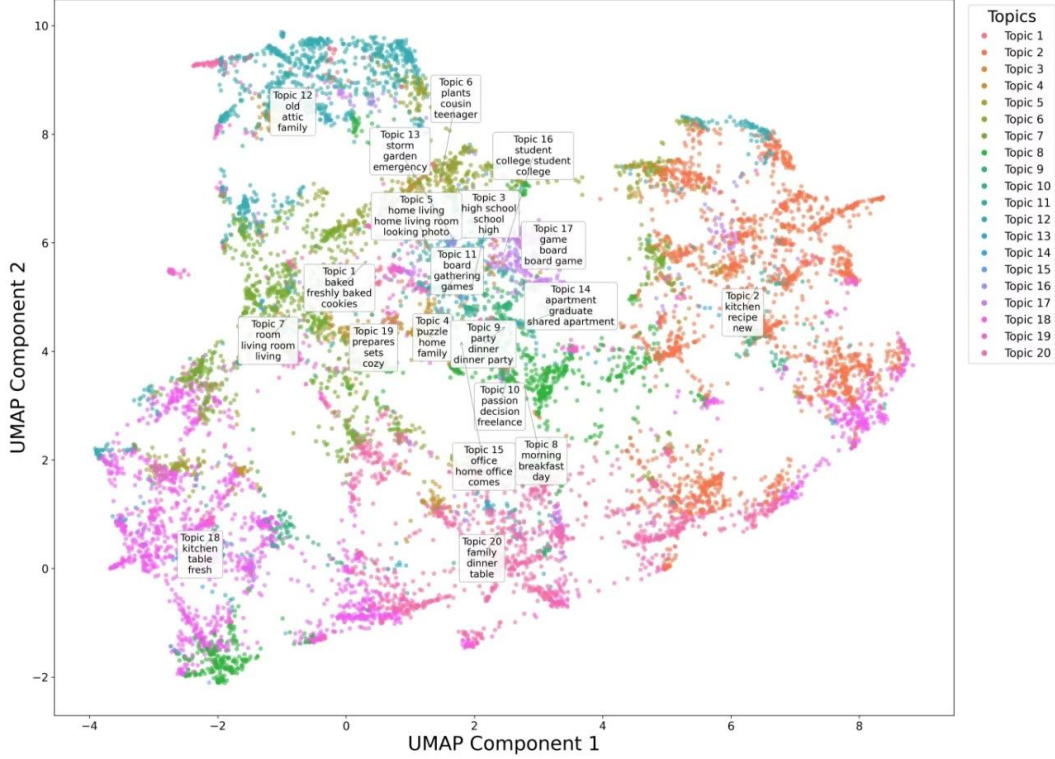


Figure 3: The Topic Space of Scenarios. Using Latent Dirichlet Allocation (LDA), the model automatically groups unlabeled scenario texts into topics based on word co-occurrence patterns. Each cluster is labeled with its dominant topic and annotated with its top three keywords, revealing distinct groupings.

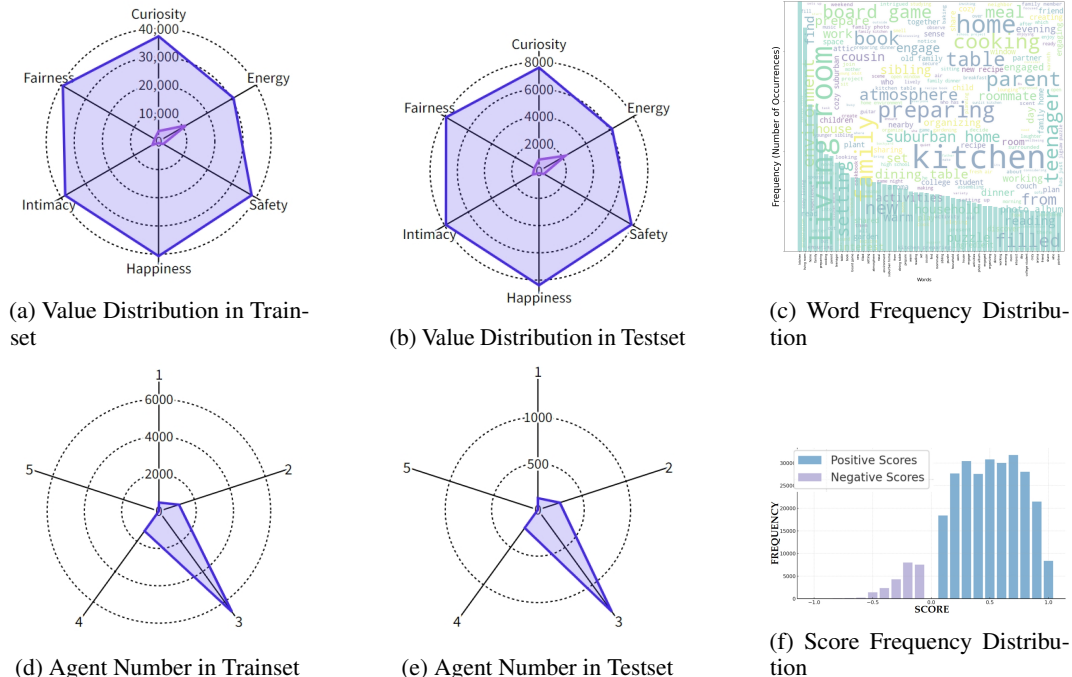
Number of Agents (Figure 4d and Figure 4e) We found that the number of agents in the scenarios of the dataset ranges from 1 to 5, and the number of scenarios containing 3 agents is the largest. In the training set and the test set, the corresponding proportion of the number of agents is basically the same.

