

# ICEBREAKER for Conversational Agents: Breaking the First-Message Barrier with Personalized Starters

Hongwei Zheng\*, Weiqi Wu\*, Zhengjia Wang\*, Guanyu Jiang,  
Haoming Li, Tianyu Wu, Yongchun Zhu†, Jingwu Chen†, Feng Zhang  
ByteDance

{zhenghongwei, wuweiqi, wangzhengjia.jia, jiangguanyu.jgy,

lihaoming.cs, wutianyu.23, zhuyongchun.zyc, chenjingwu, feng.zhang}@bytedance.com

## Abstract

Conversational agents, such as ChatGPT and Doubao, have become essential daily assistants for billions of users. To further enhance engagement, these systems are evolving from passive responders to proactive companions. However, existing efforts focus on activation within ongoing dialogues, while overlooking a key real-world bottleneck. In the conversation initiation stage, users may have a vague need but no explicit query intent, creating a first-message barrier where the conversation holds before it begins. To overcome this, we introduce **Conversation Starter Generation**: generating personalized starters to guide users into conversation. However, unlike in-conversation stages where immediate context guides the response, initiation must operate in a cold-start moment without explicit user intent. To pioneer in this direction, we present **ICEBREAKER** that frames human ice-breaking via a two-step handshake: (i) *evoked resonance* via Resonance-Aware Interest Distillation from session summaries to capture trigger interests, and (ii) *stimulate interaction* via Interaction-Oriented Starter Generation, optimized with personalized preference alignment and a self-reinforced loop to maximize engagement. Online A/B tests on one of the world’s largest conversational agent products show that ICEBREAKER improves user active days by +1.84% and click-through rate by +94.25%, and has been deployed in production.

## 1 Introduction

Conversational agents are rapidly becoming an indispensable part of everyday life (Hadi et al., 2023), people turn to them not only for information seeking (Zhu et al., 2025) and task assistance (Qin et al., 2024), but also for planning and emotional companionship (Kasneji et al., 2023; Li et al., 2023). In this work, we study a deployment setting on one of

\*Equal contribution.

†Corresponding author.

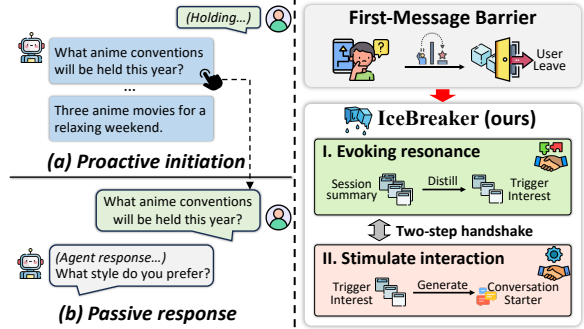


Figure 1: Paradigm comparison between (a) proactive initiation that lowers the first-message barrier by providing personalized starters and (b) passive response that occurs within an ongoing dialogue. **ICEBREAKER** mimics human ice-breaking via a two-step handshake: evoking resonance and stimulating interaction.

the world’s largest conversational agent products, serving hundreds of millions of users. In real-world deployments, the serving paradigm is shifting from passive answering to proactive participation, where the system anticipates needs and helps users move the interaction forward (Deng et al., 2025b). Prior studies on such *in-conversation* proactivity operate after the user has already started the conversation (as depicted in Figure 1(b)), for example by generating follow-up or clarifying questions (Deng et al., 2023), and have been shown to improve session depth (Andukuri et al., 2024; Li et al., 2025b).

In practice, however, a fundamental product bottleneck remains under-explored: the first-message barrier in the *before-conversation* stage. At this cold-start moment, users often have only a vague goal and a narrow sense of the agent’s actionable scope (Norman, 2013; Zamfirescu-Pereira et al., 2023), making it difficult to start the conversation and highlighting the need for proactive guidance. Production statistics also show that this “*first-message barrier*” correlates with a lower conversation start rate: *roughly 20% of users enter the product but leave without starting the conversation*. To address this bottleneck, we study **proactive ini-**

**tiation**, which provides personalized guidance that lowers the cost of conversation initiation, makes the agent’s capabilities immediately actionable, and more importantly, shapes the subsequent trajectory to improve downstream conversation quality.

However, proactive initiation poses challenges beyond standard response generation. First, in the before-conversation stage, inferring user intent without explicit user signals is challenging. Second, user preferences are highly personalized and long-tailed, so a one-size-fits-all alignment objective tends to bias generation toward generic starters that fail to resonate with individuals. To pioneer in this direction, we study proactive initiation as the task of **Conversation Starter Generation**: producing a set of personalized starter questions that guide users to start a conversation. Motivated by how people initiate conversations in cold-start situations, first surfacing a few interests likely to resonate and then phrasing a starter to elicit interaction, we propose **ICEBREAKER**, which operationalizes proactive initiation as a two-step handshake: (i) *evoke resonance*, where Resonance-Aware Interest Distillation (RID) learns a personalized resonance scorer to distill long session summaries into a compact set of trigger interests, with activity-aware gating controlling distillation strength; and (ii) *stimulate interaction*, where Interaction-Oriented Starter Generation (ISG) conditions on distilled interests to generate a small, diverse starter list, warm-started with distilled instruction data and then aligned via list-wise multi-dimensional preference optimization with periodically augmented preference pairs for self-reinforced refinement.

Online A/B tests show a +1.84% user active days increase and +94.25% CTR improvement. ICEBREAKER has been deployed at scale on one of the world’s largest conversational agent products. Our main contributions include:

- **Paradigm shift:** By formalizing Conversation Starter Generation, we move from responsive execution to proactive initiation to address the first-message barrier in real-world applications.
- **ICEBREAKER:** We propose a two-step handshake framework that bridges the resonance gap in cold-start initiation by coupling resonance-aware interest distillation with preference-aligned starter generation.
- **Real-world deployment:** Extensive offline and online experiments demonstrate that ICEBREAKER not only improves conversation ini-

tiation experience but also increases long-term engagement. ICEBREAKER has been deployed at scale, serving hundreds of millions of users.

## 2 Proposed Method: ICEBREAKER

In this section, we present ICEBREAKER, a two-stage framework for proactive initiation under cold-start moments. As illustrated in Figure 2, ICEBREAKER consists of a distiller that captures trigger interests from session summaries (§ 2.2), and a generator that produces a list of conversation starters conditioned on these interests (§ 2.3).

