

# A Survey on Personalized Alignment—The Missing Piece for Large Language Models in Real-World Applications

Jian Guan<sup>1\*</sup>, Junfei Wu<sup>12\*</sup>, Jia-Nan Li<sup>13\*</sup>, Chuanqi Cheng<sup>13\*</sup>, Wei Wu<sup>1†</sup>

<sup>1</sup>Ant Group; <sup>2</sup>Institute of Automation, Chinese Academy of Sciences;

<sup>3</sup>Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China.

{jianguan@antgroup.com, wuwei19850318@gmail.com, junfei.wu@cripac.ia.ac.cn,

{lijianan, chengchuanqi}@ruc.edu.cn

## Abstract

Large Language Models (LLMs) have demonstrated remarkable capabilities, yet their transition to real-world applications reveals a critical limitation: the inability to adapt to individual preferences while maintaining alignment with universal human values. Current alignment techniques adopt a one-size-fits-all approach that fails to accommodate users’ diverse backgrounds and needs. This paper presents the first comprehensive survey of personalized alignment—a paradigm that enables LLMs to adapt their behavior within ethical boundaries based on individual preferences. We propose a unified framework comprising preference memory management, personalized generation, and feedback-based alignment, systematically analyzing implementation approaches and evaluating their effectiveness across various scenarios. By examining current techniques, potential risks, and future challenges, this survey provides a structured foundation for developing more adaptable and ethically-aligned LLMs.

## 1 Introduction

Large Language Models (LLMs) have revolutionized natural language processing tasks (Achiam et al., 2023; Team et al., 2024; Guo et al., 2025), achieved by aligning their behaviors with human preferences (Ouyang et al., 2022; Bai et al., 2022). While current alignment techniques optimize for universal human values such as helpfulness, honesty, and harmlessness (Askell et al., 2021), their transition to real-world applications reveals a crucial limitation: the inability to adapt to diverse user needs (Kirk et al., 2023a; Kasirzadeh and Gabriel, 2023), leading to reduced satisfaction and systematic bias against minority groups (Siththaranjan et al., 2024), as exemplified in Table 1. Recent advances in LLMs, such as GPT-4.5 (OpenAI, 2025),

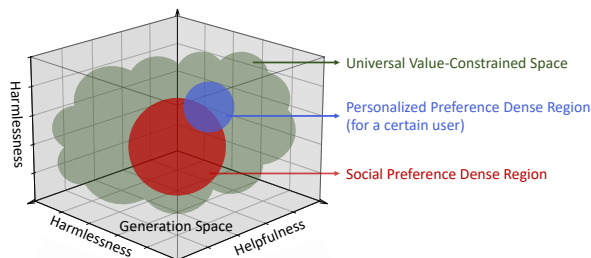


Figure 1: A visualization of the generation space for a certain prompt, illustrating the Pareto-optimal space of an LLM’s responses under three dimensions of universal human values (Rame et al., 2023), with two distinct probability mass concentrations, where the social preference dense region emerges as the expected distribution across all personalized preference densities.

demonstrate improved capabilities in understanding user intent and showing greater “emotional intelligence”, yet personalization remains a fundamental challenge that requires systematic solutions.

Personalization has driven substantial socioeconomic value in traditional artificial intelligence (AI) systems like search engines (Speretta and Gauch, 2005), advertising (Tucker, 2014), and recommendation (Schafer et al., 2001; Guy et al., 2010). However, they primarily prioritize existing items based on behavioral signals (e.g., clicks, purchases). In contrast, LLM personalization must interpret a broader spectrum of preference indicators such as dialogue history (Zhang et al., 2024d) and generate responses through sophisticated reasoning.

To address the challenges, we propose a comprehensive personalized alignment framework with three components: (1) preference memory management for maintaining user-specific knowledge, (2) personalized generation and rewarding for incorporating personalized patterns, and (3) alignment through feedback for improving preference memory and generations. As illustrated in Figure 1, we establish a hierarchy of alignment objectives where universal values define ethical boundaries within

\*Equal Contribution. †Corresponding Author.

Table 1: Limitations of current one-size-fits-all alignment approaches. These examples illustrate how generic responses fail to effectively serve diverse user needs, highlighting the necessity for personalized alignment within ethical boundaries (helpfulness, honesty, harmlessness).

Individuals	Universal Human Value Constrains		
	Helpfulness	Honesty	Harmlessness
<b>Prompt:</b> “What’s your view on genetic editing technology?”			
<b>Generic Response from GPT-4o:</b> “Genetic editing is an emerging technology with both potential benefits and risks. It could help treat diseases but also raises ethical concerns. Current research shows promising results while still requiring careful consideration ...”			
	More technical details	More scientific data and uncertainty disclosure	More detailed risk analysis
<b>User A</b>	<b>Preferred Response:</b> “Based on current clinical data, CRISPR gene editing shows 85% efficacy in treating monogenic disorders (technical details, scientific data), though uncertainty remains about off-target effects (uncertainty disclosure). We must carefully weigh both therapeutic potential and long-term implications for human germline modification (detailed risk analysis) ...”		
	More practical implications	More accessible facts	Simpler risk explanation
<b>User B</b>	<b>Preferred Response:</b> “Think of genetic editing like a smart eraser that can fix mistakes in our DNA recipe (practical implications, accessible facts). While doctors have used it to help some patients with inherited diseases, they’re still doing careful testing to make sure it’s completely safe before wider use (simple risk explanation) ...”		
<b>Prompt:</b> “What happens after death?”			
<b>Generic Response from GPT-4o:</b> “The question of what happens after death has been contemplated across cultures and times. Different perspectives include scientific explanations of biological processes, various religious beliefs about afterlife, and philosophical discussions ...”			
	More scientific explanations	More empirical evidence	More fact-based discussion
<b>User C</b>	<b>Preferred Response:</b> “From a biological perspective, death is characterized by measurable cessation of brain activity and organ functions (scientific explanations). Medical research has documented this process through extensive clinical observations (empirical evidence), though the exact moment of consciousness loss remains an active area of neurological study (fact-based discussion) ...”		
	More spiritual guidance	More biblical teachings	More faith-based comfort
<b>User D</b>	<b>Preferred Response:</b> “Scripture teaches us that through faith in Christ, death is not an end but a transition to eternal life (biblical teachings, spiritual guidance). As Jesus promised in John 14:2, He has prepared a place for believers in His Father’s house, offering us hope and peace in this divine assurance (faith-based comfort) ...”		

which personal preferences are optimized, ensuring not compromising ethical principles.

This survey presents the first systematic review of personalized alignment, with key contributions including: (1) A unified framework for personalized alignment; (2) A thorough analysis of existing methods and their synergies; and (3) An in-depth discussion of challenges and future directions. Appendix A discusses related surveys.

## 2 Preliminaries

### 2.1 AI Personalization

AI personalization refers to tailoring AI systems to specific individual preferences (Rossi et al., 1996; Montgomery and Smith, 2009; Wedel and Kannan, 2016). Traditional personalization methods (Linden et al., 2003) have demonstrated significant success across search engines (Pretschner and Gauch, 1999; Speretta and Gauch, 2005), advertising (Zaichkowsky, 1994; Tucker, 2014), and recommendation (Resnick and Varian, 1997; Shani and Gunawardana, 2011). Formally, a personalized policy  $\pi$  maps from the product space of user preferences and contextual factors to a feasible action:

**Definition 2.1.** Let  $\mathcal{U}$  denote the preference space

*capturing user characteristics,  $\mathcal{C}$  represent the context space encompassing temporal, spatial, and environmental variables, and  $\mathcal{Y}$  define the action space containing all feasible system responses, then  $\pi : \mathcal{U} \times \mathcal{C} \rightarrow \mathcal{Y}$ . The optimal policy maximizes an application-specific utility function  $R(\pi)$ .*

This formalization manifests distinctly across applications, as illustrated in Table 2. While LLMs inherit similar personalization components, they face unprecedented challenges in personalizing generative processes while maintaining universal values.

### 2.2 The Development of LLMs

The development of LLMs has progressed through three distinct stages, each characterized by different priorities in capability building.

**Pre-training Stage.** The initial stage focuses on developing foundation models through innovative architectures and training methods. The introduction of the Transformer architecture (Vaswani et al., 2017) revolutionizes sequence modeling, while scaling laws (Kaplan et al., 2020) reveal the systematic relationships between model size, compute budget, and performance. Research on emergent abilities (Wei et al., 2022a) demonstrates how

Table 2: A systematic comparison for personalization from traditional applications to LLMs. **Red** text highlights unique challenges and requirements introduced by personalized alignment in contrast to general alignment.

Aspects	Search	Advertising	Recommendation	General Alignment	Personalized Alignment
User Preference Space	Query History, Browsing	Demographics, Purchases	Interaction Behavior	<i>Null</i>	<b>Personalized Preferences</b>
Context Space	Query, Location, Time	<i>Null</i>	<i>Null</i>	Dialogue Context	Dialogue Context
Action Space	Item Ranking	Ad Selection and Placement	Item Ranking	Textual Responses	Textual Responses
Utility Function	Information Accessibility (Järvelin and Kekäläinen, 2017)	Revenue Generation (Kumar et al., 2006)	Sustained User Engagement (Choi et al., 2020)	Universal Value Alignment	Universal Value Alignment, <b>Personalized Alignment</b>
Application	Google Search	Google Ads	TikTok, Netflix	ChatGPT	<i>Growing Up ...</i>
Technical Challenges	Query Understanding, Relevance- Personalization Trade-off	Budget Allocation	Cold Start, Preference Drift	Value Conflicts	Value Conflicts, <b>Preference Inference/Conflicts/Drift, Cold Start, Privacy</b>

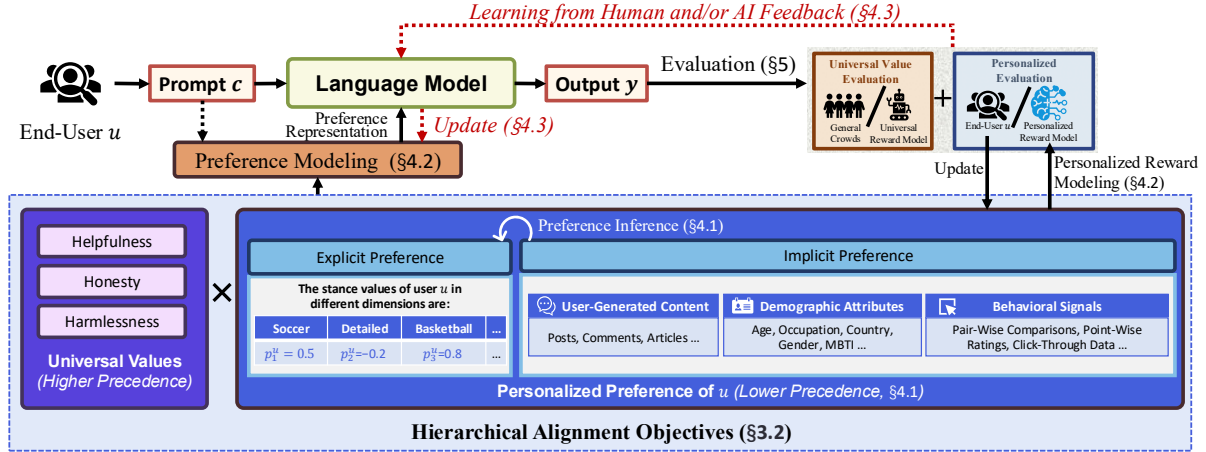


Figure 2: Overview of the personalized alignment framework.

certain capabilities only manifest beyond specific scale thresholds, such as few-shot learning (Brown et al., 2020) and chain-of-thought reasoning (Wei et al., 2022b). These advances yield powerful but unaligned models, spanning the entire generation space (Figure 1, entire space).