Score Frequency Distribution (Figure 4f) We calculated the score distribution of all value dimensions in our dataset.

B Action Selection Module

Algorithm 1 details the procedure used by the Action Selection Module to integrate the Value Assessment Network’s objective value scores with personalized user preferences, culminating in a ranked sequence of actions.

A key parameter in this integration is the weight w , which balances the influence of preference discrepancy scores against objective value scores (as seen in Eq. (3) in the main paper). The adjustability of w allows the system to be tailored for different balances between objective data and subjective alignment, conceptually similar to an LLM’s temperature setting. For tasks requiring logical reasoning with some flexibility, a w range of 0.25–0.5 is generally considered appropriate (see Table 4). Our empirical testing, which included a grid search and evaluation against human decision patterns, confirmed this, with $w = 0.3$ specifically yielding the best alignment with human choices in our experiments. represents a deliberate balance, giving notable weight to subjective alignment without overshadowing objective assessments. For all experiments presented in this paper, w was consistently set to 0.3.

Table 4: Conceptual Guidance for Subjective Weight (w) Settings

Task Type / Focus	Suggested w Range	Behavioral Characteristic
Objective Focus	0.0–0.25	Deterministic or fact-based outcomes
Analytical Reasoning	0.25–0.5	Logical reasoning with flexibility
Balanced Consideration	0.5–0.75	Well-rounded, considerate decisions
Subjective Emphasis	0.75–1.0	Outcomes strongly driven by personal preference

C Model Training and Experiment

C.1 Implementation Details

In Experiment 1, we train the networks on a train set comprising 80% of the total dataset, and then test it on 20% of the total dataset, aiming to show the efficiency of Value Assessment Network. The experiments are conducted on a system with a single NVIDIA GeForce RTX 4090 GPU (24 GB memory), using CUDA Version 12.6 and driver version 560.70. The models were optimised using the Adam optimizer [Kingma and Ba, 2014]. The hyperparameters in this optimal configuration consisted of a learning rate range of $[1e-3, 1e-5]$ and a weight decay range of $[1e-5, 1e-4]$ for the Value Assessment Network during Phases 1 and 2. The total training time was approximately 15 minutes.

C.2 Model Training Procedure

We use a curriculum learning approach [Bengio et al., 2009] to train our models, helping our model build a more robust understanding of basic concepts before tackling more challenging scenarios. To ensure balanced representation, the entire dataset is divided into l (4 in our experiment) chunks. Each chunk contains a mix of data generated from all combinations of n (6 in our experiment) value dimensions. The training process follows a predefined order of these groups. This sequence

Algorithm 1 Action Selection Module (DMM)

Input:

Scenario Objective Scores: $\rho^s = [\rho_1^s, \dots, \rho_m^s]$
Action Objective Scores (for N actions, m values): $\rho^a = [\rho^{a_1}, \dots, \rho^{a_N}]$, where
 $\rho^{a_i} = [\rho_1^{a_i}, \dots, \rho_m^{a_i}]$
Transformed User Subjective Preferences: $\mathbf{p}' = [p'_1, \dots, p'_m]$
Weight for subjective factor: w

Parameter:

Scenario (text)
Candidate Actions (list of N actions): $\mathcal{A} = \{action_1, \dots, action_N\}$
Value Dimensions (list of m values): $\mathcal{V} = \{value_1, \dots, value_m\}$

Output: *RankedActions***# 1. Contextualized Scoring**

```
1: for  $i \in \{1, \dots, N\}$  (for each action  $action_i$ ) do
2:   for  $j \in \{1, \dots, m\}$  (for each value dimension  $value_j$ ) do
3:      $d_j^s \leftarrow 1 - ||\rho_j^s| - p'_j|$ ;  $d_j^{a_i} \leftarrow 1 - ||\rho_j^{a_i}| - p'_j|$ 
4:      $r_j^s \leftarrow wd_j^s + (1-w)\rho_j^s$ ;  $r_j^{a_i} \leftarrow wd_j^{a_i} + (1-w)\rho_j^{a_i}$ 
5:      $r_{i,j} \leftarrow \frac{1}{1+e^{-|r_j^s|}} \times r_j^{a_i}$  {Final score for action  $i$  on value  $j$ }
6:   end for
7: end for

# 2. Multi-Criteria Decision-Making (PROMETHEE)
8: for  $i \in \{1, \dots, N\}$  do
9:   for  $i' \in \{1, \dots, N\}$  where  $i' \neq i$  do
10:     $V_{ii',j} \leftarrow \frac{1}{1+e^{-(r_{i,j}-r_{i',j})}}$  for each  $j \in \{1, \dots, m\}$ 
11:     $\tilde{V}_{ii'} \leftarrow \sum_{j=1}^m p'_j \cdot V_{ii',j}$  {Aggregated preference of  $action_i$  over  $action_{i'}$ }
12:   end for
13: end for

# 3. Calculate Net Outranking Flows
14: for  $i \in \{1, \dots, N\}$  do
15:    $\phi_i^+ \leftarrow \frac{1}{N-1} \sum_{i' \neq i} \tilde{V}_{ii'}$ ;  $\phi_i^- \leftarrow \frac{1}{N-1} \sum_{i' \neq i} \tilde{V}_{i'i}$ 
16:    $\phi_i \leftarrow \phi_i^+ - \phi_i^-$ 
17: end for

# 4. Rank Actions
18:  $RankedActions \leftarrow \text{Sort}(\mathcal{A}, \text{key} = \phi, \text{descending}=\text{True})$ 
```