### 2.1 Problem Formulation

We first formalize the **Conversation Starter Generation** task. Given a user’s session summaries, the goal is to generate a small list of starters that maximize subsequent engagement. Let  $\mathcal{H}_u = \{h_1, h_2, \dots, h_T\}$  denote the historical session summaries of user  $u$ , where each  $h_t$  is a short textual summary of the  $t$ -th session. The goal is to learn a policy  $\pi_\theta$  that generates a list of  $K$  conversation starters  $\mathcal{S} = [s_1, s_2, \dots, s_K]$ . The objective is to maximize the expected list-wise utility:

$$\mathcal{S}^* = \arg \max_{\mathcal{S}} \mathbb{E}_{\mathcal{S} \sim \pi_\theta(\cdot | \mathcal{H}_u)} V(\mathcal{S}), \quad (1)$$

where  $V(\mathcal{S})$  denotes the overall value of the list, which depends on both its user-specific utility and the redundancy within the list. In ICEBREAKER, this objective is approximated with a learned interaction-utility ranker and an explicit within-list diversity signal (§ 2.3).

### 2.2 Resonance-Aware Interest Distillation (RID)

As the first handshake for resonance evoking, RID captures trigger interests from session summaries by mimicking how humans initiate conversations: instead of enumerating all possible topics, people surface a small set of personal resonance cues that are most likely to elicit interaction. We implement RID as a resonance-guided distillation pipeline: we first learn a personalized resonance scorer from user actions, and then conduct adaptive gating to produce trigger interests  $\mathcal{I}^*$  from  $\mathcal{H}_u$ .

**Personalized Resonance Scorer** To quantify user-specific resonance potential, we learn a personalized resonance scorer  $s_\phi(\cdot)$  that maps a candidate session summary  $h_t$  of user  $u$  to a scalar score. Following Deng et al. (2025a) and Guo et al. (2025),

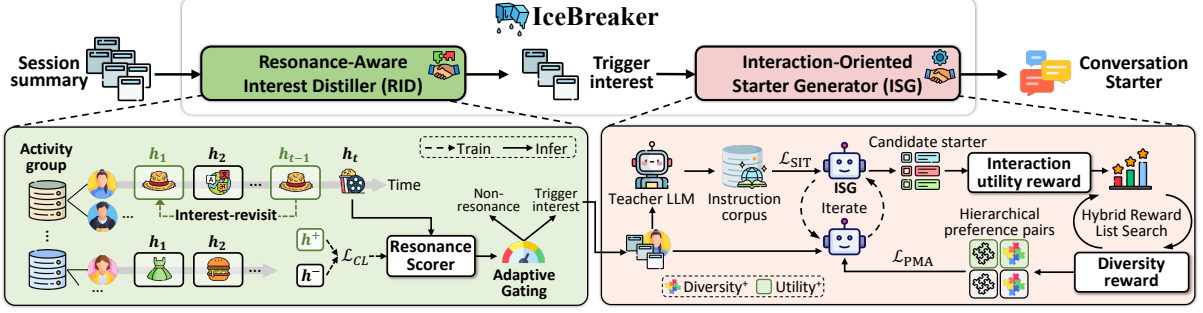


Figure 2: Overall architecture of **ICEBREAKER**. (a) **Resonance-Aware Interest Distillation** identifies trigger interests from session summaries via personalized resonance scoring and adaptive gating. (b) **Interaction-Oriented Starter Generation** produces a small set of first-turn starters conditioned on these interests, and optimizes them for personalized interaction utility and within-list diversity through preference alignment.

the historical session summaries together with user profile (a crowd portrait fitted by the platform) are encoded as user features  $\mathbf{u} = E_{\text{user}}(\mathcal{H}_u, u)$ , session summary is encoded as  $\mathbf{z}_t = E_{\text{text}}(h_t)$ :

$$s_\phi(u, h_t) = \cos(\mathbf{u}, \mathbf{z}_t), \quad (2)$$

where  $\cos(\cdot, \cdot)$  denotes cosine similarity between  $\ell_2$ -normalized embeddings. This design enables efficient retrieval over large summary pools, which is important for real-world production.

For  $s_\phi(\cdot)$  training, an *interest-revisit* signal is leveraged as a proxy for resonance, which provides direct evidence of what a user is willing to revisit. Concretely, we identify resonance by matching a user’s later session summaries to their historical session summaries using a prompting-based approach (detailed in Appendix B.2.1). A session summary is considered as positive  $h^+$  if a later session revisits the same interest. Two types of negatives are constructed: (i) within-user negatives  $\mathcal{N}_{\text{intra}}$ , i.e., other session summaries from the same user that are not matched by any later session, and (ii) cross-user negatives  $\mathcal{N}_{\text{inter}}$ , i.e., session summaries from other users in the same minibatch. Finally,  $\mathcal{N}_u = \mathcal{N}_{\text{intra}} \cup \mathcal{N}_{\text{inter}}$  is used as the negative set.

Let  $\mathcal{D}_{\text{RID}}$  denote the resulting set of training tuples  $(u, h^+, \mathcal{N}_u)$ . We train the encoders in  $s_\phi$  by minimizing the following contrastive loss:

$$\mathcal{L}_{\text{CL}}(\phi) = -\frac{1}{|\mathcal{D}_{\text{RID}}|} \sum_{\mathcal{D}_{\text{RID}}} \log \frac{\exp(s_\phi(u, h^+))}{\sum_h \exp(s_\phi(u, h))}, \quad (3)$$

where  $h \in \{h^+\} \cup \mathcal{N}_u$ ,  $\phi$  denotes the parameters of scorer  $s_\phi$ ,  $|\mathcal{D}_{\text{RID}}|$  is the number of training tuples.

**Adaptive Gating** At inference, we compute scores  $\{s_\phi(u, h_t)\}_{t=1}^T$  and distill trigger interests via an user activity-aware threshold. Concretely,

we map each user to an activity group  $b(u) \in \{1, \dots, B\}$  based on user activity, and apply a group-specific threshold  $\tau_b$  that increases with activity. Formally,  $\tau_u = \tau_{b(u)}$ , where  $\tau_b$  represents the learned score threshold for group  $b$ :

$$\mathcal{I}^* = \{h_t \in \mathcal{H}_u \mid s_\phi(u, h_t) \geq \tau_u\}. \quad (4)$$

The resulting  $\mathcal{I}^*$  serves as resonance-aware cues for the downstream starter generator.

### 2.3 Interaction-Oriented Starter Generator (ISG)

As the second handshake for stimulating reaction, conditioned on the distilled trigger interests  $\mathcal{I}^*$ , ISG generates a small list of first-turn starters  $\mathcal{S} = [s_1, \dots, s_K]$  optimized for both personalized interaction utility and within-list diversity.