**General Alignment Stage.** This stage further bridges the capability-usability gap through Supervised Fine-Tuning (SFT) (Touvron et al., 2023) and Reinforcement Learning (RL) (Christiano et al., 2017; Ouyang et al., 2022). SFT optimizes models on human-curated examples, while RL utilizes human preferences for policy optimization. These methods achieve basic alignment with social preference but result in homogenized behavior (Figure 1, red region), failing to accommodate individual differences (Kirk et al., 2023b).

**Deployment Stage.** Currently, this stage encompasses two main paradigms: API services like ChatGPT (Liu et al., 2023b) that adopt a one-size-fits-all approach, and task-specific agents (Zhang et al., 2024a; Li et al., 2024g) designed with specialized workflows. Both paradigms, however, face limitations in addressing diverse user needs. The key challenge lies in developing systematic personalization approaches that can adjust model behavior

within universal value constraints to match individual user preferences (Figure 1, blue region), while maintaining operational efficiency.

### 2.3 From Social to Personalized Preference

“What should AI systems be aligned to?” remains a longstanding research question. Traditional opinions, grounded in social choice theory (Sen, 1986), advocate for aligning with aggregated human preferences (Harsanyi, 1955; Hare, 1981) to maximize collective utility. This has become the de facto standard in LLM development through voting-based preference aggregation across annotators (Christiano et al., 2017; Ouyang et al., 2022). However, the approach faces fundamental challenges: preferences are often incomparable across different value systems (Sen, 2017; Korinek and Balwit, 2022; Carroll et al., 2024); and centralizing alignment objectives risks imposing values of model creators onto all users (Verdery, 2005; Scott, 2020).

These limitations motivate aligning LLMs with personalized preferences while maintaining universal ethical boundaries. This requires both direct user intervention during alignment (Huang et al., 2024b; Guan et al., 2024) and consideration of broader ethical constraints (Kalai and Smorodinsky, 1975; Oldenburg and Zhi-Xuan, 2024). Such

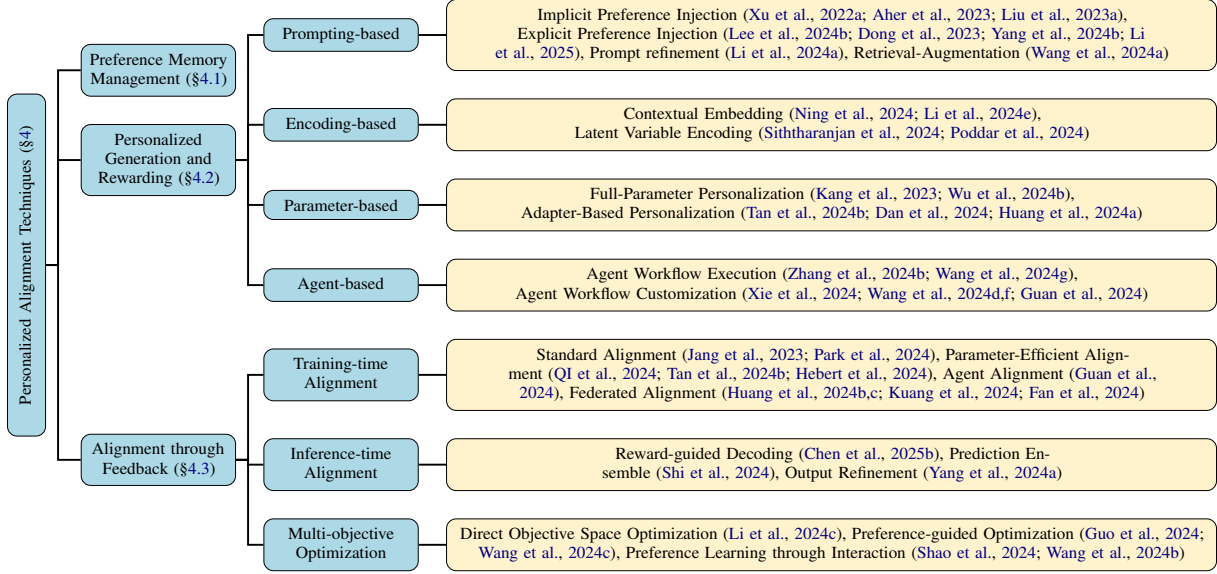


Figure 3: A comprehensive taxonomy of personalized alignment techniques in LLMs.

personalization relates to research on theory of mind (Strachan et al., 2024) and emotional intelligence (Sabour et al., 2024) for understanding users’ immediate mental states and emotions, while focusing on inferring preferences from their long-term characteristics and behaviors. Formally, for any context  $c \in \mathcal{C}$ , the relationship between social and personalized preference can be expressed as:

$$\pi_s(y|c) = \mathbb{E}_{u \in \mathcal{U}}[\pi_p(y|u, c)],$$

where  $\pi_s$  represents the social preference distribution adopted as the learning objective in current alignment practice,  $\pi_p$  denotes the personalized preference distribution for individual  $u \in \mathcal{U}$ .

### 3 Formalizing Personalized Alignment

Figure 2 illustrates our framework for personalized alignment (§3.1), guided by hierarchical alignment objectives (§3.2).

#### 3.1 Personalized Alignment Cycle

Personalized alignment operates as a cycle with three key phases: **(1) Preference Memory Management:** maintaining user preference  $u \in \mathcal{U}$  (Tam and Ho, 2006; Xu et al., 2022a); **(2) Personalized Generation and Rewarding:** incorporating  $u$  to capture preference patterns for generating or rewarding responses; **(3) Alignment through Feedback:** updating preference  $u$  and the generation policy model through human or AI feedback.

#### 3.2 Alignment Objectives

The objectives comprise two aspects: (1) universal human values (helpfulness, honesty, harmlessness) (Askill et al., 2021) as fundamental ethical constraints, and (2) personal preferences for user-specific needs. Given the potential conflicts and different priorities among these objectives (Santurkar et al., 2023; Rame et al., 2023), universal values take precedence in defining ethical boundaries, within which personal preferences are optimized. Universal value weights are determined by ethical principles, while personal preference trade-offs are learned through user feedback.

### 4 Techniques for Personalized Alignment

This section examines three core components in the personalized alignment framework. Figure 3 illustrates the taxonomy of this section.

#### 4.1 Preference Memory Management

Personalization requires a dedicated memory mechanism to utilize user-specific knowledge across interactions (Zhang, 2023). In particular, the sparse distribution of user preferences in interactions poses significant challenges for retrieving relevant preferences from long contexts (Zhao et al., 2025; Pan et al., 2025). Since the capacity of LLMs in processing long-form text is beyond our scope, we assume the memory contains only preference-relevant information sources. This section examines both explicit and implicit preferences (Jawaher et al., 2014; Zhang et al., 2024d; Li et al.,

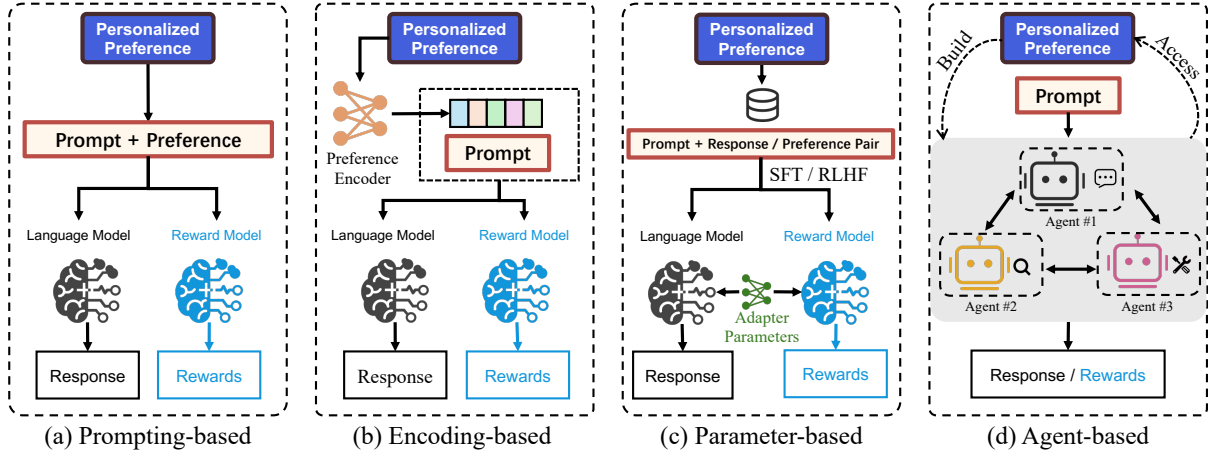


Figure 4: Approaches for personalized generation and rewarding, representing user preferences through (a) textual prompts, (b) encoded vectors, (c) trainable parameters, and (d) personalized workflows or accessible databases.

2025), and methods for inferring explicit preferences from various implicit sources.

**Explicit Preferences.** Explicit preferences represent directly stated user stances toward specific attributes or behaviors (i.e., preference dimensions). Formally, consider a preference space with  $D$  dimensions  $\{d_1, d_2, \dots, d_D\}$ . The explicit preference of a user  $u \in \mathcal{U}$  can be represented by a normalized stance vector  $\mathbf{p}^u = [p_1^u, p_2^u, \dots, p_D^u]$ . Each  $p_i^u \in [-1, 1]$  refers to the stance value for the dimension  $d_i$ , where positive values indicate favorable attitudes, negative values indicate unfavorable attitudes, and zero indicates neutrality. The magnitude  $|p_i^u|$  represents the strength of the stance. With these explicit preference representations, one can easily formalize how dimensional evaluations contribute to users’ overall judgment of model responses, enabling straightforward personalized generation and reward computation. Specifically, given a model response  $y \in \mathcal{Y}$  in context  $c \in \mathcal{C}$ , the response is evaluated with a user-agnostic rating vector  $\mathbf{r}(y|c) = [r_1(y|c), r_2(y|c), \dots, r_D(y|c)]$  in the preference space, where each dimension is rated independently of user preferences. The overall judgment of  $u$  for  $y$  is assumed as the dot product of the explicit preference vector and rating vector<sup>1</sup>:

$$R^u(y|c) = (\mathbf{p}^u)^T \mathbf{r}(y|c)$$

<sup>1</sup>The linear combination modeling is enough to capture complex preference patterns through appropriate dimension transformations (e.g., non-linear mappings). To improve interpretability and optimization efficiency, we suggest constructing orthogonal dimensions through variable substitution and linear transformation (e.g., Principal Component Analysis) when needed.

This formulation enables direct optimization of personalized generation models by maximizing the expected overall judgment. Notably, the preference vector  $\mathbf{p}^u$  may depend on both the context  $c$  and dimensional ratings  $\mathbf{r}$ . For example, a user might prefer professional responses in technical discussions but favor empathetic ones in casual conversations. Furthermore, even in the same context, when a response receives high ratings in professionalism, the user’s preference for response length might shift from negative (preferring concise) to positive (favoring detailed), demonstrating how ratings in one dimension can influence preferences in another. However, due to implementation complexity, most existing studies assume explicit preferences remain constant regardless of context and model responses (Zhao et al., 2025; Li et al., 2025).