is designed to help the model to adapt and learn efficiently, minimizing the risk of overfitting or underfitting.

Algorithm 2 provides a detailed overview of the curriculum learning approach used to train our models, guiding them through progressively complex data to enhance learning efficiency and performance. In the context of the algorithm you provided, the sequence of dataset S_{data} refers to the predefined order in which different subsets of your data are used during the model training process. This sequence is crucial in the curriculum learning approach because it dictates the order in which the model encounters data of varying difficulty or complexity.

Training Order	Phase 1		Phase 2		Time (min:sec)	AvgAcc (%)
	lr	wd	lr	wd		
3,4,2,5,1,6	1e-4	1e-5	1e-5	1e-4	22:43	66.06
6,1,5,2,4,3	1e-4	1e-5	1e-5	1e-4	23:25	66.15
1,2,3,4,5,6	1e-4	1e-5	1e-5	1e-4	22:23	66.00
6,5,4,3,2,1	1e-4	1e-5	1e-5	1e-4	24:53	63.80
6,1,5,2,4,3	1e-4	1e-5	1e-4	1e-4	23:25	65.00
6,1,5,2,4,3	1e-3	1e-5	1e-5	1e-4	14:16	66.70
6,1,5,2,4,3	1e-5	1e-5	1e-5	1e-4	133:50	66.06

Table 5: Comparison of Different Training Orders and Hyperparameter Configurations. **lr** stands for learning rate. **wd** stands for weight decay. Threshold $t = 0.2$ for AvgAcc.

Encoder	AvgAcc (%)
T5-base	66.70
Flan-T5	66.00
BERT	65.45
RoBERTa	65.69

Table 6: Selection study of encoder models. Threshold $t = 0.2$ for Average Accuracy.

Algorithm 2 Model Training Procedure

```

1: Set the sequence of dataset:  $S_{data} = [6,1,5,2,4,3]$ 
2: Initialize number of chunks to  $l$ 
3: for each chunk  $i$  from 0 to  $l - 1$  do
4:   Initialize Datasets
5:   for each  $n$  in training order do
6:     Combine_Dataset.concat(Dataset[n])
7:     Initialize early stopping parameters
8:     for each epoch do
9:       Train model, validate model, log metrics, and check for early stopping
10:      if early stopping criteria are met then
11:        Save model
12:      end if
13:    end for
14:  end for
15: end for

```

C.3 Additional Results

C.3.1 Hyperparameters

As shown in Table 5, Our Networks are trained with different training order and hyperparameters. The results show that a training order of [6, 1, 5, 2, 4, 3] and [lr, wd] sets of [1e-3, 1e-5], [1e-5, 1e-4] reaches the highest accuracy and the shortest training time.

C.3.2 Encoder Model

To prove that our decision making network works well for a variety of relatively small language models, We replaced the encoder’s model, and the experimental results are shown in Table 6.

The T5-base model outputs the highest Average Accuracy. However, upon transforming different encoder models, the final similarity remains relatively consistent, which proves that our fine-tuning method is applicable to different pre-trained language models.

D Human Study

D.1 Demographic Analysis

The study cohort (N=40) exhibited substantial demographic diversity across age, ethnicity, and socioeconomic dimensions. Subjects spanned adolescents (12.5%, n=5), young adults (47.5% aged 18-25), and middle-aged individuals (2.5% aged 51-60), with balanced gender representation (55.0% female, 45.0% male). Educational attainment ranged from secondary education to advanced degrees, with 75% holding bachelor’s qualifications or higher. Occupational diversity spanned 11 distinct professional categories, including academic researchers (15.0%), technical specialists (7.5%), and commercial professionals (5.0%), though students predominated (52.5%). While the sample is skewed toward younger, educated, urban populations—a common limitation in survey-based studies—this demographic orientation does not substantially compromise the internal validity of findings, though broader generalizability to rural or less-educated groups requires confirmation through targeted sampling in future studies.

In addition to these demographic characteristics, we collected subjects’ preference ratings on six value dimensions: *curiosity*, *energy*, *safety*, *happiness*, *intimacy*, and *fairness*. As shown in Figure 5 and Figure 6, these preference distributions were notably broad, indicating that our sample captured a wide range of individual differences in value orientations. The pairwise correlation coefficients among the dimensions are also presented, revealing varying degrees of interrelationship across curiosity, energy, safety, happiness, intimacy, and fairness.