We train ISG in two stages: (i) Supervised Interest-Expanded Instruction Tuning (SIT) to warm-start generation quality and topical coverage, and (ii) Personalized Multi-Dimensional Alignment (PMA) to align list-level utility and diversity with user-specific feedback.

#### 2.3.1 Supervised Interest-Expanded Instruction Tuning

We warm-start the generator using an interest-starter instruction corpus distilled from a teacher LLM. Given trigger interests  $\mathcal{I}^*$ , the teacher generates qualified starters under multiple constraints, yielding a coverage-oriented instruction corpus  $\mathcal{D}_{\text{cov}} = \{(\mathcal{I}^*, \mathcal{S})\}$ . Let  $\pi_\theta(\cdot|\mathcal{I}^*)$  denotes the conversation starter generator distribution. The instruction-tuning objective is:

$$\mathcal{L}_{\text{SIT}}(\theta) = -\mathbb{E}_{(\mathcal{I}^*, \mathcal{S}) \sim \mathcal{D}_{\text{cov}}} [\log \pi_\theta(\mathcal{S}|\mathcal{I}^*)]. \quad (5)$$

SIT expands the training distribution beyond the limited observation and provides a stable initialization for subsequent preference optimization.

---

**Algorithm 1: Hybrid Reward List Search**

---

**Input:** user  $u$ , trigger interests  $\mathcal{I}^*$ , generator  $\pi_\theta$ ,  
 $R_{\text{util}}(u, \cdot)$ ,  $R_{\text{div}}(\cdot, \cdot)$ , list length  $K$ , pool size  
 $N$ , beam width  $J$ , expansion size  $M$ ,  
diversity weight  $\lambda$

- 1 Iteratively sample the candidate pool from:  
 $\mathcal{C} \leftarrow \text{SAMPLE}(\pi_\theta(\cdot|\mathcal{I}^*), N)$ ;
- 2 Initialize:  $\mathcal{B} \leftarrow \{\langle \rangle\}$ ;
- 3  $V(\langle \rangle) \leftarrow 0$ ;
- 4 **for**  $\ell \leftarrow 1$  **to**  $K$  **do**
- 5     Expand beam with utility and diversity reward;
- 6      $\mathcal{B}' \leftarrow \emptyset$ ;
- 7     **foreach**  $\mathcal{P} \in \mathcal{B}$  **do**
- 8          $\mathcal{T} \leftarrow \text{TOP}(\mathcal{C} \setminus \mathcal{P}, R_{\text{util}}(u, \cdot), M)$ ;
- 9         **foreach**  $s \in \mathcal{T}$  **do**
- 10              $\mathcal{P}' \leftarrow \mathcal{P} \oplus s$ ;
- 11              $V(\mathcal{P}') \leftarrow$   
               $V(\mathcal{P}) + R_{\text{util}}(u, s) + \lambda \cdot R_{\text{div}}(\mathcal{P}, s)$ ;
- 12             **end**
- 13         **end**
- 14     Keep top- $J$  partial lists:  $\mathcal{B} \leftarrow \text{TOP}(\mathcal{B}', V(\cdot), J)$ ;
- 15 **end**
- 16  $\mathcal{S}^+ \leftarrow \arg \max_{\mathcal{S} \in \mathcal{B}} \text{Score}(\mathcal{S})$
- 17  $\mathcal{S}_{\text{util}}^-, \mathcal{S}_{\text{div}}^-, \mathcal{S}_{\text{jf}}^- \leftarrow \text{SELECTNEGATIVES}(\mathcal{B}, \mathcal{S}^+)$

**Output:**  $\mathcal{S}^+, \mathcal{S}_{\text{util}}^-, \mathcal{S}_{\text{div}}^-, \mathcal{S}_{\text{jf}}^-$

---

### 2.3.2 Personalized Multi-Dimensional Alignment

SIT teaches the generator to produce well-formed starters, but it does not align generation with user-specific interaction utility. To align the generator with personalized feedback under sparse and evolving signals, PMA optimizes the generator with Direct Preference Optimization (DPO) and mines user-conditioned preference supervision tailored to both interaction utility and within-list diversity:

$$\mathcal{L}_{\text{PMA}}(\theta) = -\mathbb{E}_{(\mathcal{S}_w, \mathcal{S}_l) \sim \mathcal{P}} \left[ \log \sigma \left( \beta \left( \log \pi_\theta(\mathcal{S}_w | \mathcal{I}^*) - \log \pi_\theta(\mathcal{S}_l | \mathcal{I}^*) \right) \right) \right], \quad (6)$$

where  $\beta$  is a temperature parameter and  $\mathcal{S}_w \succ \mathcal{S}_l$  denotes list  $\mathcal{S}_w$  is preferred to  $\mathcal{S}_l$  for a given user.

**Hybrid Reward List Search** To obtain such user-conditioned preference pairs, a straightforward approach is to derive them from user feedback data. However, feedback from a specific user is extremely sparse, resulting in an insufficient candidate space that limits both the quantity and diversity of preference supervision. To alleviate this personalized-feedback sparsity, we iteratively construct user-conditioned preference pairs from model-generated candidates during PMA.

Specifically, as shown in Algorithm 1, we sample a candidate pool  $\mathcal{C}$  from the current generator

$\pi_\theta(\cdot|\mathcal{I}^*)$  for each user. Guided by a hybrid reward signal (interaction utility reward and diversity reward; denoted as  $R_{\text{util}}$  and  $R_{\text{div}}$  and detailed below), the algorithm incrementally grows partial lists: at each position, it expands a beam of partial lists with top- $M$  utility-ranked candidates and re-scores each expansion by combining utility with marginal diversity gain. As  $\pi_\theta$  is updated across iterative preference alignment rounds, we re-sample  $\mathcal{C}$  from the latest  $\pi_\theta$  to mine fresh preference pairs.

**Utility–Diversity Reward** We guide list search and preference mining with a hybrid reward signal that jointly captures two dimensions:

- Interaction utility reward  $R_{\text{util}}(u, s)$ : the personalized likelihood that starter  $s$  triggers interaction for user  $u$ .
- Diversity reward  $R_{\text{div}}(\mathcal{P}, s)$ : the marginal diversity gain of adding  $s$  to a partial list  $\mathcal{P}$ , computed from sentence embeddings.