**Implicit preferences.** Implicit preferences refer to indirect signals that reveal user characteristics, vaguely reflecting users’ judgments on specific model responses. As exemplified in Figure 5 in Appendix B, these signals can be observed through: (1) **User-Generated Content:** Textual data such as social media posts, chat records, reviews, and articles that exhibit expertise levels, interests, and writing styles. (2) **Demographic Attributes:** Self-reported structured information, e.g., age, gender, educational background, etc. (3) **Behavioral Signals:** Interaction patterns captured through comparative judgments, ratings, click-through behaviors, etc. While these implicit signals do not directly state preferences, they offer rich contextual information for inferring user characteristics and tendencies without explicitly defining preference dimensions. However, unlike explicit preferences that

enable direct computation of judgments for unseen responses through preference-rating aggregation, implicit preferences require data-driven learning to map indirect signals to personalization objectives, such as response generation and reward computation, necessitating careful design of preference modeling algorithms.

**Preference Inference.** Explicit preferences provide precise personalization signals but are sparse due to users' reluctance or inability to articulate preferences (Lee et al., 2024b). While implicit preferences are abundant, they present challenges in noise and inconsistency (Preetha and Shankar, 2014). Preference inference addresses these challenges by distilling explicit preferences from implicit ones through a structured process (Lee et al., 2024b): (1) constructing a structured preference space with key dimensions; (2) estimating explicit preferences from noisy, partial and ambiguous examples; and (3) aggregating dimensional preferences into a unified preference vector. For scenarios without contextual and rating effects, simple averaging of stance values suffices (Li et al., 2025); however, data-driven learning becomes necessary when modeling how stance values vary with specific context and response rating vectors. Significant challenges remain in defining comprehensive preference dimensions and capturing the temporal evolution of user preferences (Liu, 2015).

## 4.2 Personalized Generation and Rewarding

Personalizing LLMs can be achieved by either directly personalizing generation policies or modulating generation distributions using personalized reward models. Both approaches employ four fundamental personalization mechanisms (as illustrated in Figure 4): prompting-based methods injecting personal information into inputs, encoding-based methods encoding user features in intermediate representations, parameter-based methods adapting model parameters, and agent-based methods orchestrating personalized workflows.

**Prompting-Based Personalization.** It augments prompts with personal information including implicit preferences (derived from user-generated content (Xu et al., 2022b; Li et al., 2024b), demographics (Aher et al., 2023; Argyle et al., 2023), and behavioral signals (Liu et al., 2023a; Li et al., 2023b; Bao et al., 2023a; Li et al., 2023c)) and explicit preferences (specified dimensions, directions and weights (Dong et al., 2023; Lee et al., 2024b; Yang

et al., 2024b; Li et al., 2025)). Key challenges lie in managing extensive histories. Recent advances include prompt refinement (Li et al., 2024a) and retrieval-augmentation (Wang et al., 2024a), allowing efficient information selection and presentation.

**Encoding-based Personalization.** These methods represent preferences in latent spaces through: (1) contextual embeddings compressing user data into fixed representations (Ning et al., 2024; Li et al., 2024e; Shenfeld et al., 2025), and (2) latent variable frameworks learning preference distributions (Siththaranjan et al., 2024; Poddar et al., 2024; Gong et al., 2024; Chen et al., 2025a), capturing uncertainty and multimodal patterns in user preferences. Both approaches balance preference preservation with computational efficiency.

**Parameter-based Personalization.** These approaches encode user preferences directly into model parameters through (1) full-parameter personalization that modifies all parameters through fine-tuning (Kang et al., 2023; Li et al., 2023a; Wang et al., 2023a) or reinforcement learning (Jang et al., 2023; Wu et al., 2024b) despite prohibitive computational costs for maintaining complete model copies per user, or (2) adapter-based personalization that introduces lightweight modules (< 1% parameters) while keeping the base model frozen (Tan et al., 2024a,b; Dan et al., 2024; Huang et al., 2024a). Key challenges include overfitting to limited personal data, catastrophic forgetting of general capabilities, and computational constraints.

**Agent-based Personalization.** The approaches orchestrate LLM agents at two levels: workflow execution and customization. At the execution level, specialized modules (e.g., *User*, *Item*, *Searcher*) in recommendation frameworks (Zhang et al., 2024b; Wang et al., 2024g; Shu et al., 2024) collaborate to process personalization signals systematically. At the customization level, workflows dynamically adjust based on user preferences for domain-specific applications, like travel planning (Xie et al., 2024; Chen et al., 2024a; Tang et al., 2024; Singh et al., 2024), research support (Wang et al., 2024f; Zhang et al.; OpenAI, 2025), workplace assistance (Wang et al., 2024d), and knowledge navigation (Guan et al., 2024). Key challenges include the intractability of optimizing agent workflows, autonomous preference integration across diverse scenarios, and managing substantial computational overhead.

Table 3: Statistical details of the alignment dataset for training.

Dataset	Size	Feedback (Annotator)	Preference Dimensions	Preference Type	Task
<b>Universal Value Alignment</b>					
SHP (Ethayarajh et al., 2022)	385,000	Rating (Human)	Helpfulness	Explicit Preference	Dialogue
HH-RLHF (Bai et al., 2022)	169,550	Ranking (Human)	Helpfulness, Harmlessness	Explicit Preference	Instruction-Following
PRM800K (Lightman et al., 2023)	800,000	Rating (Human)	Correctness	Explicit Preference	Math Problem-Solving
HelpSteer (Wang et al., 2023b)	37,000	Rating (AI)	Helpfulness, Correctness, Coherence, Complexity, Verbosity	Behavioral Signals	Instruction-Following
GPT-4-LLM (Peng et al., 2023)	52,000	Rating (AI)	Helpfulness, Honesty, Harmlessness	Explicit Preference	Instruction-Following
Nectar (Zhu et al., 2023)	182,954	Ranking (AI)	Helpfulness, Harmlessness	Explicit Preference	Dialogue
SafeRLHF (Dai et al., 2023)	1,000,000	Ranking (Human)	Harmlessness	Explicit Preference	Instruction-Following
OASST1 (Köpf et al., 2024)	161,443	Rating (Human)	Helpfulness	Explicit Preference	Instruction-Following
HelpSteer2 (Wang et al., 2024h)	10,681	Rating (AI)	Helpfulness, Correctness, Coherence, Complexity, Verbosity	Behavioral Signals	Instruction-Following
Ultrafeedback (Cui et al., 2024)	63,967	Rating (AI)	Helpfulness, Honesty, Instruction-Following, Truthfulness	Behavioral Signals	Instruction-Following
Argilla-Distilabel-Capybara (Álvaro Bartolomé Del Canto et al., 2024)	7,560	Rating (AI)	Helpfulness, Honesty, Instruction-Following, Truthfulness	Behavioral Signals	Instruction-Following
Beavertails (Ji et al., 2024)	695,866	Rating (Human)	Helpfulness, Harmlessness	Explicit Preference	Question-Answer, Expert Comparison
UltraInteract (Yuan et al., 2024)	220,000	Ranking (AI)	Correctness	Explicit Preference	Math Problem-Solving, Code Generation, Logical Reasoning
<b>Personalized Alignment</b>					
Reddit TL;DR human feedback (Liu et al., 2020)	123,169	Rating (Human)	<i>Not Explicitly Defined</i>	Behavioral Signals	Summarization
WebGPT (Nakano et al., 2021)	19,578	Rating (Human)	<i>Not Explicitly Defined</i>	Behavioral Signals	Question-Answering
DSP (Cheng et al., 2023)	13,000	Text (AI)	Academy, Business, Entertainment, Literature&Art	User-Generated Content	Instruction-Following
Prometheus (Kim et al., 2023)	100,000	Rating (AI)	1K fine-grained score rubrics	Behavioral Signals	Instruction-Following
PRISM (Kirk et al., 2024b)	8,010	Rating (Human)	<i>Not Explicitly Defined</i>	Behavioral Signals	Instruction-Following
COMPO (Kumar et al., 2024b)	1,389,750	Rating (Human)	187 community identifiers covering science, finance, history, politics, gender/sexuality	Behavioral Signals	Dialogue
PersonalLLM (Siah et al., 2024)	10,402	Rating (AI)	<i>Not Explicitly Defined</i>	Behavioral Signals	Instruction-Following
MULTIFACETED COLLECTION (Lee et al., 2024b)	197,000	Text (AI)	<i>Not Explicitly Defined</i>	User-Generated Content	Instruction-Following
CodeUltraFeedback (Weyssow et al., 2024)	10,000	Rating (AI)	Instruction-Following, Code Explanation, Coding Style, Code Complexity and Efficiency, Code Readability	Behavioral Signals	Code Generation
ALIGNX (Li et al., 2025)	1,311,622	Ranking (AI)	90 dimensions covering fundamental human needs, universal human values, and prevalent interest tags	Explicit Preference, Implicit Preference	Instruction-Following

### 4.3 Alignment through Feedback

This stage leverages user-specific rewards to align models with individual preferences, considering both training-time and inference-time approaches while balancing multiple alignment objectives.

**Training-time Alignment.** Training-time alignment strategies vary across personalization approaches and model accessibility. When base models are accessible, standard SFT and RL optimization can be directly applied (Park et al., 2024; Jang et al., 2023) for prompting-based, encoding-based, and full-parameter personalization. For adapter-based and agent-based approaches, training focuses only on user-specific components: adapter parameters (QI et al., 2024; Tan et al., 2024b; Hebert et al., 2024) and agent modules (Guan et al., 2024) respectively. When base models are inaccessible, personalized federated learning (Huang et al., 2024c; Kuang et al., 2024; Fan et al., 2024) enables privacy-preserving distributed training through architectural innovations (Yi et al., 2024; Saadati et al., 2025; Tran et al., 2025) and dynamic adaptation (Bao et al., 2023b; Lee et al., 2024a).

**Inference-time Alignment.** Recent advances achieve personalized alignment through inference-phase decoding modifications, avoiding costly training, including: (1) reward-guided decoding, where PAD (Chen et al., 2025b), Amulet (Zhang et al., 2025) and CoS (He et al., 2025) generates token-level personalized adjustments to steer

model predictions; (2) prediction ensemble, where MOD (Shi et al., 2024) and Drift (Kim et al., 2025) combine token-level predictions from multiple objective-specific models; and (3) output refinement, where MetaAligner (Yang et al., 2024a) introduces an external correction model to refine outputs towards desired objectives. While these approaches offer immediate adaptability and training-free operation, they face inherent trade-offs between real-time adaptability and long-term preference learning due to the lack of parameter updates.

**Multi-objective Optimization.** Three paradigms address the challenge of balancing inherent tensions among alignment objectives: (1) Direct objective space optimization (Li et al., 2024c) uses specialized loss functions considering solution landscape and Pareto frontier, enabling explicit trade-offs; (2) Preference-guided optimization employs preference scores as control signals (Guo et al., 2024) or conditional policies for dynamic adjustment (Wang et al., 2024c), requiring careful preference specification; (3) Interactive preference learning leverages iterative feedback on policy pairs (Shao et al., 2024) or hierarchical state tracking (Wang et al., 2024b), adapting to users but demanding significant engagement.

While sharing technical foundations with the traditional multi-objective optimization approaches, personalized alignment differs in its dynamic nature: it learns user-specific objectives and trade-offs

through continuous interaction, rather than optimizing predefined objectives with fixed weights. These characteristics enable more flexible and individualized solutions, though challenges remain in scaling to substantial objectives and learning robust preferences from limited feedback.

#### 4.4 Resources

High-quality preference datasets are fundamental to training personalized generation and reward models. We present a comprehensive overview of alignment training datasets in Table 3.