Furthermore, Figure 7 provides a two-dimensional t-SNE visualization of the subjects’ preference profiles. The relatively dispersed clustering in this plot suggests that there is no single dominant pattern; instead, subjects’ preferences exhibit substantial heterogeneity. Taken together, the breadth of preference distributions, alongside the demographic diversity described earlier, underscores the representativeness of the sample in terms of varying backgrounds and value priorities.

D.2 Comparative Experiment

We designed a test questionnaire (see D.3) containing 11 formal test questions. The questionnaire included a detailed introduction to six value dimensions at the beginning, where subjects rated the importance of these dimensions in their daily lives on a scale from 0 to 1, with higher scores indicating greater importance.

To ensure subjects fully understood the experiment’s design, a simulation test with three preliminary questions was conducted. This allowed subjects to familiarize themselves with the process, enabling them to provide more accurate data during the formal test. Following the simulation test, subjects re-scored the importance of the six value dimensions, generating a final scoring vector. For each scenario described in the questionnaire, subjects were asked to rank their willingness to choose different actions based on their value dimension scores. This produced a six-dimensional value scoring vector and corresponding action ranking data. Subjects were also asked to identify the value dimensions associated with each action in the scenarios, which facilitated the analysis of the rationality of the data.

We then input the scenario descriptions and corresponding action lists into our model, along with each subject’s six-dimensional value scoring vector. The model processed this data and outputted action selection rankings based on the subject’s value vector. As a comparative experiment, we also input the six-dimensional scores of the subjects into large language models. For each participant, we obtained action selection rankings from both our model and the large language models. We then used the proposed sequence similarity measurement method OS-Sim and First-Acc to compare the ranking lists generated by our model and the large language models with the actual lists provided by the subjects. Then we calculate the Mean OS-Sim, based on the comparative results of different models, fitting the choice of human subjects. Through confidence analysis, we found that the fitting degree of our model to human action selection is better than that of the current mainstream large language model, which proves the effectiveness of our model in analyzing the value dynamics of agents.

We also conducted a experiment to compare the effectiveness of different Multi-Criteria Decision Analysis (MCDA) methods in accurately ranking actions within a given scenario. Based on the experimental results, PROMETHEE shows the highest average accuracy at 73.16% with a low error margin of $\pm 0.43\%$, making it the most reliable method in this scenario, as shown in Table 7.

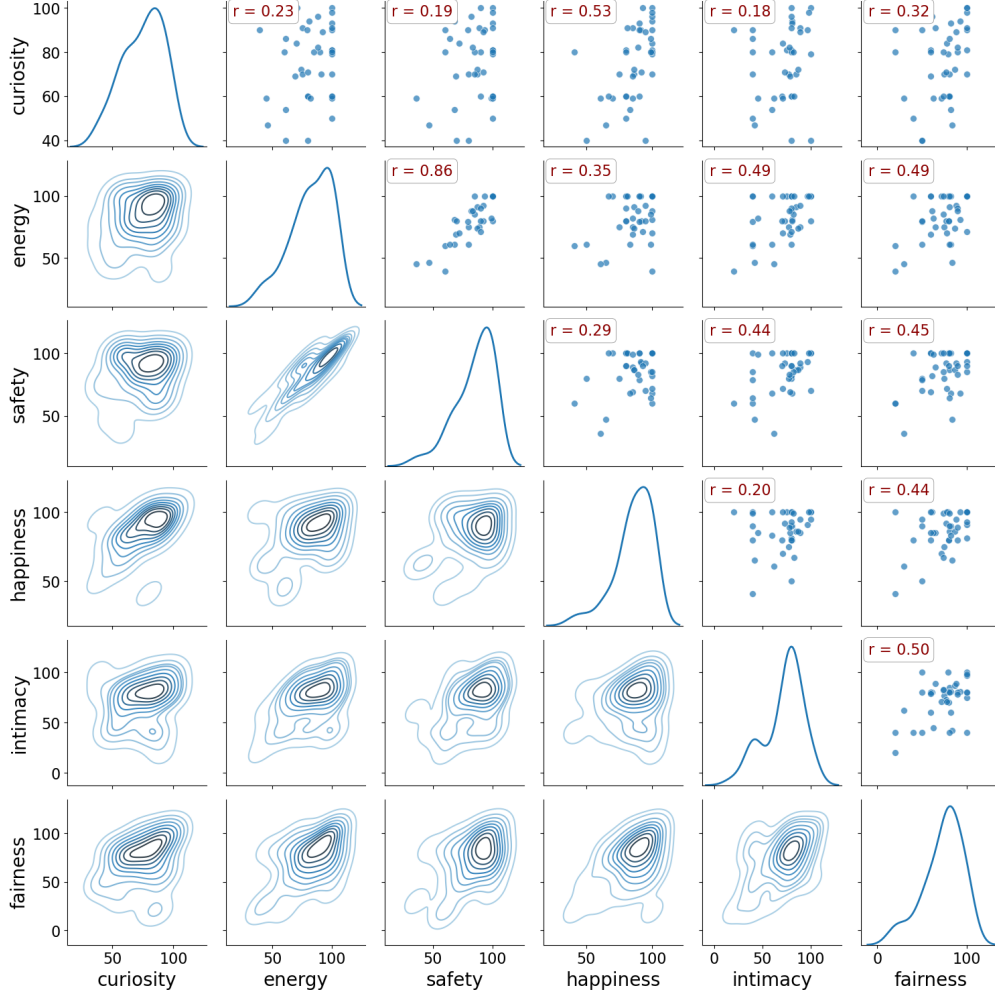


Figure 5: Pairwise Correlations and Distributions of Value Preferences Across the Six Dimensions with Pearson correlation coefficient (r)