Algorithm 1 maintains a running list value  $V(\mathcal{P})$  by accumulating these stepwise rewards. For interaction utility,  $R_{\text{util}}(u, s)$  captures both short-term and long-term value: (i) predicting whether the user engages with the presented starter, and (ii) predicting the depth of subsequent conversations measured by the number of sessions. The model outputs a point-wise score in  $(0, 1)$ :

$$R_{\text{util}}(u, s) = \text{Proj}([E_{\text{user}}(\mathcal{H}_u, u); E_{\text{text}}(s)]), \quad (7)$$

where  $\text{Proj}(\cdot)$  is the projection layer and  $E(\cdot)$  denotes the text encoder. For a completed list  $\mathcal{S} = [s_1, \dots, s_K]$ , we compute its normalized utility  $\hat{v}_{\text{util}}(\mathcal{S}) = \frac{1}{K} \sum_{k=1}^K R_{\text{util}}(u, s_k)$  and its diversity  $\hat{v}_{\text{div}}(\mathcal{S}) \in [0, 1]$  as the average pairwise cosine dissimilarity among sentence embeddings. We then select the preferred list by:

$$\mathcal{S}^+ = \arg \max_{\mathcal{S}} V(\mathcal{S}), \quad (8)$$
$$V(\mathcal{S}) = \hat{v}_{\text{util}}(\mathcal{S}) + \lambda \cdot \hat{v}_{\text{div}}(\mathcal{S}),$$

where  $\lambda$  controls the utility–diversity trade-off.

Three types of dispreferred lists are constructed for hierarchical supervision:

- Utility negatives ( $\mathcal{S}_{\text{util}}^-$ ): lists with substantially lower  $\hat{v}_{\text{util}}$  than  $\mathcal{S}^+$  while keeping  $\hat{v}_{\text{div}}$  comparable, isolating the signal for user-specific utility.
- Diversity negatives ( $\mathcal{S}_{\text{div}}^-$ ): lists with substantially lower  $\hat{v}_{\text{div}}$  while keeping  $\hat{v}_{\text{util}}$  comparable, providing the signal for within-list non-redundancy.
- Joint-failure negatives ( $\mathcal{S}_{\text{jf}}^-$ ): lists that are poor on both  $\hat{v}_{\text{util}}$  and  $\hat{v}_{\text{div}}$ , providing easy negatives that stabilize preference learning.

Finally,  $\mathcal{S}^- \in \{\mathcal{S}_{\text{util}}^-, \mathcal{S}_{\text{div}}^-, \mathcal{S}_{\text{jf}}^-\}$ , which disentangle utility and diversity trade-offs for DPO training, enabling the generator to improve engagement without collapsing into generic or repetitive starters.

**Self-Reinforced Iterative Optimization** To continuously harvest informative supervision under sparse feedback, we alternate between preference mining and policy optimization by iteratively mining hierarchical preference pairs from the latest  $\pi_\theta$ , augmenting them with teacher-distilled starters, and updating  $\pi_\theta$  on the accumulated preference set. In deployment, we periodically execute this procedure to track drifting user preferences.

### 3 Experiments

We evaluate ICEBREAKER with offline benchmarks and online A/B tests, and conduct ablations to quantify the contribution of each component.

#### 3.1 Experimental Setup

**Datasets** We construct the offline benchmark from anonymized user actions collected in a large-scale conversational system, using activity stratification to ensure balanced representation across engagement levels. For each user, historical session summaries are retrieved to generate candidate starters, which are then evaluated with the interaction utility ranker. Final effectiveness is validated through online A/B tests for more than one month.

**Baselines** We compare ICEBREAKER against two categories of LLM-based generation. For training-free baselines: • **PE** directly prompts the backbone model with session summaries to generate starters; • **PE + RID** prompts the model with distilled trigger interests by Resonance-Aware Interest Distillation. For fine-tuned baselines: • **SFT** fine-tunes the starter generator on interest–starter pairs to improve generation quality and coverage; • **SFT + DPO** further trains the generator with vanilla direct preference optimization based on user interaction signals.

#### 3.2 Offline Evaluation and Model Selection

Offline evaluations compare ICEBREAKER against baselines and determine the optimal backbone for deployment. As shown in Table 1, we evaluate Qwen2.5-7B and Doubao1.5-Lite across two dimensions: (i) **Utility** measures alignment with the deployed learning-to-rank model through user-level ranking consistency (fraction of improved

Method	Utility		Diversity	
	R-User $\uparrow$	R-Score $\uparrow$	Lexical $\uparrow$	Semantic $\uparrow$
Qwen2.5-7B				
PE	–	–	<b>29.03</b>	<b>6.11</b>
PE + RID	+0.69	+0.32	24.12	4.99
SFT	+0.73	+0.39	26.46	5.17
SFT + DPO	<u>+0.75</u>	<u>+0.42</u>	13.75	2.64
<b>ICEBREAKER</b>	<b>+0.82</b>	<b>+0.74</b>	<u>27.74</u>	5.01
Doubao1.5-Lite				
PE	+0.56	+0.08	<b>29.45</b>	<b>6.23</b>
PE + RID	+0.71	+0.38	25.13	4.86
SFT	+0.78	+0.44	<u>28.97</u>	<u>5.59</u>
SFT + DPO	<u>+0.79</u>	<u>+0.52</u>	12.94	2.37
<b>ICEBREAKER</b>	<b>+0.89</b>	<b>+0.80</b>	28.83	5.28

Table 1: Main offline results across backbones ( $\uparrow$  indicates higher is better). Utility reflects relative improvement over PE. The best and second-best results are indicated in **bold** and underlined, respectively.

Method	Active $\uparrow$	Avg.S. $\uparrow$	CTR $\uparrow$	CSR $\uparrow$
PE	-0.01	-0.26	-16.16*	-0.17
SFT	+0.20	+0.33	+6.97*	-0.05
SFT + DPO	+1.16	+0.42*	+56.41*	+0.68
<b>ICEBREAKER</b>	<b>+1.84*</b>	<b>+1.59*</b>	<b>+94.25*</b>	<b>+1.27*</b>

Table 2: Main online A/B test results (relative lifts % over the deployed baseline). We report user active days (Active), average sessions per user (Avg.S.), click-through rate (CTR), and conversation start rate (CSR). The best results are **bolded** and \* indicates  $p < 0.05$ .

users) and score-level lift. (ii) **Diversity** measures diversity using a keyword-based lexical metric and a classifier-based semantic metric. Details are provided in Appendix B.1.

Experimental results show that: (i) **RID improves personalization**. Distilling resonance-triggering interests provides more user-specific conditioning and improves over direct prompting. (ii) **ICEBREAKER balances utility and diversity**. Through utility–diversity preference alignment, it achieves the highest ranking consistency without sacrificing variety or collapsing into repetitive outputs. Doubao-based ICEBREAKER is selected for online A/B testing due to its superior performance.