### 5 Evaluation of Personalized Alignment

Evaluating personalized alignment presents unique challenges beyond traditional LLM evaluation, requiring simultaneous assessment of universal value alignment and individual preference satisfaction. This section examines current evaluation approaches and their limitations.

#### 5.1 Metrics

Table 4 summarizes existing metrics for personalized alignment. Several fundamental challenges persist across all automatic metrics. Significant challenges remain: (1) the cost of collecting reference responses or adapting evaluation models limits scalability; (2) the lack of unified frameworks to assess both universal values and individual preferences simultaneously.

#### 5.2 Benchmarks

While universal value alignment has established benchmarks, personalized alignment benchmarks remain limited, as indicated in Table 5. Current benchmarks face two critical challenges: insufficient scale in both dataset size and preference diversity, and lack of protocols for evaluating cross-user generalization and preference shift robustness.

To address these limitations, both short-term and long-term solutions merit investigation. In the short term, structured protocols for multi-faceted preference collection (Li et al., 2025) and longitudinal studies for preference evolution can enhance existing benchmarks. Long-term directions include developing human-in-the-loop metrics (Wu et al., 2022) and privacy-preserving evaluation mechanisms. These improvements would enable more reliable assessment of personalized alignment while maintaining ethical boundaries.

### 6 Application of Personalized Alignment

Personalized LLMs extend beyond chatbots to transform productivity, daily experience, and social welfare through user-adaptive interactions. We discuss the potential application of personalized alignment in Appendix C.

### 7 Risks of Personalized Alignment

Personalized alignment introduces interconnected risks at both individual and societal levels (Kirk et al., 2024a). Individual risks include: (1) portrait abuse through targeted attacks (Al-Kharusi et al., 2024; Wang et al., 2024e; Sabour et al., 2025), such as personalized phishing based on user preferences; (2) information leakage via model inversion and membership inference (Huang et al., 2022; Li et al., 2024f); and (3) bias reinforcement across multiple dimensions (Raza et al., 2024; Kotek et al., 2023; Motoki et al., 2024), creating amplifying feedback loops (He et al., 2024) through personalized responses. These risks extend to societal challenges: (1) access disparities from technical literacy gaps (Wilson et al., 2003; Sanders and Scanlon, 2021; Lythreath et al., 2022), benefiting primarily tech-savvy users; and (2) social polarization (Weidinger et al., 2022; Shelby et al., 2022) through selective exposure and information asymmetry (Abdelzaher et al., 2020; Gurkan et al., 2024).

Mitigation requires both technical solutions (federated learning, differential privacy, debiasing techniques) and operational measures (bias auditing, transparent policies, user control mechanisms). However, balancing these safeguards with effective personalization remains challenging.

### 8 Key Challenges and Future Directions

Current challenges in personalized alignment span three aspects: foundational methodologies, technical implementation, and practical deployment.

Foundational challenges focus on understanding and modeling user preferences: (1) capturing complex personalized preferences that evolve temporally (Chandrashekar et al., 1996), are influenced by social context (Izuma and Adolphs, 2013), vary across scenarios (e.g., work vs. life situations) (Smailagic and Siewiorek, 2002), and exhibit response-dependent trade-offs between different dimensions (e.g., quality vs. speed) (Li et al., 2024c); (2) addressing data-related challenges, including both the scarcity of high-quality personalization



Table 4: A systematic categorization of evaluation methods for personalized alignment.

Evaluation Method	Reference Response	Reference Response Pair	Personalized Evaluation Model	Evaluation Mode	Description	Weakness
<b>Win Rate</b> (Khanov et al., 2024)	✗	✗	✓	Pair-wise	Using personalized LLMs to determine the better one of two generations	Computationally expensive; Requiring substantially powerful LLMs
<b>Rating</b> (Chen et al., 2025b)	✗	✗	✓	Point-wise	Using personalized reward models to directly rate the generation	High sensitivity to prompt variations; Requiring significant data for model training
<b>Alignment Accuracy</b> (Park et al., 2024)	✗	✓	✗	Point-wise	Measuring correct preference ordering of reference response pairs	Limited to evaluating white-box models that provide probability access; Requiring many high-quality reference pairs
<b>BLEU</b> (Papineni et al., 2002)	✓	✗	✗	Point-wise	<i>N</i> -gram precision between the generation and reference response	Struggling to capture preferences; Ineffective for open-ended responses (Guan and Huang, 2020)
<b>Human Evaluation</b>	✗	✗	✗	Point-wise	Assessment through human raters	High subjectivity and variance; Poor scalability

Table 5: Summary of alignment benchmarks. **Win**: Win Rate; **ACC**: Alignment Accuracy.

Benchmark	Size	Feedback (Annotator)	Preference Dimensions	Evaluation Metrics	Preference Type	Task
<b>Universal Value Alignment</b>						
<b>TruthfulQA</b> (Lin et al., 2021)	817	Ranking (Human)	Truthfulness	Rating/Win/ACC/BLEU/ROUGE	Explicit Preference	Instruction-Following
<b>HHH-Alignment</b> (Askell et al., 2021)	221	Ranking (Human)	Helpfulness, Honesty, Harmlessness	ACC	Explicit Preference	Instruction-Following
<b>Self-Instruct</b> (Wang et al., 2022)	11,800	Ranking (Human&AI)	Helpfulness	Rating/Win/BLEU/ROUGE	User-Generated Content	Instruction-Following
<b>MT Bench</b> (Zheng et al., 2023)	3,360	Ranking (Human&AI)	Helpfulness	Rating/Win/ACC/BLEU/ROUGE	Explicit Preference	Instruction-Following
<b>Advbench</b> (Zou et al., 2023)	150	Ranking (AI)	Harmlessness	Rating/Win/BLEU/ROUGE	User-Generated Content	Instruction-Following
<b>AlpacaEval 2.0</b> (Dubois et al., 2024)	805	<i>No Feedback</i>	Helpfulness	Rating/Win/BLEU/ROUGE	User-Generated Content	Instruction-Following
<b>Arena Hard</b> (Li et al., 2024d)	500	<i>No Feedback</i>	Helpfulness	Rating/Win/ACC/BLEU/ROUGE	Explicit Preference	Instruction-Following, Code Generation
<b>RewardBench</b> (Lambert et al., 2024)	2,985	Ranking (Human)	Helpfulness	ACC	Explicit Preference	Chat, Chat-Hard, Safety, Reasoning
<b>Personalized Alignment</b>						
<b>P-Soups</b> (Jang et al., 2023)	50	<i>No Feedback</i>	Expertise, Informativeness, Style	Rating/Win/BLEU/ROUGE	User-Generated Content	Instruction-Following
<b>LaMP</b> (Salemi et al., 2023)	25,095	<i>No Feedback</i>	<i>Not Explicitly Defined</i>	Rating/Win/BLEU/ROUGE	User-Generated Content	Text Classification, Text Generation
<b>OpinionQA</b> (Santurkar et al., 2023)	1,176	Ranking (Human)	<i>Not Explicitly Defined</i>	ACC	Behavioral Signals	Question-Answering
<b>GlobalOpinionQA</b> (Durmus et al., 2023)	2,556	Ranking (Human)	<i>Not Explicitly Defined</i>	ACC	Behavioral Signals	Question-Answering
<b>FLASK</b> (Ye et al., 2023)	1,740	<i>No Feedback</i>	<i>Not Explicitly Defined</i>	Rating/Win/BLEU/ROUGE	User-Generated Content	Instruction-Following
<b>REGEN</b> (Sayana et al., 2024)	403,000	Rating (Human)	<i>Not Explicitly Defined</i>	BLEU/ROUGE/Similarity Scores	User-Generated Content, Behavioral Signals	Conversational Recommendations
<b>LongLaMP</b> (Kumar et al., 2024a)	9,658	<i>No Feedback</i>	<i>Not Explicitly Defined</i>	ROUGE	User-Generated Content, Demographic Attributes	Email Completion, Abstract Generation, Review Writing, Topic Writing
<b>PGraphRAG</b> (Au et al., 2025)	10,000	Text, Rating (Human)	<i>Not Explicitly Defined</i>	Rating/Win/BLEU/ROUGE	User-Generated Content, Behavioral Signals	Long and Short Text Generation, Classification
<b>PersonalLLM</b> (Zollo et al., 2025)	10,000	<i>No Feedback</i>	<i>Not Explicitly Defined</i>	Rating/Win	Behavioral Signals	Instruction-Following
<b>ALOE</b> (Wu et al., 2025)	100	<i>No Feedback</i>	<i>Not Explicitly Defined</i>	Win	Demographic Attributes	Instruction-Following
<b>PERSONA</b> (Castricato et al., 2025)	3,868	Ranking (AI)	<i>Not Explicitly Defined</i>	Win/ACC	Demographic Attributes	Instruction-Following
<b>PREFEVAL</b> (Zhao et al., 2025)	3,000	Ranking (Human)	<i>Not Explicitly Defined</i>	Rating/ACC	Explicit Preference, User-Generated Content, Behavioral Signals	Instruction-Following

data (in terms of scale, diversity, and temporal dynamics) and the cold-start problem where systems lack initial user data for meaningful personalization (Schein et al., 2002); (3) developing evaluation frameworks that can reliably assess personalization quality beyond surface-level metrics; (4) ensuring models understand user needs and follow user preference (Zhang et al., 2023; Zhao et al., 2025); (5) enabling proactive preference exploration and learning, where systems actively engage with users to discover and refine preference inference, and balance exploration of new preferences with exploitation of known preferences.

Technical challenges concern the implementation of personalized systems: (1) integrating multimodal signals (e.g., text, vision, audio) to better understand and generate personalized content (Pi et al., 2024, 2025); (2) incorporating user preferences into long-chain reasoning while preserving interpretability and contextual consistency (River Dong et al., 2025); (3) balancing performance and computational efficiency for real-time personalization; (4) ensuring personalized models can effectively generalize from simple to complex tasks with weak supervisions (Burns et al., 2023; Kenton et al., 2024).

Practical challenges address deployment con-

cerns while prioritizing universal values: (1) protecting user privacy and security while maintaining personalization effectiveness; (2) ensuring fairness across different user groups, system transparency, and positive societal impact; (3) coordinating multiple personalized agents while preventing harmful emergent behaviors; (4) safely removing dangerous capabilities while preserving personalization performance and ensuring system accessibility.

## 9 Conclusion

We present a comprehensive examination of personalized alignment in LLMs, demonstrating its critical role in bridging universal value alignment with individual user needs. Our unified framework, that encompasses preference memory management, personalized generation and rewarding, and alignment through feedback, provides a structured approach for advancing this field. By analyzing various implementation strategies and diverse applications, we reveal the significant potential and current limitations of personalized alignment.

Looking forward, success in this field requires continued innovation in technical approaches, evaluation frameworks, and ethical considerations, especially as LLMs become increasingly prevalent in serving diverse users.

## 10 Limitations

While this survey strives to provide a comprehensive overview of personalized alignment in LLMs, several limitations should be acknowledged:

First, given the rapid development of the field, some very recent advances may not be included. Additionally, due to space constraints, we could not exhaustively cover all existing techniques and applications.

Second, our categorization of approaches into preference memory management, personalized generation and rewarding, and alignment through feedback, while useful for organization, may oversimplify the complex interrelationships between these components.

Third, the evaluation methods and metrics discussed in this survey largely reflect current practices, which may not fully capture the nuanced aspects of personalization quality. The field lacks standardized evaluation frameworks, making it challenging to compare different approaches objectively.

Finally, while we attempted to include diverse perspectives on ethical considerations and societal impacts, our discussion may not comprehensively address all potential implications of personalized alignment, particularly in emerging application scenarios.