Method	Mean OS-Sim (%)
PROMETHEE	73.16 ± 0.43
AHP	72.9 ± 0.56
MAUT	67.08 ± 0.47
TOPSIS	68.19 ± 0.53

Table 7: Comparison of Different Ranking Methods

D.3 Test Questionnaire

Test Instructions: We have constructed a value-driven decision-making model for complex environments. The input to the model consists of six value dimensions:

- **Curiosity:** An intrinsic desire for knowledge about the world around us. It motivates individuals to actively seek new information and experiences, expanding their understanding of things and enriching their perspectives through exploration. People with high curiosity tend to: actively seek knowledge, bravely explore unknown areas, ask questions to deepen understanding, maintain an open mind, and accept new and diverse views.
- **Energy:** The pursuit and satisfaction of physiological energy needs, including food, water, air, sleep, and exercise. It means striving for enough energy to maintain physical and mental

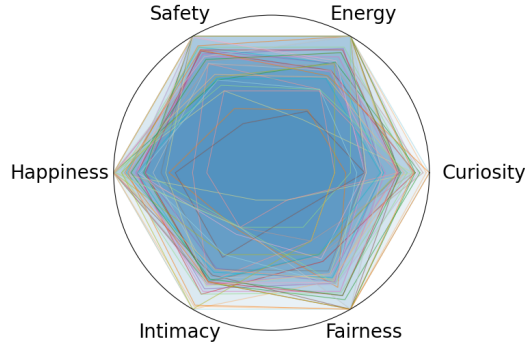


Figure 6: Radar chart showing the value preferences of all subjects. Each polygon represents a subject, and different colors of outline correspond to different subjects.

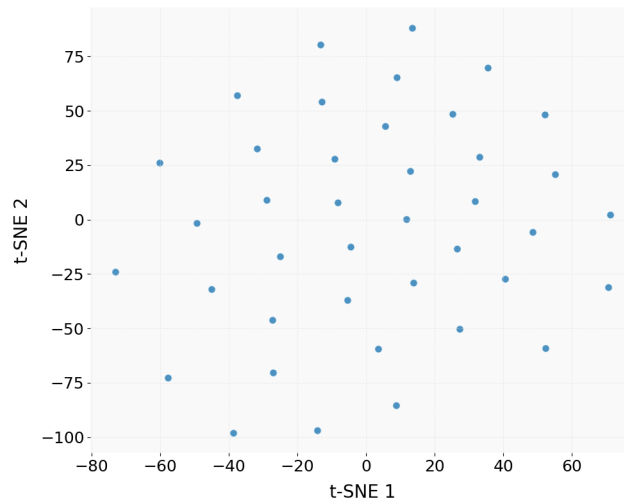


Figure 7: t-SNE Projection of All Subjects' Value Preferences

health. People with high energy needs tend to: maintain a healthy diet, ensure adequate hydration, secure enough and quality sleep, engage in moderate exercise, and breathe fresh air regularly.

- **Safety:** The concern for one's physical and psychological health and the need for stability and predictability in one's environment. This value dimension triggers actions to ensure personal safety, stability in life, and protection from pain, threats, and illnesses, as well as safeguarding personal property. People with high safety needs tend to: maintain physical safety measures, cultivate and maintain mental health, seek stable living conditions, plan for emergencies, and protect personal property.
- **Happiness:** A desire for positive emotions and a sense of fulfillment. It drives actions that bring joy, including cultivating positive relationships, achieving personal goals, and enjoying pleasant activities. Happiness manifests in the real world as pleasure and satisfaction with life. People with high happiness tend to: pursue positive emotions, establish meaningful relationships, strive towards personal goals, seek enjoyment and entertainment, and cultivate a sense of accomplishment.
- **Intimacy:** Emphasizes deep connections and emotional resonance in various relationships, including family, friendship, romantic, and subordinate relationships. It reflects actions taken to achieve mutual support, understanding, and shared experiences. In the real world, intimate relationships provide warmth and support, promoting individual psychological health and well-being. People with high needs for intimacy tend to: share thoughts, feelings, and experiences, provide support and understanding, willingly share personal privacy, engage in positive and meaningful communication, and pursue a sense of belonging.

- **Fairness:** Reflects an individual's pursuit of justice and equality. It motivates actions to create fair and equal conditions in the real world, including actions to ensure fair treatment in social and organizational environments. The goal is to ensure that every individual has equal opportunities and rights. People who highly value fairness tend to: respect equal rights and treatment, promote fair opportunities for everyone, participate in promoting equal decision-making, and treat others justly and fairly.

