Query-Focused Individual Simulation with Progressive Persona Completion

Weiwen Su[†]
The University of Tokyo
su-w@tkl.iis.u-tokyo.ac.jp

Naoki Yoshinaga and Masashi Toyoda
Institute of Industrial Science,
The University of Tokyo
[†] Institute for Digital Observatory,
The University of Tokyo
{ynaga,toyoda}@iis.u-tokyo.ac.jp

Abstract

Large language models (LLMs) enable simulating individual responses from persona information, supporting applications such as opinion elicitation and virtual character creation. However, existing approaches typically assume rich persona profiles, which are often unavailable in practice. In this work, motivated by recent findings that LLMs can identify query-relevant persona dimensions (*e.g.*, *whether a user is price-sensitive*), we study query-focused individual simulation under cold-start settings, where relevant persona information is identified and requested on demand for each query. To solve this task while minimizing the number of persona requests, we explore a progressive method that iteratively predicts the most critical relevant persona dimension and uses self-reported confidence as a stopping signal to determine when sufficient information has been collected. Experiments on two dialogue datasets[1] show that this query-driven paradigm achieves simulation performance comparable to approaches that rely on rich persona information extracted from dialogue history, using only a few persona dimensions (up to five per query), and this number is further reduced by our progressive method while maintaining or improving simulation quality.

1 Introduction

Large language models (LLMs) have enabled us to simulate human behavior (Park et al., 2023, 2024; Zhang et al., 2025) based on textual descriptions of individuals. Such human behavior simulation capabilities enable useful applications such as opinion elicitation (Chuang et al., 2024), medical personnel training (Du et al., 2025), and realistic virtual character creation (Wang et al., 2024). By conditioning response generation on persona information, LLM-based simulators can move beyond generic responses and better capture individual differences in beliefs, values, and preferences.

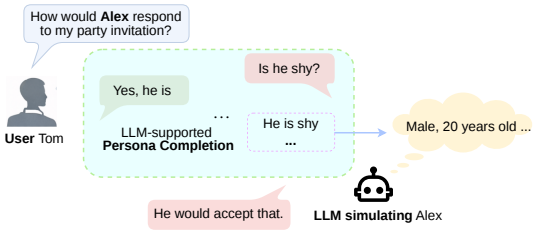


Figure 1: Query-focused individual simulation in cold-start settings: starting from a basic persona, the LLM infers relevant persona dimensions, requests their values, and complements the persona to generate a response.

Previous studies have primarily relied on rich persona information obtained from encyclopedic biographies (Shao et al., 2023) or interviews and questionnaires (Park et al., 2024; Ran et al., 2024; Ng et al., 2024), enabling faithful simulation of public figures. However, in real-world scenarios, the target individuals are often ordinary users whose personas are only partially observable. Preparing extensive persona profiles in advance is costly or even infeasible, making such methods difficult to apply in practice. Meanwhile, recent work (Su et al., 2026) reported that LLMs can identify, for a given query, missing but relevant persona dimensions (*e.g.*, *whether a user is price-sensitive*), pointing to the feasibility of simulation in *cold-start settings* that request query-relevant persona values on demand. However, key questions remain: how to acquire the values of such persona dimensions, how to recognize the sufficiency of acquired persona values, and how to effectively request minimum yet sufficient persona values.

In this study, we thus study **query-focused individual simulation** in cold-start settings (Figure 1), where only basic persona information is available in advance, and additional information is acquired on demand for each query. We assume that the user

[1] https://github.com/NioHww/PICQ-simulation

requesting the simulation can provide persona information about the target individual when prompted. Since this setup involves collecting persona information from the user, we aim to minimize the number of persona requests. In addition to a basic approach that requests values for a fixed number of relevant persona dimensions at once, we explore a progressive approach that identifies and acquires the most critical missing dimension at each step. A key challenge is determining when sufficient persona information has been collected. To address this, we use the model's self-reported confidence as a dynamic stopping signal. By monitoring confidence across iterations, the approach adaptively decides when to stop acquisition, avoiding unnecessary requests while maintaining simulation quality.

We evaluate our methods on the queries in the PICQ-drama dataset (Su et al., 2026) and those we newly extracted from the Nagoya University Conversation Corpus (Fujimura et al., 2012) across multiple simulator LLMs (OpenAI, 2024; Qwen-Team, 2025) and baselines, including context-only simulation and query-agnostic persona.[2] Experimental results show that our progressive persona completion achieves comparable simulation quality while substantially minimizing persona acquisition cost. These results suggest that effective simulation can be achieved with minimal, query-relevant persona information rather than exhaustive profiles.