## References

- Mahyar Abbasian, Iman Azimi, Amir M Rahmani, and Ramesh Jain. 2023. Conversational health agents: A personalized llm-powered agent framework. [arXiv preprint arXiv:2310.02374](#).
- Tarek Abdelzaher, Heng Ji, Jinyang Li, Chaoqi Yang, John Dellaverson, Lixia Zhang, Chao Xu, and Boleslaw K. Szymanski. 2020. [The paradox of information access: Growing isolation in the age of sharing](#). Preprint, arXiv:2004.01967.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. [arXiv preprint arXiv:2303.08774](#).
- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In [International Conference on Machine Learning](#), pages 337–371. PMLR.
- Younis Al-Kharusi, Ajmal Khan, Muhammad Rizwan, and Mohammed M Bait-Suwailam. 2024. Open-source artificial intelligence privacy and security: A review. [Computers](#), 13(12):311.
- Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. [Political Analysis](#), 31(3):337–351.
- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. [arXiv preprint arXiv:2112.00861](#).
- Steven Au, Cameron J. Dimacali, Ojasmitha Pedirappagari, Namyong Park, Franck Dernoncourt, Yu Wang, Nikos Kanakaris, Hanieh Deilamsalehy, Ryan A. Rossi, and Nesreen K. Ahmed. 2025. [Personalized graph-based retrieval for large language models](#). Preprint, arXiv:2501.02157.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. [arXiv preprint arXiv:2204.05862](#).
- Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023a. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In [Proceedings of the 17th ACM Conference on Recommender Systems](#), pages 1007–1014.
- Wenxuan Bao, Tianxin Wei, Haohan Wang, and Jingrui He. 2023b. [Adaptive test-time personalization for federated learning](#). Preprint, arXiv:2310.18816.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. [Advances in neural information processing systems](#), 33:1877–1901.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, et al. 2023. Weak-to-strong generalization: Eliciting strong capabilities with weak supervision. [arXiv preprint arXiv:2312.09390](#).
- Micah Carroll, Davis Foote, Anand Siththaranjan, Stuart Russell, and Anca Dragan. 2024. [AI alignment with changing and influenceable reward functions](#). In [ICML 2024 Workshop on Models of Human Feedback for AI Alignment](#).
- Marco Cascella, Jonathan Montomoli, Valentina Bellini, and Elena Bignami. 2023. Evaluating the feasibility of chatgpt in healthcare: an analysis of multiple clinical and research scenarios. [Journal of medical systems](#), 47(1):33.

- Louis Castricato, Nathan Lile, Rafael Rafailov, Jan-Philipp Fränken, and Chelsea Finn. 2025. [PERSONA: A reproducible testbed for pluralistic alignment](#). In [Proceedings of the 31st International Conference on Computational Linguistics](#), pages 11348–11368, Abu Dhabi, UAE. Association for Computational Linguistics.
- Murali Chandrashekar, Beth A Walker, James C Ward, and Peter H Reingen. 1996. Modeling individual preference evolution and choice in a dynamic group setting. [Journal of Marketing Research](#), 33(2):211–223.
- Aili Chen, Xuyang Ge, Ziquan Fu, Yanghua Xiao, and Jiangjie Chen. 2024a. Travelagent: An ai assistant for personalized travel planning. [arXiv preprint arXiv:2409.08069](#).
- Daiwei Chen, Yi Chen, Aniket Rege, Zhi Wang, and Ramya Korlakai Vinayak. 2025a. [PAL: Sample-efficient personalized reward modeling for pluralistic alignment](#). In [The Thirteenth International Conference on Learning Representations](#).
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024b. [From persona to personalization: A survey on role-playing language agents](#). [Transactions on Machine Learning Research](#). Survey Certification.
- Ruizhe Chen, Xiaotian Zhang, Meng Luo, Wenhao Chai, and Zuozhu Liu. 2025b. [PAD: Personalized alignment at decoding-time](#). In [The Thirteenth International Conference on Learning Representations](#).
- Pengyu Cheng, Jiawen Xie, Ke Bai, Yong Dai, and Nan Du. 2023. Everyone deserves a reward: Learning customized human preferences. [arXiv preprint arXiv:2309.03126](#).
- Hana Choi, Carl F Mela, Santiago R Balseiro, and Adam Leary. 2020. Online display advertising markets: A literature review and future directions. [Information Systems Research](#), 31(2):556–575.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. [Advances in neural information processing systems](#), 30.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, et al. 2024. Ultrafeedback: Boosting language models with scaled ai feedback. In [Forty-first International Conference on Machine Learning](#).
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2023. Safe rlhf: Safe reinforcement learning from human feedback. [arXiv preprint arXiv:2310.12773](#).
- Zhenlong Dai, Chang Yao, WenKang Han, Yuanying Yuanying, Zhipeng Gao, and Jingyuan Chen. 2024. Mpcoder: Multi-user personalized code generator with explicit and implicit style representation learning. In [Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 3765–3780.
- Yuhao Dan, Jie Zhou, Qin Chen, Junfeng Tian, and Liang He. 2024. P-tailor: Customizing personality traits for language models via mixture of specialized lora experts. [arXiv preprint arXiv:2406.12548](#).
- Yi Dong, Zhilin Wang, Makesh Sreedhar, Xianchao Wu, and Oleksii Kuchaiev. 2023. [SteerLM: Attribute conditioned SFT as an \(user-steerable\) alternative to RLHF](#). In [Findings of the Association for Computational Linguistics: EMNLP 2023](#), pages 11275–11288, Singapore. Association for Computational Linguistics.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpacaeval: A simple way to debias automatic evaluators. [arXiv preprint arXiv:2404.04475](#).
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. Towards measuring the representation of subjective global opinions in language models. [arXiv preprint arXiv:2306.16388](#).
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with v-usable information. In [International Conference on Machine Learning](#), pages 5988–6008. PMLR.
- Flint Xiaofeng Fan, Cheston Tan, Yew-Soon Ong, Roger Wattenhofer, and Wei-Tsang Ooi. 2024. [Fedrlhf: A convergence-guaranteed federated framework for privacy-preserving and personalized rlhf](#). [Preprint, arXiv:2412.15538](#).
- Senay A Gebreab, Khaled Salah, Raja Jayaraman, Muhammad Habib ur Rehman, and Samer Ella-ham. 2024. Llm-based framework for administrative task automation in healthcare. In [2024 12th International Symposium on Digital Forensics and Security \(ISDFS\)](#), pages 1–7. IEEE.
- Zhuocheng Gong, Jian Guan, Wei Wu, Huishuai Zhang, Dongyan Zhao, and Rui Yan. 2024. Latent preference coding: Aligning large language models via discrete latent codes.
- Jian Guan and Minlie Huang. 2020. [UNION: An Unreferenced Metric for Evaluating Open-ended Story Generation](#). In [Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing \(EMNLP\)](#), pages 9157–9166, Online. Association for Computational Linguistics.
- Jian Guan, Wei Wu, zujie wen, Peng Xu, Hongning Wang, and Minlie Huang. 2024. [AMOR:](#)

- [A recipe for building adaptable modular knowledge agents through process feedback](#). In [The Thirty-eighth Annual Conference on Neural Information Processing Systems](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. [arXiv preprint arXiv:2501.12948](#).
- Yiju Guo, Ganqu Cui, Lifan Yuan, Ning Ding, Zexu Sun, Bowen Sun, Huimin Chen, Ruobing Xie, Jie Zhou, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2024. [Controllable preference optimization: Toward controllable multi-objective alignment](#). In [Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing](#), pages 1437–1454, Miami, Florida, USA. Association for Computational Linguistics.
- Necdet Gurkan, Mohammed Almarzouq, and Pon Rahul Murugaraj. 2024. [Personalized content moderation and emergent outcomes](#). [Preprint, arXiv:2405.09640](#).
- Ido Guy, Naama Zwerdling, Inbal Ronen, David Carmel, and Erel Uziel. 2010. Social media recommendation based on people and tags. In [Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval](#), pages 194–201.
- Richard Mervyn Hare. 1981. [Moral Thinking: Its Levels, Method and Point](#), volume 156. Clarendon Press.
- John C Harsanyi. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. [Journal of political economy](#), 63(4):309–321.
- Jerry Zhi-Yang He, Sashrika Pandey, Mariah L Schrum, and Anca Dragan. 2024. [Cos: Enhancing personalization and mitigating bias with context steering](#). [arXiv preprint arXiv:2405.01768](#).
- Jerry Zhi-Yang He, Sashrika Pandey, Mariah L Schrum, and Anca Dragan. 2025. [Context steering: Controllable personalization at inference time](#). In [The Thirteenth International Conference on Learning Representations](#).
- Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In [Proceedings of the 26th international conference on world wide web](#), pages 173–182.
- Liam Hebert, Krishna Sayana, Ambarish Jash, Alexandros Karatzoglou, Sukhdeep Sodhi, Sumanth Doddapaneni, Yanli Cai, and Dima Kuzmin. 2024. [Persona: Personalized soft prompt adapter architecture for personalized language prompting](#). [Preprint, arXiv:2408.00960](#).
- Rob High. 2012. The era of cognitive systems: An inside look at ibm watson and how it works. [IBM Corporation, Redbooks](#), 1:16.
- Buntaro Hiraki, Kensei Hamamoto, Ami Kimura, Masateru Tsunoda, Amjed Tahir, Kwabena Ebo Benin, Akito Monden, and Keitaro Nakasai. 2024. Personalization of code readability evaluation based on llm using collaborative filtering. [arXiv preprint arXiv:2411.10583](#).
- Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. 2022. [Are large pre-trained language models leaking your personal information?](#) In [Findings of the Association for Computational Linguistics: EMNLP 2022](#), pages 2038–2047, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Qiushi Huang, Xubo Liu, Tom Ko, Bo Wu, Wenwu Wang, Yu Zhang, and Lilian Tang. 2024a. Selective prompting tuning for personalized conversations with llms. [arXiv preprint arXiv:2406.18187](#).
- Saffron Huang, Divya Siddarth, Liane Lovitt, Thomas I Liao, Esin Durmus, Alex Tamkin, and Deep Ganguli. 2024b. Collective constitutional ai: Aligning a language model with public input. In [The 2024 ACM Conference on Fairness, Accountability, and Transparency](#), pages 1395–1417.
- Wenke Huang, Mang Ye, Zekun Shi, Guancheng Wan, He Li, Bo Du, and Qiang Yang. 2024c. [Federated learning for generalization, robustness, fairness: A survey and benchmark](#). [IEEE Transactions on Pattern Analysis and Machine Intelligence](#), 46(12):9387–9406.
- Keise Izuma and Ralph Adolphs. 2013. Social manipulation of preference in the human brain. [Neuron](#), 78(3):563–573.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. [arXiv preprint arXiv:2310.11564](#).
- Kalervo Järvelin and Jaana Kekäläinen. 2017. [Ir evaluation methods for retrieving highly relevant documents](#). [SIGIR Forum](#), 51(2):243–250.
- Gawesh Jawaheer, Peter Weller, and Patty Kostkova. 2014. Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback. [ACM Transactions on Interactive Intelligent Systems \(TiIS\)](#), 4(2):1–26.
- Jaeho Jeon and Seongyong Lee. 2023. Large language models in education: A focus on the complementary relationship between human teachers and chatgpt. [Education and Information Technologies](#), 28(12):15873–15892.