We hope you fill in the questionnaire. We ensure that the questionnaire is filled anonymously and the obtained data is only used to test the generalization effect of the model in the home environment without commercial use.

Simulation Test: Before addressing specific questions, there is a simulation test to help you understand the specific format of the test.

Please rate your attention to the six value dimensions in your daily life. Ratings are on a scale from 0 to 1, with higher scores indicating greater attention to the value dimension.

- **Curiosity:**_____
- **Energy:**_____
- **Safety:**_____
- **Happiness:**_____
- **Intimacy:**_____
- **Fairness:**_____

Below are three simulation test questions below. Please bring the first-person perspective into the following scene. What actions do you tend to choose in this scene?

Please fill in the form below each question. The first step, for each action, please give what you think the value dimension of the action is. **For each action, please indicate the value dimensions it represents by placing a checkmark (✓) in the appropriate boxes.**

Action	Curiosity	Energy	Safety	Happiness	Intimacy	Fairness	Other Values
1	✓			✓			
2				✓	✓		
3		✓		✓			
...							

Examples of other value dimensions:

1 Responsibility: Taking the initiative to pick up waste paper falling on the ground and throw it into the trash can may be a test of responsibility.

2 Ambitiousness: Choosing to read at home on a good weekend may be out of ambition.

3 Authority: Forcing children to do something may be out of the consideration of maintaining their own family authority. Etc.

The second step is to prioritize each action in the 'action selection' according to your willingness to choose, or the possibility of you making the action from high to low. (Blank box to fill in each action before the corresponding serial number 1234...) (Note: Different actions are selected for the parallel occurrence relationship in the same scene, and do not arrange in the order of the occurrence time in the scene.)

Action Preference Sort	1	2	3	4	...
Action Index					

Formal Test: Presumably, you already know how to test. Here is a formal test with 11 questions.

Please re-score yourself on the importance of the six value dimensions in daily life. Ratings are on a scale from 0 to 1, with higher scores indicating greater attention to the value dimension.

- **Curiosity:**_____
- **Energy:**_____
- **Safety:**_____
- **Happiness:**_____

- **Intimacy:**_____
- **Fairness:**_____

Below are 11 formal questions, please finish them in the same requirements as before.

Q1:

Scenario: You are a 20-year-old college student. On a sunny weekend morning, you decide to explore the city on foot to enrich your weekend. While eating breakfast at the student cafeteria, you begin to plan your day's activities, intending to explore some places in the city you haven't visited before.

Action list:

1. You plan to first go to an online library to collect e-books about the city's development history, so you can have a basic understanding of the city before exploring.
2. You are hesitant about exploring unknown corners of the city and think it is better to stick to visiting main and officially recommended attractions to avoid unnecessary trouble.
3. You plan to try a mobile app to experience the city's new AR historical guide application, exploring the city in a different way through virtual reality technology.
4. You believe that the true joy of hiking adventures lies in discovering the unknown. So, you decide not to make any detailed plans but instead choose a small alley or street you have never walked before, following your intuition to explore the city..

Q2:

Scenario: You, B, and C are university roommates. The dorm is usually messy, and no one cleans the common space. Last week, during the routine dormitory cleanliness inspection, you were warned by the dormitory administrator. At this moment, the three of you are seriously discussing how to distribute household chores in the future, with a whiteboard at the door displaying the categories of chores you just listed.

Action list:

1. To prevent 'free-riding', you seriously propose to establish a detailed schedule for household chores and assign tasks to each person.
2. To prevent omissions, you suggest that everyone think about whether the table listed on the whiteboard fully includes all the chores.
3. You think it is unnecessary to assign chores too specifically since you are roommates living together, so you suggest that everyone just do their chores as they go, reminding each other when someone has time.
4. You don't want to make things too complicated and propose to decide the assignment of chores by rolling dice.
5. You suggest that it's not a big deal and not to be too serious; proposing to go out for dinner together while discussing the distribution of chores.

Q3:

Scenario: You, A, and B are a family living together in a suburban area of a quiet town. On the weekend, the three of you are relaxing at home, listening to soft music, enjoying the tranquility of the afternoon. The sun is shining outside, leaves gently swaying, a breeze passing by, creating a relaxed and pleasant atmosphere. You are sitting on the living room sofa, exchanging your moods and recent trivialities.

Action list:

1. To enhance family intimacy, you suggest playing a simple riddle game together.
2. You feel this weekend should be about relaxation, so you suggest preparing some snacks and watching a light-hearted comedy movie together.
3. You think it's a good opportunity to do some housework during the weekend rest time, such as cleaning the living room or organizing clutter, to make the home more tidy and comfortable.
4. You suddenly remember a fresh activity and suggest going for a walk in the nearby park together, breathing fresh air and relaxing.
5. You suggest using this time to prepare dinner together, trying out some new dishes to enhance family bonding.

Q4:

Scenario: You are in the kitchen adjusting the stove temperature, as dinner for your family is being

cooked. B is your husband, tired from work, sitting idly in the kitchen. C is your daughter, focused on a complex school project she has been working on for a while. A severe storm has started outside.