#### 3.3 Online A/B Tests

We conduct online A/B tests on one of the world’s largest conversational agent products for more than one month. Detailed metric definitions are provided in Appendix B.1. Table 2 shows that: (i) **Breaking the first-message barrier**. ICEBREAKER achieves the largest CSR lift, indicating that more users cross the first-turn initiation threshold. (ii) **Downstream benefits**. Beyond initiation, ICEBREAKER consistently improves Active and

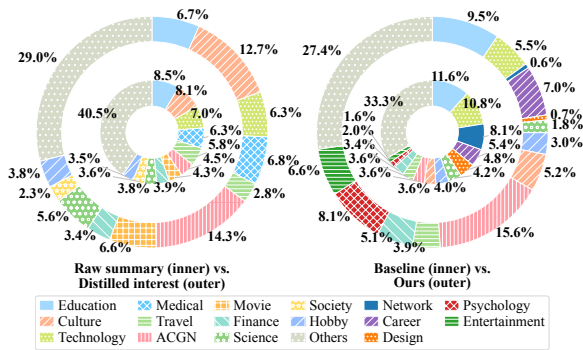


Figure 3: Distribution analysis. **Left:** RID distills away generic topics toward triggering categories. **Right:** starters generated by ICEBREAKER show broader coverage and higher interaction potential.

Avg.S., suggesting sustained engagement after the conversation started.

### 3.4 Further Analysis

**Distribution Analysis** We map each text to a fine-grained topic category with a pre-trained classifier and visualize topic proportions across pipeline stages. In Figure 3, the inner and outer rings show the distributions before and after each step (■ “Others” aggregates categories below the display threshold): (i) **RID denoises and sharpens interests.** Distilled trigger interests down-weight functional and generic head topics and shift the distribution toward resonance-prone, long-tailed cues. (ii) **ISG steers starters toward ice-breaking topics.** Compared with distilled interests, generated starters tilt further toward easy-to-engage, consumption-oriented categories (e.g., ■ “Psychology”, ■ “ACGN”, ■ “Entertainment”), which are more likely to elicit interaction. (iii) **Utility-aligned diversity.** Relative to the deployed baseline, ICEBREAKER achieves broader topical coverage with reduced head dominance while remaining aligned with interaction utility, rather than being diverse but irrelevant.

**Cases Analysis** Table 3 provides an illustrative comparison. Compared with retrieved generic starters, ICEBREAKER generates more personalized, interaction-inviting starters.

## 4 Related Work

Recent advances in LLMs have empowered conversational agents to move beyond passive responding (Li et al., 2025a; Liu et al., 2024; Deng et al., 2025b). Prior work on proactivity mainly studies *in-conversation* behaviors, including follow-up

**Session summaries:** AI trends; Weekend schedule planning; One Piece anime wallpapers; Translating a Latin article on Philosophy; Cosplay beginner tutorial; Top anime opening of 2025; JoJo character analysis; Local LLM deployment

### Retrieved conversation starters (Deployed):

- ① What are the latest LLM papers?
- ② How to spend a fulfilling vacation?
- ③ How to better cosplay female anime characters?
- ④ Recommend some popular anime openings in 2025.

**Distilled trigger interests:** *Anime openings and characters (One Piece, JoJo); cosplay ideas; AI trends and applications*

### Generated conversation starters (ICEBREAKER):

- ① What anime conventions will be held this year?
- ② Which One Piece character is beginner-friendly to cosplay?
- ③ Explain techniques for acceleration of LLM deployment.
- ④ Three anime movies for a relaxing weekend.

Table 3: Case study comparing the deployed retrieval-based starters and our ICEBREAKER. (1) Denoising: RID filters out functional noise in session summaries (e.g., translation requests) and distills trigger interests. (2) Finer-grained personalization beyond retrieval: compared to generic starters, ICEBREAKER generates more personalized starters with finer topical granularity.

question generation (Li et al., 2025b) and topic steering (Rebedea et al., 2024), which depend on dialogue context and explicit intents. For preference alignment, a growing body of work studies preference alignment for LLMs to better reflect human preferences and improve assistance quality, leveraging feedback signals via supervised fine-tuning and preference optimization (Yin et al., 2025; Sun et al., 2024; Wu et al., 2025). However, such general alignment objectives are insufficient for proactive initiation, which requires user-level personalization under sparse feedback. Appendix A provides an extended discussion.

## 5 Conclusion

Conversational agents are becoming daily assistants for billions of users. We highlight proactive initiation as a practical capability for modern LLM chat assistants to overcome the first-message barrier in cold-start moments. To this end, we study Conversation Starter Generation and propose **ICEBREAKER**, which operationalizes proactive initiation as a two-step handshake: Resonance-Aware Interest Distillation for evoking user resonance and Interaction-Oriented Starter Generation for stimulating interaction. Extensive offline evaluations and online A/B tests demonstrate that ICEBREAKER consistently improves user experiences and conversation quality, providing a scalable solution for industrial-grade proactive conversational agents.