- Jiaming Ji, Mickel Liu, Josef Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. 2024. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, et al. 2023. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*.
- Eunkyoung Jo, Daniel A Epstein, Hyunhoon Jung, and Young-Ho Kim. 2023. Understanding the benefits and challenges of deploying conversational ai leveraging large language models for public health intervention. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–16.
- Md Rayhan Kabir and Fuhua Lin. 2023. An llm-powered adaptive practicing system. In *LLM@AIED*, pages 43–52.
- Ehud Kalai and Meir Smorodinsky. 1975. Other solutions to nash’s bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 513–518.
- Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. 2023. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv preprint arXiv:2305.06474*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Atoosa Kasirzadeh and Iason Gabriel. 2023. In conversation with artificial intelligence: aligning language models with human values. *Philosophy & Technology*, 36(2):27.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. 2023. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.
- Zachary Kenton, Noah Yamamoto Siegel, Janos Kramar, Jonah Brown-Cohen, Samuel Albanie, Jannis Bulian, Rishabh Agarwal, David Lindner, Yunhao Tang, Noah Goodman, and Rohin Shah. 2024. *On scalable oversight with weak LLMs judging strong LLMs*. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Maxim Khanov, Jirayu Burapachee, and Yixuan Li. 2024. Args: Alignment as reward-guided search. *arXiv preprint arXiv:2402.01694*.
- Minbeom Kim, Kang-il Lee, Seongho Joo, and Hwaran Lee. 2025. Drift: Decoding-time personalized alignments with implicit user preferences. *arXiv preprint arXiv:2502.14289*.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, et al. 2023. Prometheus: Inducing fine-grained evaluation capability in language models. In *The Twelfth International Conference on Learning Representations*.
- Hannah Rose Kirk, Andrew Michael Bean, Bertie Vidgen, Paul Rottger, and Scott A. Hale. 2023a. *The past, present and better future of feedback learning in large language models for subjective human preferences and values*. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2024a. The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*, pages 1–10.
- Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, et al. 2024b. The prism alignment project: What participatory, representative and individualised human feedback reveals about the subjective and multi-cultural alignment of large language models. *arXiv preprint arXiv:2404.16019*.
- Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward Grefenstette, and Roberta Raileanu. 2023b. Understanding the effects of rlhf on llm generalisation and diversity. *arXiv preprint arXiv:2310.06452*.
- Roham Koohestani and Maliheh Izadi. 2025. Rethinking ide customization for enhanced hax: A hyperdimensional perspective. *arXiv preprint arXiv:2501.02491*.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.
- Osama Koraisi. 2023. Teaching english in the age of ai: Embracing chatgpt to optimize efl materials and assessment. *Language Education and Technology*, 3(1).
- Anton Korinek and Avital Balwit. 2022. Aligned with whom? direct and social goals for ai systems. Technical report, National Bureau of Economic Research.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. Gender bias and stereotypes in large language models. In *Proceedings of the ACM collective intelligence conference*, pages 12–24.

- Weirui Kuang, Bingchen Qian, Zitao Li, Daoyuan Chen, Dawei Gao, Xuchen Pan, Yuexiang Xie, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024. Federatedscope-llm: A comprehensive package for fine-tuning large language models in federated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5260–5271.
- Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A. Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka, Chien Van Nguyen, Thien Huu Nguyen, and Hamed Zamani. 2024a. *Longlamp: A benchmark for personalized long-form text generation*. Preprint, arXiv:2407.11016.
- Sachin Kumar, Chan Young Park, Yulia Tsvetkov, Noah A Smith, and Hannaneh Hajishirzi. 2024b. Compo: Community preferences for language model personalization. arXiv preprint arXiv:2410.16027.
- Subodha Kumar, Varghese S Jacob, and Chelliah Sriskandarajah. 2006. Scheduling advertisements on a web page to maximize revenue. *European journal of operational research*, 173(3):1067–1089.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, et al. 2024. Rewardbench: Evaluating reward models for language modeling. arXiv preprint arXiv:2403.13787.
- Royson Lee, Minyoung Kim, Da Li, Xinchu Qiu, Timothy Hospedales, Ferenc Huszár, and Nicholas Lane. 2024a. Fedl2p: Federated learning to personalize. *Advances in Neural Information Processing Systems*, 36.
- Seongyun Lee, Sue Hyun Park, Seungone Kim, and Minjoon Seo. 2024b. Aligning to thousands of preferences via system message generalization. arXiv preprint arXiv:2405.17977.
- Daniel Leiker, Sara Finnigan, Ashley Ricker Gyllen, and Mutlu Cukurova. 2023. Prototyping the use of large language models (llms) for adult learning content creation at scale. arXiv preprint arXiv:2306.01815.
- Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. 2024a. Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM on Web Conference 2024*, pages 3367–3378.
- Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. 2023a. Teach llms to personalize—an approach inspired by writing education. arXiv preprint arXiv:2308.07968.
- Hao Li, Chenghao Yang, An Zhang, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2024b. Hello again! llm-powered personalized agent for long-term dialogue. arXiv preprint arXiv:2406.05925.
- Jia-Nan Li, Jian Guan, Songhao Wu, Wei Wu, and Rui Yan. 2025. *From 1,000,000 users to every user: Scaling up personalized preference for user-level alignment*. Preprint, arXiv:2503.15463.
- Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. 2023b. Text is all you need: Learning language representations for sequential recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1258–1267.
- Peng Li, Lixia Wu, Chaoqun Feng, Haoyuan Hu, Lei Fu, and Jieping Ye. 2024c. *Differentiation of multi-objective data-driven decision pipeline*. Preprint, arXiv:2406.00779.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024d. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. arXiv preprint arXiv:2406.11939.
- Xinyi Li, Yongfeng Zhang, and Edward C Malthouse. 2023c. A preliminary study of chatgpt on news recommendation: Personalization, provider fairness, fake news. arXiv preprint arXiv:2306.10702.
- Xinyu Li, Ruiyang Zhou, Zachary C Lipton, and Liu Leqi. 2024e. Personalized language modeling from personalized human feedback. arXiv preprint arXiv:2402.05133.
- Yan Li, Hongyang Yan, Teng Huang, Zijie Pan, Jiewei Lai, Xiaoxue Zhang, Kongyang Chen, and Jin Li. 2024f. Model architecture level privacy leakage in neural networks. *Science China Information Sciences*, 67(3):132101.
- Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, et al. 2024g. Personal llm agents: Insights and survey about the capability, efficiency and security. arXiv preprint arXiv:2401.05459.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. arXiv preprint arXiv:2305.20050.
- Guanyu Lin, Tao Feng, Pengrui Han, Ge Liu, and Jiaxuan You. 2024. Paper copilot: A self-evolving and efficient llm system for personalized academic assistance. arXiv preprint arXiv:2409.04593.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2021. Truthfulqa: Measuring how models mimic human falsehoods. arXiv preprint arXiv:2109.07958.
- Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing*, 7(1):76–80.

- Fei Liu et al. 2020. Learning to summarize from human feedback. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics.
- Junling Liu, Chao Liu, Peilin Zhou, Renjie Lv, Kang Zhou, and Yan Zhang. 2023a. Is chatgpt a good recommender? a preliminary study. arXiv preprint arXiv:2304.10149.
- Xin Liu. 2015. Modeling users' dynamic preference for personalized recommendation. In Twenty-fourth international joint conference on artificial intelligence.
- Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, Zihao Wu, Lin Zhao, Dajiang Zhu, Xiang Li, Ning Qiang, Dingang Shen, Tianming Liu, and Bao Ge. 2023b. Summary of chatgpt-related research and perspective towards the future of large language models. Meta-Radiology, 1(2):100017.
- Sophie Lythreathis, Sanjay Kumar Singh, and Abdul-Nasser El-Kassar. 2022. The digital divide: A review and future research agenda. Technological Forecasting and Social Change, 175:121359.
- Tala Mirzaei, Leila Amini, and Pouyan Esmailzadeh. 2024. Clinician voices on ethics of llm integration in healthcare: A thematic analysis of ethical concerns and implications. BMC Medical Informatics and Decision Making, 24(1):250.
- Alan L Montgomery and Michael D Smith. 2009. Prospects for personalization on the internet. Journal of Interactive Marketing, 23(2):130–137.
- Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. 2024. More human than human: measuring chatgpt political bias. Public Choice, 198(1):3–23.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. arXiv preprint arXiv:2112.09332.
- Mohamed Nejjar, Luca Zacharias, Fabian Stiehle, and Ingo Weber. 2025. Llms for science: Usage for code generation and data analysis. Journal of Software: Evolution and Process, 37(1):e2723.
- Lin Ning, Luyang Liu, Jiaying Wu, Neo Wu, Devora Berlowitz, Sushant Prakash, Bradley Green, Shawn O'Banion, and Jun Xie. 2024. User-llm: Efficient llm contextualization with user embeddings. arXiv preprint arXiv:2402.13598.
- Ninell Oldenburg and Tan Zhi-Xuan. 2024. Learning and sustaining shared normative systems via bayesian rule induction in markov games. In Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, pages 1510–1520.
- OpenAI. 2025. Introducing deep research. Accessed: 2025-02-10.
- OpenAI. 2025. Introducing GPT-4.5: Advancing artificial intelligence. Accessed: 2025-02-27.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. Advances in neural information processing systems, 35:27730–27744.
- Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Xufang Luo, Hao Cheng, Dongsheng Li, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Jianfeng Gao. 2025. Secom: On memory construction and retrieval for personalized conversational agents. In The Thirteenth International Conference on Learning Representations.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In Proceedings of the 40th annual meeting of the Association for Computational Linguistics, pages 311–318.
- Chanwoo Park, Mingyang Liu, Dingwen Kong, Kaiqing Zhang, and Asuman E Ozdaglar. 2024. Rlhf from heterogeneous feedback via personalization and preference aggregation. In ICML 2024 Workshop: Aligning Reinforcement Learning Experimentalists and Theorists.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. arXiv preprint arXiv:2304.03277.
- Renjie Pi, Tianyang Han, Wei Xiong, Jipeng Zhang, Runtao Liu, Rui Pan, and Tong Zhang. 2024. Strengthening multimodal large language model with bootstrapped preference optimization. In European Conference on Computer Vision, pages 382–398. Springer.
- Renjie Pi, Jianshu Zhang, Tianyang Han, Jipeng Zhang, Rui Pan, and Tong Zhang. 2025. Personalized visual instruction tuning. In The Thirteenth International Conference on Learning Representations.
- Sriyash Poddar, Yanming Wan, Hamish Ivison, Abhishek Gupta, and Natasha Jaques. 2024. Personalizing reinforcement learning from human feedback with variational preference learning. In The Thirty-eighth Annual Conference on Neural Information Processing Systems.
- S Preetha and KN Vimal Shankar. 2014. Personalized search engines on mining user preferences using clickthrough data. In International Conference on Information Communication and Embedded Systems (ICICES2014), pages 1–6. IEEE.
- Alexander Pretschner and Susan Gauch. 1999. Ontology based personalized search. In Proceedings 11th International Conference on Tools with Artificial Intelligence, pages 391–398. IEEE.