Action list:

1. To help B relax, you go to talk to him and prepare a hot drink for him.
2. To prevent accidents from the storm, you check all doors and windows to ensure they are securely locked.
3. To prevent your daughter from being disturbed by the noise of the stove, you close the kitchen door.
4. To provide some suggestions and help, you discuss the project with C.
5. To create a cozy atmosphere for dinner, you start setting the table.
6. Seeing C focused, you remind her to take a break and relax, joining you for dinner later.

Q5:

Scenario: You are an enthusiastic middle school student interested in astronomy. B is your brother, a Ph.D. student, and C is your younger sister. It's holiday time, and all three of you are at home. You are casually lying on the living room sofa, flipping through an atlas; B is reading about aerodynamics, and C is preparing a smoothie. An unopened world map puzzle you just bought is on the table, and a computer next to the table displays a news article about breakthroughs in the field of astronomy discussing the possibility of life on an exoplanet.

Action list:

1. Looking at the world map on the table, you can't resist opening it and starting to puzzle.
2. You propose to C to play a quiz game about the capitals of countries using the atlas, with B acting as the judge.
3. Inspired by the article on the computer, you discuss with C whether you believe in the existence of extraterrestrial life.
4. On a whim, you build a fort of chairs and blankets on the sofa, imagining it as a spacecraft headed to the planet.

Q6:

Scenario: You are a high school student who has just moved to a new community. After school, you sit in your room feeling a bit bored. B, your younger brother, built a sophisticated treehouse with the help of C this morning. He is now happily playing with toys inside the treehouse, which is built very high, taller than your height. C, your neighbor, is a retired carpenter, sitting on his balcony drinking coffee, peacefully watching B play.

Action list:

1. To ensure the treehouse is sturdy and your brother is safe, you pick up a toolbox to check for any loose bolts and observe the branches to see if they are strong enough to support the treehouse.
2. To alleviate your boredom, you decide to bring some toys and join your brother in the treehouse to play together.
3. Curious about the interior structure of the treehouse, you decide to climb up and see how it is built.
4. Seeing the exquisitely built treehouse, you become curious about C's career and start a conversation with him.

Q7:

Scenario: You and B are college students living together off-campus. One hot summer noon on a weekend, both of you are at home. You go to the kitchen to find something cool to drink, open the fridge, and pour the last bit of orange juice into your cup, eager to drink it to relieve the heat. Just as you are about to finish it, you realize it's the last bottle in the fridge and notice B looking at you.

Action list:

1. You get another cup and pour half the juice for B.
2. You plan to pretend you didn't notice it was the last cup, drink it all now, and apologize verbally to B later.
3. You plan to drink it all first, then as an apology, go out in the heat later to buy a few more bottles.
4. You think it's just a cup of juice and not a big deal, so you drink it all and then go back to your activities.
5. You suggest playing rock-paper-scissors with B, where the winner gets the juice, to make the situation fair.
6. You stop for a moment to think, then ask B if he wants to drink, and give the juice to him.

Q8:

Scenario: You are a teenage boy living in a villa, currently using 3D modeling software in your room to create a digital model of a dinosaur skeleton for a school extracurricular research project, currently facing a bottleneck and unsure how to proceed. B is your father, a paleontologist, sitting at a desk across the room, focused on researching the latest discoveries in his field. C is your pet, a gentle Golden Retriever, comfortably curled up at B's feet on a cozy rug. Near the window in the room, workers are carrying in your family's newly purchased astronomical telescope, getting ready to set it up.

Action list:

1. Attracted by the telescope, you go over to join in the excitement, curiously asking the workers about how to operate the telescope.
2. Deciding to shift your focus, you go to pet C, covering him with a blanket to keep warm.
3. You decide to talk to B about the bottleneck in making the dinosaur model, hoping B can help open your mind.
4. Deciding to shift your focus, you take C out to play fetch, hoping to breathe some fresh air and exercise.
5. You decide to take a short nap, thinking you might just be too tired, and that waking up might bring new ideas.
6. Deciding to shift your focus, you walk around the large house, checking to ensure all windows and doors are secure.

Q9:

Scenario: In a cozy family setting on a typical Thursday evening, you, B, and C, three members of the family, are discussing your weekend plans together. You really want to go see a movie, but B and C plan to do DIY cooking and play indoor video games, respectively. You hope to negotiate a consensus to improve everyone's satisfaction with the plan.

Action list:

1. You plan to create a schedule for the weekend to see if there's enough time for all activities.
2. You insist on going to see the movie and try to convince B and C by explaining how worthwhile the movie is.
3. For fairness, you suggest that everyone vote on what to do over the weekend, each person voting for one other activity besides their own.
4. To accommodate everyone's ideas, you start thinking about a suitable place that might allow for watching movies, cooking, and playing video games at the same time.
5. You suggest that everyone go to their preferred places individually, without the entire family having to stick together.

Q10:

Scenario: You are a professional triathlete, just finished a grueling morning training session, and are both tired and hungry. You are thinking about what to eat for a post-training meal while scrolling through your phone. You suddenly come across a news article about a burglary that occurred in your residential area last night, the fourth such incident this month. Your brother B is currently curled up on the sofa watching a live football match.