## References

- Chinmaya Andukuri, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. 2024. [STar-GATE: Teaching language models to ask clarifying questions](#). In *First Conference on Language Modeling*.
- Jiaxin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui Zhou. 2025a. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965*.
- Yang Deng, Lizi Liao, Liang Chen, Hongru Wang, Wenqiang Lei, and Tat-Seng Chua. 2023. Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10602–10621.
- Yang Deng, Lizi Liao, Wenqiang Lei, Grace Hui Yang, Wai Lam, and Tat-Seng Chua. 2025b. Proactive conversational ai: A comprehensive survey of advancements and opportunities. *ACM Transactions on Information Systems*, 43(3):1–45.
- Xian Guo, Ben Chen, Siyuan Wang, Ying Yang, Chenyi Lei, Yuqing Ding, and Han Li. 2025. One-sug: The unified end-to-end generative framework for e-commerce query suggestion. *arXiv preprint arXiv:2506.06913*.
- Muhammad Usman Hadi, Rizwan Qureshi, Abbas Shah, Muhammad Irfan, Anas Zafar, Muhammad Bilal Shaikh, Naveed Akhtar, Jia Wu, Seyedali Mirjalili, and 1 others. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, and 1 others. 2023. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.
- Han Li, Renwen Zhang, Yi-Chieh Lee, Robert E Kraut, and David C Mohr. 2023. Systematic review and meta-analysis of ai-based conversational agents for promoting mental health and well-being. *NPJ Digital Medicine*, 6(1):236.
- Hao Li, Chenghao Yang, An Zhang, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2025a. [Hello again! LLM-powered personalized agent for long-term dialogue](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5259–5276, Albuquerque, New Mexico. Association for Computational Linguistics.
- Xiaoyu Li, Xiao Li, Li Gao, Yiding Liu, Xiaoyang Wang, Shuaiqiang Wang, Junfeng Wang, and Dawei Yin. 2025b. Proactive guidance of multi-turn conversation in industrial search. *arXiv preprint arXiv:2505.24251*.
- Jiayi Liao, Xiangnan He, Ruobing Xie, Jiancan Wu, Yancheng Yuan, Xingwu Sun, Zhanhui Kang, and Xiang Wang. 2024. [Rosepo: Aligning llm-based recommenders with human values](#). *ArXiv*, abs/2410.12519.
- Na Liu, Liangyu Chen, Xiaoyu Tian, Wei Zou, Kaijiang Chen, and Ming Cui. 2024. From llm to conversational agent: A memory enhanced architecture with fine-tuning of large language models. *arXiv preprint arXiv:2401.02777*.
- Chen Luo, Dimitri Papadimitriou, Hariharan Muralidharan, Dhineshkumar Ramasubbu, Aakash Kolekar, Wenju Xu, Cong Xu, Anirudh Srinivasan, Mukesh Jain, and Qi He. 2025. [Language model alignment for conversational shopping at amazon](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '25*, page 4314–4318, New York, NY, USA. Association for Computing Machinery.
- Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Basic books.
- Weicong Qin, Yi Xu, Weijie Yu, Chenglei Shen, Ming He, Jianping Fan, Xiao Zhang, and Jun Xu. 2025. [MAPS: Motivation-aware personalized search via LLM-driven consultation alignment](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3039–3051, Vienna, Austria. Association for Computational Linguistics.
- Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Xuanhe Zhou, Yufei Huang, Chaojun Xiao, and 1 others. 2024. Tool learning with foundation models. *ACM Computing Surveys*, 57(4):1–40.
- Traian Rebedea, Makesh Sreedhar, Shaona Ghosh, Jiaqi Zeng, and Christopher Parisien. 2024. Canttalkaboutthis: Aligning language models to stay on topic in dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12232–12252.
- Chao Sun, Yaobo Liang, Yaming Yang, Shilin Xu, Tianmeng Yang, and Yunhai Tong. 2024. [Rlrf4rec: Reinforcement learning from recsys feedback for enhanced recommendation reranking](#). *ArXiv*, abs/2410.05939.
- Jianling Wang, Yifan Liu, Yinghao Sun, Xuejian Ma, Yueqi Wang, He Ma, Zhengyang Su, Minmin Chen, Mingyan Gao, Onkar Dalal, Ed H. Chi, Lichan Hong, Ningren Han, and Haokai Lu. 2025. [User feedback alignment for LLM-powered exploration in large-scale recommendation systems](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 996–1003, Vienna, Austria. Association for Computational Linguistics.

Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. [Aligning large language models with human: A survey](#). *Preprint*, arXiv:2307.12966.

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.

Zhongheng Yang, Aijia Sun, Yushang Zhao, Yinuo Yang, Dannier Li, and Chengrui Zhou. 2025. [Rlhf fine-tuning of llms for alignment with implicit user feedback in conversational recommenders](#). *Preprint*, arXiv:2508.05289.

Junhao Yin, Haolin Wang, Peng Bao, Ju Xu, and Yongliang Wang. 2025. [From clicks to preference: A multi-stage alignment framework for generative query suggestion in conversational system](#). *Preprint*, arXiv:2508.15811.

J Diego Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why johnny can't prompt: how non-ai experts try (and fail) to design llm prompts. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, pages 1–21.

Jinhui Zhu, Xiangfeng Luo, Xin Yao, and Xiao Wei. 2026. [Lrsa: Llm-recsys alignment for time-specific next poi recommendation](#). *Information Processing & Management*, 63(2, Part B):104434.

Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Haonan Chen, Zheng Liu, Zhicheng Dou, and Ji-Rong Wen. 2025. Large language models for information retrieval: A survey. *ACM Transactions on Information Systems*, 44(1):1–54.

## A Extended Related Work

### A.1 LLM-based Conversational Agents

Recent advances in LLMs have empowered conversational agents (Li et al., 2025a; Liu et al., 2024) to move beyond passive responding (Deng et al., 2025b). Existing works on proactivity primarily focus on *in-conversation* stages, such as generating follow-up questions (Li et al., 2025b; Yin et al., 2025), clarifying ambiguous queries (Deng et al., 2023), or steering topics (Rebetea et al., 2024). However, these methods rely on ongoing dialogue contexts and explicit user intents, rendering them ineffective for cold-start users who have not yet initiated a conversation. Our work targets the *before-conversation* stage, addressing the critical bottleneck of the first-message barrier.

### A.2 LLM Alignment for User Preference

Aligning LLMs with user preferences has been extensively studied in recommendation (Zhu et al., 2026; Wang et al., 2025; Liao et al., 2024), search (Luo et al., 2025; Qin et al., 2025), and conversational systems (Wu et al., 2025; Wang et al., 2023). Recent approaches leverage human feedback, such as clicks or comparative judgments, to align LLMs with user values and interests through Supervised Fine-Tuning (SFT) (Luo et al., 2025), Reinforcement Learning from Human Feedback (RLHF) (Yin et al., 2025; Yang et al., 2025), or Direct Preference Optimization (DPO) (Sun et al., 2024; Liao et al., 2024). These methods can only align overall platform preferences, not personalized preferences for each user. To address the issue of sparse feedback in personalized preference alignment, we designed a list-wise, multi-objective formulation to bridge the generative decoder with user interests.

## B More Details of Experiment Setups

### B.1 Metrics Calculation

We provide detailed definitions and calculation procedures for the offline and online metrics used throughout our experiments.

#### B.1.1 Offline Metrics

**Diversity** measures the coverage of user interests at both lexical and semantic granularities. **Lexical diversity** employs a TF-IDF-based keyword extraction approach, computing the number of unique high-weight keywords extracted from generated starters. This captures surface-level topical variety through vocabulary distinctiveness. **Semantic diversity** employs a classifier-based approach, where we map generated starters to a hierarchical interest taxonomy using a trained topic classifier and compute the number of distinct interest categories activated. This captures deeper semantic variation beyond surface keywords.

**Utility** Utility evaluates the alignment between generated starters and the deployed online learning-to-rank (LTR) model. Let  $\mathcal{U}$  denote the set of evaluated users. For each user  $u \in \mathcal{U}$ , let  $\mathcal{G}_u$  and  $\mathcal{R}_u$  denote the sets of generative starters and baseline candidates, respectively.  $f(\cdot)$  denotes the scoring function of the deployed learning-to-rank model.