- Jiaxing QI, Zhongzhi Luan, Shaohan Huang, Carol Fung, Hailong Yang, and Depei Qian. 2024. [Fdlora: Personalized federated learning of large language model via dual lora tuning](#). [Preprint](#), arXiv:2406.07925.
- Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. 2023. [Rewarded soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards](#). In [Thirty-seventh Conference on Neural Information Processing Systems](#).
- Shaina Raza, Shardul Ghuge, Chen Ding, and Deval Pandya. 2024. Fair enough: How can we develop and assess a fair-compliant dataset for large language models' training? [arXiv preprint arXiv:2401.11033](#).
- Paul Resnick and Hal R Varian. 1997. Recommender systems. [Communications of the ACM](#), 40(3):56–58.
- Yijiang River Dong, Tiancheng Hu, Yinhong Liu, Ahmet Üstün, and Nigel Collier. 2025. When personalization meets reality: A multi-faceted analysis of personalized preference learning. [arXiv e-prints](#), pages arXiv–2502.
- Peter E Rossi, Robert E McCulloch, and Greg M Allenby. 1996. The value of purchase history data in target marketing. [Marketing Science](#), 15(4):321–340.
- Yasaman Saadati, Mohammad Rostami, and M. Hadi Amini. 2025. [pmixed: Efficient personalized federated learning through adaptive layer-wise mixup](#). [Preprint](#), arXiv:2501.11002.
- Sahand Sabour, June M. Liu, Siyang Liu, Chris Z. Yao, Shiyao Cui, Xuanming Zhang, Wen Zhang, Yaru Cao, Advait Bhat, Jian Guan, Wei Wu, Rada Mihalcea, Tim Althoff, Tatia M. C. Lee, and Minlie Huang. 2025. [Human decision-making is susceptible to ai-driven manipulation](#). [Preprint](#), arXiv:2502.07663.
- Sahand Sabour, Siyang Liu, Zheyuan Zhang, June Liu, Jinfeng Zhou, Alvionna Sunaryo, Tatia Lee, Rada Mihalcea, and Minlie Huang. 2024. Emobench: Evaluating the emotional intelligence of large language models. In [Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 5986–6004.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2023. Lamp: When large language models meet personalization. [arXiv preprint arXiv:2304.11406](#).
- Malik Sallam. 2023. Chatgpt utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. In [Healthcare](#), volume 11, page 887. MDPI.
- Cynthia K Sanders and Edward Scanlon. 2021. The digital divide is a human rights issue: Advancing social inclusion through social work advocacy. [Journal of human rights and social work](#), 6(2):130–143.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In [International Conference on Machine Learning](#), pages 29971–30004. PMLR.
- Krishna Sayana, Raghavendra Vasudeva, Yuri Vasilevski, Kun Su, Liam Hebert, James Pine, Hubert Pham, Ambarish Jash, and Sukhdeep Sodhi. 2024. [Beyond retrieval: Generating narratives in conversational recommender systems](#). [Preprint](#), arXiv:2410.16780.
- J Ben Schafer, Joseph A Konstan, and John Riedl. 2001. E-commerce recommendation applications. [Data mining and knowledge discovery](#), 5:115–153.
- Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In [Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval](#), pages 253–260.
- James C Scott. 2020. [Seeing like a state: How certain schemes to improve the human condition have failed](#). yale university Press.
- Amartya Sen. 1986. Social choice theory. [Handbook of mathematical economics](#), 3:1073–1181.
- Amartya Sen. 2017. [Collective choice and social welfare: Expanded edition](#). Penguin UK.
- Guy Shani and Asela Gunawardana. 2011. Evaluating recommendation systems. [Recommender systems handbook](#), pages 257–297.
- Han Shao, Lee Cohen, Avrim Blum, Yishay Mansour, Aadirupa Saha, and Matthew Walter. 2024. Eliciting user preferences for personalized multi-objective decision making through comparative feedback. [Advances in Neural Information Processing Systems](#), 36.
- Renee Shelby, Shalaleh Rismani, Kathryn Henne, Ajung Moon, Negar Rostamzadeh, Paul Nicholas, YILLA-AKBARI N'MAH, Jess Gallegos, Andrew Smart, and Gurleen Virk. 2022. Identifying sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction. [arXiv preprint arXiv:2210.05791](#).
- Idan Shenfeld, Felix Faltings, Pulkit Agrawal, and Aldo Pacchiano. 2025. Language model personalization via reward factorization. [arXiv preprint arXiv:2503.06358](#).
- Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hananeh Hajishirzi, Noah A Smith, and Simon S Du. 2024. Decoding-time language model alignment with multiple objectives. [arXiv preprint arXiv:2406.18853](#).



- Yubo Shu, Haonan Zhang, Hansu Gu, Peng Zhang, Tun Lu, Dongsheng Li, and Ning Gu. 2024. Rah! recsys–assistant–human: A human-centered recommendation framework with llm agents. IEEE Transactions on Computational Social Systems.
- Andrew Siah, Thomas P. Zollo, Naimeng Ye, Ang Li, and Hongseok Namkoong. 2024. Personallm: Tailoring llms to individual preferences. Columbia University. {andrew.siah, tpz2105, ny2336, a14263, hn2369}@columbia.edu.
- Harmanpreet Singh, Nikhil Verma, Yixiao Wang, Manasa Bharadwaj, Homa Fashandi, Kevin Ferreira, and Chul Lee. 2024. Personal large language model agents: A case study on tailored travel planning. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track, pages 486–514.
- Anand Siththaranjan, Cassidy Laidlaw, and Dylan Hadfield-Menell. 2024. Distributional preference learning: Understanding and accounting for hidden context in RLHF. In The Twelfth International Conference on Learning Representations.
- Asim Smailagic and Daniel Siewiorek. 2002. Application design for wearable and context-aware computers. IEEE Pervasive Computing, 1(4):20–29.
- Mirco Speretta and Susan Gauch. 2005. Personalized search based on user search histories. In The 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI'05), pages 622–628. IEEE.
- James WA Strachan, Dalila Albergio, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, et al. 2024. Testing theory of mind in large language models and humans. Nature Human Behaviour, pages 1–11.
- Xiaoyuan Su. 2009. A survey of collaborative filtering techniques.
- Yuqian Sun, Phoebe J Wang, John Joon Young Chung, Melissa Roemmele, Taewook Kim, and Max Kreminski. 2025. Drama llama: An llm-powered storylets framework for authorable responsiveness in interactive narrative. arXiv preprint arXiv:2501.09099.
- Kar Yan Tam and Shuk Ying Ho. 2006. Understanding the impact of web personalization on user information processing and decision outcomes. MIS quarterly, pages 865–890.
- Zhaoxuan Tan, Zheyuan Liu, and Meng Jiang. 2024a. Personalized pieces: Efficient personalized large language models through collaborative efforts. arXiv preprint arXiv:2406.10471.
- Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. 2024b. Democratizing large language models via personalized parameter-efficient fine-tuning. arXiv preprint arXiv:2402.04401.
- Yihong Tang, Zhaokai Wang, Ao Qu, Yihao Yan, Zhaofeng Wu, Dingyi Zhuang, Jushi Kai, Kebin Hou, Xiaotong Guo, Jinhua Zhao, et al. 2024. Itinera: Integrating spatial optimization with large language models for open-domain urban itinerary planning. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track, pages 1413–1432.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530.
- Lukas Teufelberger, Xintong Liu, Zhipeng Li, Max Moebus, and Christian Holz. 2024. Llm-for-x: Application-agnostic integration of large language models to support personal writing workflows. arXiv preprint arXiv:2407.21593.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrusti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.
- Linh Tran, Wei Sun, Stacy Patterson, and Ana Milanova. 2025. Privacy-preserving personalized federated prompt learning for multimodal large language models. In The Thirteenth International Conference on Learning Representations.
- Catherine E Tucker. 2014. Social networks, personalized advertising, and privacy controls. Journal of marketing research, 51(5):546–562.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc.
- Katherine Verdery. 2005. What was socialism, and why did it fall? In The Revolutions of 1989, pages 73–94. Routledge.
- Hongru Wang, Wenyu Huang, Yang Deng, Rui Wang, Zezhong Wang, Yufei Wang, Fei Mi, Jeff Z Pan, and Kam-Fai Wong. 2024a. Unims-rag: A unified multi-source retrieval-augmented generation for personalized dialogue systems. arXiv preprint arXiv:2401.13256.
- Jie Wang, Alexandros Karatzoglou, Ioannis Arapakis, Xin Xin, Xuri Ge, and Joemon M. Jose. 2024b. Sparks of surprise: Multi-objective recommendations with hierarchical decision transformers for diversity, novelty, and serendipity. In Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24, page 2358–2368, New York, NY, USA. Association for Computing Machinery.

- Kaiwen Wang, Rahul Kidambi, Ryan Sullivan, Alekh Agarwal, Christoph Dann, Andrea Michi, Marco Gelmi, Yunxuan Li, Raghav Gupta, Kumar Avinava Dubey, Alexandre Rame, Johan Ferret, Geoffrey Cideron, Le Hou, Hongkun Yu, Amr Ahmed, Aranyak Mehta, Leonard Hussenot, Olivier Bachem, and Edouard Leurent. 2024c. [Conditional language policy: A general framework for steerable multi-objective finetuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2153–2186, Miami, Florida, USA. Association for Computational Linguistics.
- Sitong Wang, Xuanming Zhang, Jenny Ma, Alyssa Hwang, and Lydia B Chilton. 2024d. [Jumpstarter: Getting started on personal goals with ai-powered context curation](#). [arXiv preprint arXiv:2410.03882](#).
- Tianyang Wang, Ziqian Bi, Yichao Zhang, Ming Liu, Weiche Hsieh, Pohsun Feng, Lawrence K. Q. Yan, Yizhu Wen, Benji Peng, Junyu Liu, Keyu Chen, Sen Zhang, Ming Li, Chuanqi Jiang, Xinyuan Song, Junjie Yang, Bowen Jing, Jintao Ren, Junhao Song, Hong-Ming Tseng, Silin Chen, Yunze Wang, Chia Xin Liang, Jiawei Xu, Xuanhe Pan, Jinlang Wang, and Qian Niu. 2024e. [Deep learning model security: Threats and defenses](#). [Preprint, arXiv:2412.08969](#).
- Xintao Wang, Jiangjie Chen, Nianqi Li, Lida Chen, Xinfeng Yuan, Wei Shi, Xuyang Ge, Rui Xu, and Yanghua Xiao. 2024f. [Surveyagent: A conversational system for personalized and efficient research survey](#). [arXiv preprint arXiv:2404.06364](#).
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. [Self-instruct: Aligning language models with self-generated instructions](#). [arXiv preprint arXiv:2212.10560](#).
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, et al. 2023a. [Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models](#). [arXiv preprint arXiv:2310.00746](#).
- Zhefan Wang, Yuanqing Yu, Wendi Zheng, Weizhi Ma, and Min Zhang. 2024g. [Macrec: A multi-agent collaboration framework for recommendation](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2760–2764.
- Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. 2024h. [Helpsteer2: Open-source dataset for training top-performing reward models](#). [arXiv preprint arXiv:2406.08673](#).
- Zhilin Wang, Yi Dong, Jiaqi Zeng, Virginia Adams, Makesh Narsimhan Sreedhar, Daniel Egert, Olivier Delalleau, Jane Polak Scowcroft, Neel Kant, Aidan Swope, et al. 2023b. [Helpsteer: Multi-attribute helpfulness dataset for steerlm](#). [arXiv preprint arXiv:2311.09528](#).
- Zilong Wang, Yuedong Cui, Li Zhong, Zimin Zhang, Da Yin, Bill Yuchen Lin, and Jingbo Shang. 2024i. [Officebench: Benchmarking language agents across multiple applications for office automation](#). [arXiv preprint arXiv:2407.19056](#).
- Michel Wedel and PK Kannan. 2016. [Marketing analytics for data-rich environments](#). *Journal of marketing*, 80(6):97–121.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022a. [Emergent abilities of large language models](#). [arXiv preprint arXiv:2206.07682](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. [Taxonomy of risks posed by language models](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229.
- Martin Weyssow, Aton Kamanda, Xin Zhou, and Houari Sahraoui. 2024. [Codeultrafeedback: An llm-as-a-judge dataset for aligning large language models to coding preferences](#). [arXiv preprint arXiv:2403.09032](#).
- Kenneth R Wilson, Jennifer S Wallin, and Christa Reiser. 2003. [Social stratification and the digital divide](#). *Social Science Computer Review*, 21(2):133–143.
- Shujin Wu, Yi R. Fung, Cheng Qian, Jeonghwan Kim, Dilek Hakkani-Tur, and Heng Ji. 2025. [Aligning LLMs with individual preferences via interaction](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7648–7662, Abu Dhabi, UAE. Association for Computational Linguistics.
- Weiqi Wu, Hongqiu Wu, Lai Jiang, Xingyuan Liu, Jiale Hong, Hai Zhao, and Min Zhang. 2024a. [From role-play to drama-interaction: An llm solution](#). [arXiv preprint arXiv:2405.14231](#).
- Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. [A survey of human-in-the-loop for machine learning](#). *Future Generation Computer Systems*, 135:364–381.
- Zequi Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2024b. [Fine-grained human feedback gives better rewards for](#)