Our contributions are as follows:

- We formulate query-focused individual simulation in a cold-start setting, where persona information is acquired on demand (§ 3).
- We propose a progressive persona acquisition method with confidence-based stopping to minimize acquisition cost (§ 4).
- We confirmed that a few query-focused personas can match or outperform large sets of query-agnostic personas in simulation (§ 5).

2 Related Work

This section reviews prior studies on conversational user simulation, which focus on generating persona-conditioned responses in a dialogue. In what follows, we first review persona-conditioned response generation and then discuss methods for preparing persona information for simulation.

[2]To promote reproducibility, we use an LLM-based agent with the dialogue history as a proxy for a user who provides persona information about the simulation target.

Persona-conditioned Response Generation Initial studies focus on the feasibility of response generation conditioned on short persona descriptions (Sato et al., 2017; Zhang et al., 2018), while later studies (Gao et al., 2023) enhance personas using commonsense knowledge. Subsequent work leverages role-playing, memory, and interaction histories to steer agent behavior (Park et al., 2023; Shao et al., 2023). Recent work further explores task-specific modeling, such as structured belief dependencies for opinion alignment (Chuang et al., 2024) and cognitive prototypes for student simulation (Wu et al., 2025). Most approaches, however, rely on rich, predefined, query-agnostic persona information, which limits their applicability to simulating ordinary individuals.

Persona Construction The common sources of persona information for simulation are encyclopedic web data (Shao et al., 2023) or one-shot interviews and questionnaires (Park et al., 2024; Ng et al., 2024; Ran et al., 2024). Recent efforts, such as Persona Hub (Ge et al., 2025) and population-aligned generation (Hu et al., 2025), synthesize persona information to improve coverage and demographic realism; however, they do not support simulating responses of specific real individuals. Moreover, these methods determine persona elements *before* any concrete query is given and are query-agnostic. Li et al. (2025) noted that unverified static personas often contain hallucinations or biases, suggesting the need for verification-aware and relevance-driven acquisition.

Query-Focused Relevant Persona Prediction Su et al. (2026) introduced the PICQ-drama dataset, which contains persona-influenced choice questions, and thereby demonstrated that recent LLMs can identify necessary yet missing persona dimensions for a given query. However, it is still unclear how we obtain and utilize the values of predicted relevant persona dimensions for downstream simulation tasks.

Building upon the work by Su et al. (2026), this study explores a pipeline task comprising dimension prediction, persona acquisition, and simulation. Unlike the above static construction, we model persona construction as a coupled and adaptive process driven by the evolving simulation state. The progressive method supplements one dimension at a time until it acquires only the minimum sufficient persona, bridging the prior prediction-only method and practical individual simulation.

3 Query-focused Individual Simulation

In this study, we investigate query-focused individual simulation, where the goal is to simulate how a specific individual would respond to a given query without assuming rich persona information in advance. As stated in § 2, existing persona-based simulation approaches typically rely on a fixed set of persona values collected beforehand. However, different queries often depend on different persona dimensions, suggesting that accurately simulating diverse queries may require extensive persona information. In practice, assuming access to such a comprehensive persona is often unrealistic, especially at the individual level. This mismatch can lead to either unnecessary information acquisition or insufficient persona coverage when using static persona collections. To address this issue, we propose a formulation in which persona information is acquired *on demand*, conditioned on the query.

3.1 Problem Formulation

We consider two persons involved in the individual simulation: a user A and a target individual B, where A asks an LLM to simulate B's response to a query (User Tom for A and Alex for B in Figure 1). We assume that A is familiar with B and can provide B's persona information if requested, since B's persona is typically unavailable in public resources. This setting arises in scenarios such as decision support, opinion elicitation, and character-driven content creation.

A central difficulty is that the persona dimensions required for faithful simulation vary across queries (Hu and Collier, 2024). Here, a *persona dimension* refers to an individual attribute expressed as a verifiable question (*e.g.*, *whether he is shy*) (Su et al., 2026). In a query-focused individual simulation, the premise is not to assume a complete set of persona attributes in advance, but to identify and acquire only those dimensions that are necessary for answering a given query. Accordingly, we model each persona dimension as an atomic attribute d, whose value may be known or unknown. Unknown values correspond to missing persona information. Instead of specifying a complete persona upfront, relevant persona dimensions can be elicited from A on demand. Since acquiring persona information may incur effort for A to recall and provide personas during interaction, it is desirable to minimize the number of requested dimensions and focus only on those that are necessary for the given query.

Ideally, persona information about B would be provided directly by the real user A. However, to promote the reproducibility of our experiments, we adopt an LLM-based agent as a controlled proxy for user A in our experiments. This choice is supported by prior findings (Yuan et al., 2024), which show that LLMs can extract persona attributes of fictional or narrative characters from textual context. While our setting differs in granularity and the use of contextual evidence, these results provide methodological grounding for employing LLMs to infer persona information from dialogue history. The proxy agent is given the dialogue history between A and B, which serves as contextual evidence of their past interactions.