**R-User** measures the fraction of users whose generated starters achieve higher average scores than baseline candidates:

$$\text{R-User} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{I}(\bar{f}(\mathcal{G}_u) > \bar{f}(\mathcal{R}_u)), \quad (9)$$

**R-Score** measures the average LTR score improvement of generative starters over baseline candidates:

$$\text{R-Score} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{\bar{f}(\mathcal{G}_u) - \bar{f}(\mathcal{R}_u)}{\bar{f}(\mathcal{R}_u)}, \quad (10)$$

### B.1.2 Online Metrics

For online A/B tests, all metrics are reported as relative lifts over the deployed baseline. Let  $\mathcal{U}$  denote the set of users exposed to the system during the evaluation period, and  $\mathcal{S}_u$  denote the set of sessions initiated by user  $u \in \mathcal{U}$ . Let  $\mathbb{I}(\cdot)$  be the indicator function.

**Active** (user active days) measures the number of days a user engages with the system within a 7-day window.

**Avg.S.** (average sessions per user) measures the average number of conversation sessions per user:

$$\text{Avg.S.} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} |\mathcal{S}_u|. \quad (11)$$

**CTR** (click-through rate) measures the proportion of recommended starters that receive clicks:

$$\text{CTR} = \frac{N_{\text{clicked}}}{N_{\text{shown}}}, \quad (12)$$

**CSR** (conversation start rate) measures the proportion of active users who have initiated conversations:

$$\text{CSR} = \frac{|\{u \in \mathcal{U} : \text{start}(u) > 0\}|}{|\mathcal{U}|}, \quad (13)$$

where  $\mathcal{U}$  denotes the set of active users during the evaluation period and  $\text{start}(u)$  indicates whether user  $u$  starts a conversation, including entering a conversation via clicking a starter.

## B.2 More Implementation Details

### B.2.1 Prompts

We summarize the prompts used in our pipeline. Table 4 is used to generate a coverage-oriented interest-starter corpus for Supervised Interest-Expanded Instruction Tuning (SIT), as well as for

---

#### Task Objective.

Based on the information provided, generate several questions and corresponding reasoning processes that **conform to common sense, share similar interest subjects, are relevant to the input, but do not repeat specific content.**

---

#### Key Rules.

1. **Length Constraint:** Each question must be strictly limited to 20 characters;
  2. **Number of Questions:** Output 5–20 questions;
  3. **Diversity Requirement:** Cover diverse topic types, no medical topics;
  4. **Long- and Short-term Interest Balance:** Focus on both core and immediate interests;
  5. **Generality:** Generalization and association are required;
  6. **Fine-grained Requirement:** Questions should be specific, not generic;
  7. **User Expertise Matching:** Adjust professionalism based on user interest;
  8. **Factuality:** Subjects must objectively exist;
  9. **Safety:** No sensitive, negative, or invasive topics.
- 

#### Output Format.

```
<think> Reasoning process </think>
<sug> Question 1 && Question 2 && ... && Question n
</sug>
```

---

#### Input.

```
{user_interest}
```

---

Table 4: Prompt used for distilled data generation.

the prompt-based baseline. Table 5 is used to detect interest revisits by matching later session summaries to historical ones, providing samples for training the personalized resonance scorer in RID.

In Table 4, the input {user\_interest} is a set of distilled trigger interests. The model is instructed to output multiple short, diverse, and qualified conversation starters under explicit constraints; we parse the questions from the <sug> field and ignore the intermediate reasoning.

In Table 5, we apply the prompt to pairs of historical and later session summaries from the same user. We treat MATCH pairs as interest-revisit positives for RID training and discard the optional <rationale> in downstream processing.

### B.2.2 Implementation Details

**Model Architecture** We experiment with two backbone models as the base generator  $\pi_\theta$ :

- **Qwen2.5-7B:** A 7-billion-parameter open-source model with 32 transformer layers, 4096 hidden dimensions, and 32 attention heads. The model supports a context window of 32K tokens.
- **Doubao1.5-Lite:** A proprietary lightweight model optimized for low-latency inference with comparable capacity. The model employs effi-

---

**Task Objective.**

Given an earlier session summary and a later session summary from the same user, decide whether the later session revisits the same underlying interest as the earlier one.

---

**Key Rules.**

1. Use only the provided session summaries.
  2. Output MATCH only if both summaries refer to the same specific interest/topic (paraphrases allowed); otherwise output NO\_MATCH.
  3. Shared generic themes (e.g., “music”, “travel”) are insufficient for MATCH unless the specific subject aligns.
- 

**Output Format.**

<label> MATCH or NO\_MATCH </label>

<rationale> One-sentence justification </rationale>

---

**Input.**

[Earlier session summary]: {h\_earlier}

[Later session summary]: {h\_later}

---

Table 5: Prompt used for interest-revisit detection via session-summary matching.

cient attention mechanisms and is trained with multi-task instruction tuning. Model details are available at <https://console.volcengine.com/ark>.

**Deployment Latency** In production deployment, ICEBREAKER meets real-world latency requirements: the median end-to-end latency for generating a full starter list is 2.9 seconds, and the first-token latency is 0.28 seconds. ICEBREAKER has been deployed at scale in the product.

## C Additional Experimental Results

We conduct extensive studies to quantify the contribution of key design choices in our system. All experiments are evaluated using the same offline metrics as in Section 3.2. The results validate that our method designs are complementary and jointly necessary to improve both utility and diversity.

### C.1 Analysis on Resonance-Aware Interest Distillation

The RID component is designed to distill high-resonance trigger interests. We further investigate its effectiveness through a two-step evaluation: (1) a direct quality assessment of the distilled interests, and (2) their downstream impact on starter generation. We compare RID against two baselines: (i) **Non-Distillation**, which feeds all historical interests into the generator without filtering; and (ii) **Rule-based**, which selects interests based on static heuristics such as frequency of occurrence and recency within the user’s history summary.

**Interest Quality Evaluation** To quantify the quality of distilled interests directly, we employ both LLM-as-judge and human evaluations in a pairwise comparison setting. Given a user’s raw interests (Non-Distillation), both LLMs (DeepSeek-R1-671B and Doubao-Seed-1.8-High) and human annotators are asked to judge which distillation strategy produces interests that are more likely to serve as high-resonance triggers for a future conversation. As shown in Table 6, **RID achieves a dominant win rate over rule-based distillation across all evaluators**. This confirms that static heuristics often surface generic or repetitive interests, while RID accurately identifies specific, high-reward points aligned with user resonance.