- language model training. *Advances in Neural Information Processing Systems*, 36.
- Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and Yu Su. 2024. Travelplanner: A benchmark for real-world planning with language agents. In *Forty-first International Conference on Machine Learning*.
- Jing Xu, Arthur Szlam, and Jason Weston. 2022a. [Beyond goldfish memory: Long-term open-domain conversation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.
- Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang, and Shihang Wang. 2022b. Long time no see! open-domain conversation with long-term persona memory. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2639–2650.
- Gautam Yadav, Ying-Jui Tseng, and Xiaolin Ni. 2023. Contextualizing problems to student interests at scale in intelligent tutoring system using large language models. [arXiv preprint arXiv:2306.00190](#).
- Kailai Yang, Zhiwei Liu, Qianqian Xie, Jimin Huang, Tianlin Zhang, and Sophia Ananiadou. 2024a. Metaaligner: Towards generalizable multi-objective alignment of language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Rui Yang, Xiaoman Pan, Feng Luo, Shuang Qiu, Han Zhong, Dong Yu, and Jianshu Chen. 2024b. Rewards-in-context: Multi-objective alignment of foundation models with dynamic preference adjustment. [arXiv preprint arXiv:2402.10207](#).
- Seonghyeon Ye, Doyoung Kim, Sungdong Kim, Hyeonbin Hwang, Seungone Kim, Yongrae Jo, James Thorne, Juho Kim, and Minjoon Seo. 2023. Flask: Fine-grained language model evaluation based on alignment skill sets. [arXiv preprint arXiv:2307.10928](#).
- Liping Yi, Han Yu, Chao Ren, Heng Zhang, Gang Wang, Xiaoguang Liu, and Xiaoxiao Li. 2024. [pfdmoe: Data-level personalization with mixture of experts for model-heterogeneous personalized federated learning](#). Preprint, arXiv:2402.01350.
- Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin Chen, Ruobing Xie, Yankai Lin, et al. 2024. Advancing llm reasoning generalists with preference trees. [arXiv preprint arXiv:2404.02078](#).
- Judith Lynne Zaichkowsky. 1994. The personal involvement inventory: Reduction, revision, and application to advertising. *Journal of advertising*, 23(4):59–70.
- An Zhang, Yuxin Chen, Leheng Sheng, Xiang Wang, and Tat-Seng Chua. 2024a. On generative agents in recommendation. In *Proceedings of the 47th international ACM SIGIR conference on research and development in Information Retrieval*, pages 1807–1817.
- Gangyi Zhang. 2023. [User-centric conversational recommendation: Adapting the need of user with large language models](#). In *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys '23*, page 1349–1354, New York, NY, USA. Association for Computing Machinery.
- Jizhi Zhang, Keqin Bao, Wenjie Wang, Yang Zhang, Wentao Shi, Wanhong Xu, Fuli Feng, and Tat-Seng Chua. 2024b. Prospect personalized recommendation on large language model-based agent platform. [arXiv preprint arXiv:2402.18240](#).
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. [arXiv preprint arXiv:2308.10792](#).
- Starkson Zhang, Alfredo Pearson, and Zhenting Wang. Autonomous generalist scientist: Towards and beyond human-level automatic research using foundation model-based ai agents and robots (a position).
- Xiaokang Zhang, Jing Zhang, Zeyao Ma, Yang Li, Bohan Zhang, Guanlin Li, Zijun Yao, Kangli Xu, Jinchang Zhou, Daniel Zhang-Li, et al. 2024c. Tablellm: Enabling tabular data manipulation by llms in real office usage scenarios. [arXiv preprint arXiv:2403.19318](#).
- Zhaowei Zhang, Fengshuo Bai, Qizhi Chen, Chengdong Ma, Mingzhi Wang, Haoran Sun, Zilong Zheng, and Yaodong Yang. 2025. [Amulet: Realignment during test time for personalized preference adaptation of LLMs](#). In *The Thirteenth International Conference on Learning Representations*.
- Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. 2024d. Personalization of large language models: A survey. [arXiv preprint arXiv:2411.00027](#).
- Siyuan Zhao, Mingyi Hong, Yang Liu, Devamanyu Hazarika, and Kaixiang Lin. 2025. [Do LLMs recognize your preferences? evaluating personalized preference following in LLMs](#). In *The Thirteenth International Conference on Learning Representations*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. [arXiv preprint arXiv:2303.18223](#).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023.

Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

Jinfeng Zhou, Zhuang Chen, Dazhen Wan, Bosi Wen, Yi Song, Jifan Yu, Yongkang Huang, Libiao Peng, Jiaming Yang, Xiyao Xiao, et al. 2023. Character-glm: Customizing chinese conversational ai characters with large language models. *arXiv preprint arXiv:2311.16832*.

Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. 2023. Starling-7b: Improving llm helpfulness & harmlessness with rlaiif.

Thomas P Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, and Hongseok Namkoong. 2025. *PersonalLLM: Tailoring LLMs to individual preferences*. In *The Thirteenth International Conference on Learning Representations*.

Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.

Álvaro Bartolomé Del Canto, Gabriel Martín Blázquez, Agustín Piqueres Lajarín, and Daniel Vila Suero. 2024. Distilabel: An ai feedback (aif) framework for building datasets with and for llms. <https://github.com/argilla-io/distilabel>.

## A Related Survey

Prior surveys have extensively studied different aspects of LLMs and personalization. Zhao et al. (2023) review the general development of LLMs, while Ji et al. (2023) specifically examine AI alignment techniques. Zhang et al. (2024d) investigate personalized applications. Chen et al. (2024b) focus on role-playing LLMs that simulate specific personas. While these works lay important foundations, personalized alignment presents fundamentally different challenges that have not been systematically addressed. Unlike traditional alignment that optimizes for universal values (Ji et al., 2023), personalized alignment must handle diverse objectives that vary across users and evolve over time while ensuring ethical boundaries. Unlike conventional personalization that focuses on surface-level preferences (Zhang et al., 2024d), our framework addresses deeper preference inference covering fundamental values and manages complex trade-offs between universal values and personal preferences. Distinctively, our survey provides the first comprehensive examination of these unique challenges, from privacy-preserving preference learning and real-time adaptation techniques to evaluation methods and risk mitigation strategies.

## B Case Study

Figure 5 shows an example of how a personalized alignment system infers user preferences from various sources.

## C Application of Personalized Alignment

This section discusses the application of personalized alignment.

### C.1 Personal Assistants

Personalized LLMs advance traditional AI assistance (e.g., open-domain QA (High, 2012)) by providing domain-specific support, including: intelligent coding companions with tailored completions and optimizations (Dai et al., 2024; Hiraki et al., 2024; Nejjar et al., 2025; Koohestani and Izadi, 2025), research aids generating customized literature reviews (Wang et al., 2024f; Lin et al., 2024), and workplace assistants for personalized task management (Zhang et al., 2024c; Wang et al., 2024i,d; Teufelberger et al., 2024).

**Task:** Given the following preference memory, infer the user's preferences and generate preference-aligned responses.

**Preference Memory:**

**User-Generated Content:**

**Post (User):** "Can you explain quantum computing?"

**Response (LLM):** "Quantum computing leverages quantum mechanical phenomena like superposition and entanglement. In quantum bits or qubits, states can exist simultaneously in multiple configurations, enabling parallel computations through quantum gates and ..."

**Post (User):** "That's too theoretical. Could you give real-world examples of how it's used in business or daily life? Theory without application isn't helpful." (prefers practical applications)

**Demographic Attributes:** Professional Background: MBA with 5 years consulting experience; Learning Style: Action-oriented, prefers case studies over theoretical frameworks; Content Interaction: High engagement with business case analyses and implementation guides; Reading Pattern: Frequently accesses professional development and practical business strategy content (prefers practical over theoretical content)

**Behavioral Signals:**

**Post (User):** "I'm planning a trip to Japan next month. Any suggestions welcome!"

**Response A (LLM):** "Here's a comprehensive cultural guide: Japan has 47 prefectures, 8 major regions, and a history dating back to 30,000 BCE ..."

**Response B (LLM):** "Let's focus on practical tips: Learn basic greetings like 'arigatou' (thank you), avoid tipping as it's not customary, and remove shoes before entering homes ..."

**User's Preference:** Response B > Response A (Prefers practical, accessible information)

**Post (User):** "How can I improve my public speaking skills?"

**Response A (LLM):** "Public speaking originated in ancient Greece with rhetoric principles developed by Aristotle. The fundamental elements include ethos, pathos, and logos..."

**Response B (LLM):** "Here are specific techniques you can use: 1) Start with a relevant story, 2) Practice the 10-20-30 rule: 10 slides, 20 minutes, 30-point font..."

**User's Preference:** Response B > Response A (prefers actionable advice)

**System's Preference Inference:** Based on the user's consistent preference for practical explanations over theoretical details across different domains, the user strongly prefers practical, accessible explanations with real-world applications over theoretical details

**Current Post (User):** "What's the best way to understand blockchain?"

**Candidate Response 1 (LLM):** "Blockchain is a distributed ledger technology utilizing cryptographic hashing functions and consensus mechanisms..."

**Candidate Response 2 (LLM):** "Think of blockchain like a shared digital notebook - everyone has a copy, and when someone writes something new, everyone's copy gets updated automatically..."

**System's Alignment Objective:** Response 2 > Response 1

Figure 5: Example showing how a personalized alignment system infers user preferences from multiple information sources and generates preference-aligned responses.

## C.2 Consumer Applications

Personalized LLMs enhance digital services via user-generated content analysis. In recommendations, they transcend traditional approaches (Su, 2009; He et al., 2017) via instruction-based generation (Liu et al., 2023a; Li et al., 2023b,c) or multi-agent frameworks (Wang et al., 2024g). Entertainment applications include adaptive companions (Zhou et al., 2023), role-playing (Chen et al., 2024b), and interactive narratives (Wu et al., 2024a; Sun et al., 2025), though raising ethical concerns regarding emotional manipulation.

## C.3 Public Services

In education, Personalized LLMs enable individualized learning experiences (Yadav et al., 2023; Kabir and Lin, 2023) and support educators with automated analytics and material generation (Jeon and Lee, 2023; Leiker et al., 2023; Koraiishi, 2023; Kasneci et al., 2023). In healthcare, they advance patient care (Casella et al., 2023; Gebreab et al., 2024) by integrating medical knowledge with individual profiles (Abbasian et al., 2023; Jo et al., 2023; Sallam, 2023), though requiring rigorous validation (Mirzaei et al., 2024).