Dataset For the main evaluation (§ 5.2.1), we use the PICQ-drama dataset (Su et al., 2026) derived from TVShowGuess scripts (Sang et al., 2022), which contains manually annotated, context-aware choice QA pairs with relevant but missing persona dimensions. Refer to Table 3 and Table 12 in Appendix A for examples of QA pairs. In each instance, the character pair serves as the proxy user A and the simulation target B. As the source for the LLM proxy to acquire persona values, all dialogue segments in which both characters appear are extracted, temporally ordered, and concatenated with scene timestamps to form a dialogue history. For each query, only the dialogue history preceding the query is used to acquire persona values, reflecting the information available to A at that time.

Settings Formally, the missing persona predictor and simulator take as input a query q requiring B's response, and a dialogue context c consisting of up to 20 turns of interaction between A and B. For simulation, the missing persona predictor may request additional persona information from A (an LLM-based proxy conditioned on dialogue history H in this study). The simulator then generates a simulated response generated from the acquired persona information, dialogue context, and query. We make two practical assumptions: (1) Persona values for B are provided via an LLM-based proxy for A. (2) Following prior work (Su et al., 2026), we initialize missing persona dimension prediction with a basic persona (gender, age, and basic relationship between participants) to avoid inefficient warm-up on trivial attributes.

The evaluation focuses on two aspects: simulation quality and persona acquisition cost. The objective is to achieve high simulation quality while

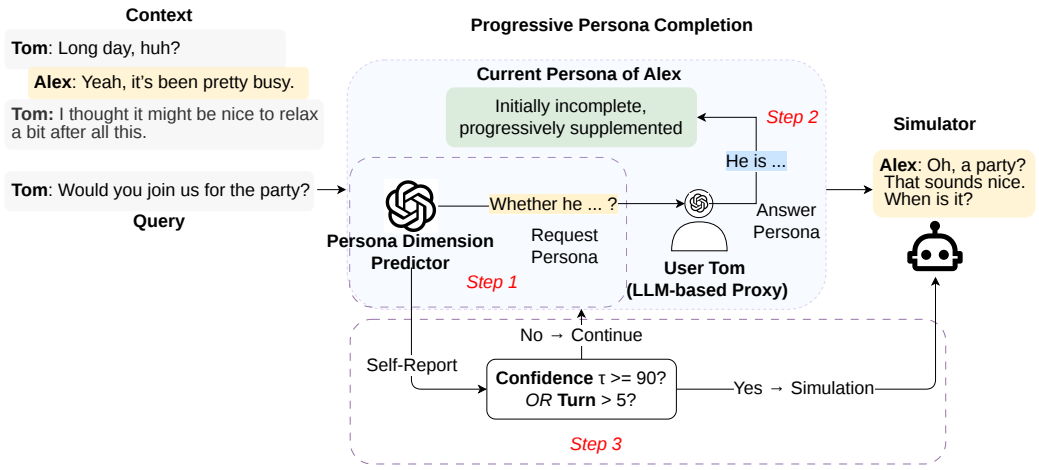


Figure 2: Overview of the progressive persona completion for query-focused individual simulation. The predictor identifies the most influential missing persona dimension conditioned on the dialogue context, query, $\mathcal{P}_{\text{known}}$, and $\mathcal{P}_{\text{banned}}$. The LLM proxy for the user who requests the simulation (here, Tom) answers the requested dimension, and the persona of the target (here, Alex) is updated progressively. A self-reported confidence monitor sufficiency and triggers stopping; simulation starts once confidence exceeds the threshold or the maximum iteration limit is reached.

minimizing the number of persona requests. We operationalize these aspects using concrete evaluation metrics, described in detail in § 5.1.

3.2 Research Questions

The above formulation of query-focused persona acquisition raises two research questions.

How to recognize the persona sufficiency? Although Su et al. (2026) identified a fixed number of relevant persona dimensions, such a query-agnostic upper bound may lead to either unnecessary information acquisition or insufficient persona coverage for individual queries. We need a method that recognizes the sufficiency of persona for faithful simulation.

How to effectively request persona values? We aim to reduce the number of persona requests for efficient simulation. Two challenges arise: the interdependencies among persona dimensions and the difficulty of acquiring certain values. We should avoid requesting values for overlapping dimensions while considering their varying acquisition difficulty.

These questions motivate the need for an adaptive, interactive, and progressive persona acquisition strategy. We will explore a practical solution in § 4.1 to address these questions and compare the solution with Su et al. (2026)'s approach that identifies the fixed number of missing persona dimensions for a given query.