**Effect on Starters Generation** We further investigate how the quality of distilled interests translates into final generation performance using the offline evaluation metrics described in §3.2. The non-distillation setting, while having the highest diversity due to its broad input, underperforms at utility, indicating that a lack of distillation introduces significant noise. Rule-based retrieval improves utility but is still outperformed by RID. Our **reward-based strategy (RID) achieves the best utility while maintaining good diversity**, demonstrating that by accurately distilling high-resonance seeds, RID effectively guides the generator to produce conversation starters that are both personalized and engaging.

### C.2 Analysis on Self-Reinforced Iterative Optimization

During the self-reinforced iterative optimization process, the model is continually updated using time-sliced data to harvest the latest supervision signals. Each slice corresponds to user interests collected within a fixed window (e.g., one week) before a specific date, paired with candidates scored by the ranking model deployed at that time. During evaluation, all models are assessed using the latest ranking model and the same user inputs, ensuring that performance reflects adaptability to the current online environment.

As shown in Table 7, **generators trained on more recent data slices consistently achieve higher utility**, confirming that iterative optimization enables the generator to better adapt to evolving ranking preferences and user behavior. However, we also observe a gradual decrease in diversity as the training slice becomes more recent. We at-

Distillation Strategy	Interest Quality			Utility		Diversity	
	DeepSeek-R1	Doubao-1.8	Human	R-User ↑	R-Score ↑	Lexical ↑	Semantic ↑
Non-Distillation	–	–	–	–	–	<b>29.45</b>	<b>6.23</b>
Rule-based	0.07	0.08	0.14	+0.70	+0.35	25.98	4.87
Reward-based (RID)	<b>0.93</b>	<b>0.92</b>	<b>0.86</b>	<b>+0.78</b>	<b>+0.58</b>	28.83	5.28

Table 6: Evaluation of interest distillation strategies. **Interest Quality** reports the pairwise win rate from LLM-as-judge (DeepSeek-R1-671B and Doubao-Seed-1.8-High) and human annotators. **Utility** and **Diversity** evaluate the downstream generation quality.

Data Slice	Utility		Diversity	
	R-User ↑	R-Score ↑	Lexical ↑	Semantic ↑
Early	+0.79	+0.40	30.76	<b>6.25</b>
Middle	+0.85	+0.70	<b>30.82</b>	6.14
Latest	<b>+0.89</b>	<b>+0.80</b>	28.83	5.28

Table 7: Effect of Self-Reinforced Iterative Optimization using different training periods. All models are evaluated with the latest ranking model and compared against the prompt engineering model as the baseline.

tribute this phenomenon to the increased concentration of supervision signals from the current ranking model, which tends to emphasize a narrower set of high-reward interests.

### C.3 Analysis on Training Data Usage

To understand how training data usage affects generation quality, we ablate the data composition in both the SIT and PMA stages while keeping the backbone and evaluation protocol unchanged. (i) For SIT, we compare our distilled-only setting that instruction-tunes the generator on teacher-distilled interest-starter pairs (ICEBREAKER) against a variant that additionally mixes in self-reinforced data (**w/ Self-Reinforced**). (ii) For PMA, we compare our default preference-pair construction that mixes distilled starters with self-reinforced candidates (ICEBREAKER) against removing the distilled starters (**w/o Distilled**).

Results shown in Table 8 validate our design: SIT should prioritize coverage expansion through teacher distillation to establish a broad sample space and raise the model’s capability ceiling, while PMA refines utility-diversity trade-offs through personalized preference learning.

**SIT Data Composition** Using distilled supervision for SIT yields the best performance. In contrast, adding self-reinforced data at the SIT stage (w/ Self-Reinforced) degrades both utility and diversity, suggesting that early instruction tuning ben-

Data Source	Utility		Diversity	
	R-User ↑	R-Score ↑	Lexical ↑	Semantic ↑
SIT Data Composition				
<b>ICEBREAKER</b>	<b>+0.89</b>	<b>+0.80</b>	<b>28.83</b>	<b>5.28</b>
w/ Self-Reinforced	+0.83	+0.32	19.6	4.79
PMA Data Composition				
<b>ICEBREAKER</b>	<b>+0.89</b>	<b>+0.80</b>	<b>28.83</b>	<b>5.28</b>
w/o Distilled	+0.85	+0.34	16.12	3.20

Table 8: Analysis of training data composition for SIT and PMA stages. The best results of each stage are indicated in **bold**.

Negative Sampling	Utility		Diversity	
	R-User ↑	R-Score ↑	Lexical ↑	Semantic ↑
<b>Utility-Only</b>	<b>+0.85</b>	<b>+0.36</b>	15.18	2.94
<b>Diversity-Only</b>	+0.58	+0.08	<b>33.34</b>	<b>6.50</b>
<b>Joint-failure</b>	+0.72	+0.20	21.3	4.59
<b>Hierarchical</b>	<u>+0.84</u>	<u>+0.35</u>	<u>28.83</u>	<u>5.28</u>

Table 9: Ablation study of hierarchical negative sampling strategies in the PMA stage. All results are evaluated using the same ranking model and compared against the PE baseline. The best result is shown in **bold** and the second best is underlined.

efits more from clean, high-coverage distilled signals than from on-policy self-generated data.

**PMA Data Composition** For PMA, mixing distilled starters with self-reinforced candidates achieves the best results. Removing distilled starters (w/o Distilled) leads to clear drops in both utility and diversity, indicating that distilled data provides complementary coverage that improves preference learning and avoids collapsing to a narrow set of repetitive patterns.

### C.4 Analysis on Hierarchical Negative Sampling

We ablate the hierarchical negative sampling design in the DPO stage by selectively using differ-

ent types of negative preference pairs. As shown in Table 9, only using utility negatives improves utility but significantly harms diversity, while diversity negatives enhance diversity with limited gains in ranking-related metrics. Simpler negatives like joint-failure negatives provide weak but stable supervision across all dimensions. In contrast, the hierarchical combination of all negative types achieves the most balanced performance, indicating that structured negative sampling is crucial for jointly optimizing engagement and diversity.

## D Limitations

While ICEBREAKER demonstrates strong performance in generating personalized conversation starters, a few limitations remain. (1) **Hallucination**: Like other generative models, ICEBREAKER may occasionally produce starters that contain factually incorrect information or hallucinated content. In deployment, we mitigate this risk through safety classifiers that detect potentially problematic content. However, fully eliminating hallucination remains an open challenge, particularly for long-tail interests where training supervision is sparse. (2) **Cold-Start for New Users**: The Resonance-Aware Interest Distillation (RID) component requires historical session summaries to identify trigger interests effectively. Future work could explore incorporating zero-shot or few-shot learning techniques to better handle cold-start scenarios.