Action list:

1. You decide to grill some lean meat as your meal, considering the protein and nutrition it provides.
2. Just finished exercising, you decide to relax first, sitting on the sofa with B to enjoy the live broadcast of the football match and predict which team will win.
3. You ask B to wait a moment from watching TV, discussing spending some money on a more advanced security door and adding security cameras.
4. Suddenly feeling too tired to do anything, you decide to go to sleep right away to recover your strength.

Q11:

Scenario: You are a 30-year-old sports enthusiast, currently enjoying a weekend morning with your wife and kids at your comfortable suburban home. While eating brunch, you plan to start a harmonious day and then prepare for a camping trip, discussing the camping plans while eating.

Action list:

1. You encourage each family member to drink more water, focusing on their hydration.
2. You suggest that the family engage in light stretching activities after eating, to prepare for camping.
3. While discussing the details of the camping plans, you ensure every family member has a chance to speak.
4. You reminisce about the last time your family went camping, recalling fun and precious moments.
5. You delegate the preparation of camping supplies, discussing and deciding on the division of labor fairly with everyone.

D.4 Examples of OS-Sim

We demonstrate the unique value of Order-Sensitive Similarity (OS-Sim) through comparative examples contrasting its behavior with traditional rank correlation metrics.

D.4.1 Example 1: Positional Sensitivity

Consider two action sequences:

$$S = [5, 3, 1, 4, 2]$$

$$T = [3, 1, 5, 4, 2]$$

The OS-Sim calculation process reveals progressive alignment:

d	S_d	T_d	Jaccard(S_d, T_d)
1	{5}	{3}	$0/1 = 0$
2	{5, 3}	{3, 1}	$1/2 = 0.5$
3	{5, 3, 1}	{3, 1, 5}	$3/3 = 1$
4	{5, 3, 1, 4}	{3, 1, 5, 4}	$4/4 = 1$
5	{5, 3, 1, 4, 2}	{3, 1, 5, 4, 2}	$5/5 = 1$

$$\text{OS-Sim} = \frac{0 + 0.5 + 1 + 1 + 1}{5} = 0.7$$

Key Insight: Although there is full alignment at depths 3-5, the initial mismatches at depths 1-2 reduce the overall score a lot. This reflects real-world scenarios where the ranking of actions in the beginning holds more significance, influencing the outcome more than actions that come later.

D.4.2 Example 2: Differentiation Capacity

Consider three recommendation sequences:

$$A = [1, 2, 3, 4, 5] \quad (\text{Ground Truth})$$

$$B = [1, 2, 3, 5, 4] \quad (\text{Back-end Divergence})$$

$$C = [2, 1, 3, 4, 5] \quad (\text{Front-end Divergence})$$

Metric	A vs B	A vs C	Sensitivity
Spearman's ρ	0.90	0.90	×
Kendall's τ	0.80	0.80	×
OS-Sim	0.95	0.80	✓

Key Insight: While traditional rank correlations fail to distinguish between back-end (B) and front-end (C) divergences, OS-Sim successfully quantifies their operational differences through progressive prefix analysis:

- *Back-end Divergence:* Less penalty (0.95 vs 1.0) for swapping final elements
- *Front-end Divergence:* High penalty (0.80 vs 1.0) from initial misranking

This demonstrates OS-Sim's unique capacity to model real-world decision dynamics, where top-ranked actions have higher execution likelihood. Misprioritization of high-stakes actions incurs disproportionately severe consequences compared to lower-priority permutations—a fundamental operational distinction captured by OS-Sim but invisible to conventional sequence-agnostic metrics.

E Broader Impacts

The development of our value-driven decision-making framework presents a novel strategy in the creation of AI systems that are not only intelligent but also deeply aligned with human values and ethics. By embedding value dimensions directly into the decision-making process, our approach enhances the ability of AI to make socially responsible and contextually appropriate decisions. This has profound implications for applications in areas like healthcare, education, and personalized services, where the alignment of AI actions with ethical standards and individual preferences is critical. Moreover, our framework’s robust capabilities enable AI systems to adapt to a wide range of real-world scenarios, reducing the need for constant retraining and ensuring that these systems remain versatile and effective over time. This adaptability is particularly beneficial in enhancing user interactions, allowing AI to provide more personalized and satisfying experiences. Although potential risks exist, such as biases and privacy concerns associated with misuse of the technology, we believe that the careful design and implementation of our framework, coupled with ongoing vigilance, can mitigate these issues. Ultimately, the positive societal impacts of our work—such as improved decision-making, enhanced personalization, and ethical AI development—far outweigh the challenges, paving the way for AI systems that contribute meaningfully to human well-being and social progress.

F Limitations

While ValuePilot demonstrates a promising approach to personalized, value-driven AI, we acknowledge certain limitations inherent in our current methodology. The framework primarily emphasizes a novel engineering integration of existing techniques to enable explicit personalized value modeling, rather than on advancing the foundational algorithms for value learning or preference elicitation themselves. Consequently, the current value modeling, while effective for the presented demonstrations, offers a relatively straightforward representation of human values, which are inherently complex, hierarchical, and context-dependent. Furthermore, the method currently employed for eliciting user value preferences relies on direct rating scales. Although practical, this approach may fall short in capturing the nuanced and relative importance of individual values as effectively as more sophisticated psychometric techniques could.