4 Progressive Persona Completion

To address the questions mentioned in § 3.2, we frame persona acquisition as a closed-loop, progressive process that tightly interleaves persona prediction, acquisition, and stopping. Rather than predicting multiple persona dimensions simultaneously (Su et al., 2026), we propose a method that iteratively identifies and acquires the most relevant missing persona dimension, conditioned on the currently known persona, as shown in Figure 2.

Given the dialogue context c and the query q, at each iteration the method updates the known persona set $\mathcal{P}_{\text{known}}$ and the unavailable persona dimensions $\mathcal{P}_{\text{banned}}$. Each iteration consists of three steps: Step 1, persona dimension prediction; Step 2, persona acquisition from user A (or its LLM-based proxy); and Step 3, confidence-based stopping.

This progressive design allows each acquisition decision to condition on $\mathcal{P}_{\text{known}}$ and $\mathcal{P}_{\text{banned}}$, reducing dependency errors across persona dimensions and preventing stagnation in the acquisition process. The stopping mechanism determines whether sufficient persona information has been acquired, avoiding unexpected acquisition costs.

4.1 Step 1: Persona Dimension Prediction

Our method predicts only the most relevant missing persona dimension at each iteration, while requesting and acquiring their values. Each prediction is thus conditioned on verified persona information and on unavailable persona dimensions, mitigat-

ing the dependency and cascading error issues described in § 3.2 while avoiding repeated prediction.

Following Su et al. (2026), we adopt a two-level representation of persona dimension to balance structural consistency and informational flexibility, where each persona dimension has a high-level category (*e.g.*, personality traits, preferences, beliefs) and a natural-language description. We also employ the multi-task prompting strategy proposed in prior work to jointly encourage relevance and interpretability in persona dimension prediction.

To improve persona acquisition rates, we augment each predicted persona dimension with two auxiliary probes that describe observable behavioral manifestations of the same underlying persona dimension. These probes are generated by conditioning on both the predicted dimension and the currently known persona. For example, a dimension like *whether the target is shy* can include probes such as *whether they avoid initiating conversations*. These probes serve as concrete, behavior-oriented cues that are easier for the proxy agent to assess based on dialogue evidence. Refer to Table 6 and Table 7 in Appendix A.1 for the prompts for generating and leveraging probes, respectively.

4.2 Step 2: Persona Acquisition

Once a persona dimension is predicted, its value is acquired by querying the LLM proxy of A. We employ an LLM with the dialogue history as the proxy agent for persona acquisition. The proxy is instructed to consider the target dimension along with its auxiliary probes, and then infer the corresponding value based on the interaction history between A and B; refer to Table 7 in Appendix A.1 for the prompt. After acquiring the value $v^{(t)}$ for the predicted dimension $p^{(t)}$, the known persona is updated as $\mathcal{P}_{\text{known}}^{(t+1)} = \mathcal{P}_{\text{known}}^{(t)} \cup (p^{(t)}, v^{(t)})$. If the value cannot be reliably inferred, the dimension is marked as unavailable, and the set of banned dimensions is updated as $\mathcal{P}_{\text{banned}}^{(t+1)} = \mathcal{P}_{\text{banned}}^{(t)} \cup p^{(t)}$.

4.3 Step 3: Confidence-Based Stopping

One key issue in progressive persona completion is determining when enough persona information has been collected. To address this, we introduce a dynamic stopping mechanism based on the model's self-reported confidence, with a fixed acquisition budget to prevent an infinite loop (*e.g.*, when decisive persona dimensions are unavailable). Prior work (Tian et al., 2023) has shown that eliciting confidence from language models can efficiently estimate uncertainty, and chain-of-thought prompting can reduce overconfidence by encouraging reflective reasoning (Xiong et al., 2024). Motivated by these studies, we design a self-report confidence prompt that asks predictors to assess whether the known persona information is sufficient for a reliable response and whether additional information would change the response. Refer to Table 8 in Appendix A.1 for the prompt.

We treat the reported confidence not as a calibrated probability, but as a relative sufficiency signal indicating whether additional persona information would likely improve simulation quality. By monitoring confidence across iterations, the method adaptively determines when to stop persona acquisition, reducing unnecessary requests and minimizing the cognitive burden on users.

5 Experiments

We evaluate our method on the query-focused individual simulation task from § 3 in dialogue-based settings. For each test instance, the simulator is given a dialogue context, a query, and persona information about the target individual B, to generate a response simulating B's. Initially, B's persona contains only age, gender, and basic relationship with A, which is gradually expanded during simulation. Evaluation focuses on the quality of simulated responses and the efficiency of persona acquisition.

5.1 Setup

Models We use two strong open- and closed-source LLMs, GPT-4.1 and Qwen3-32B,[3] to implement missing persona predictors (§ 4.1) and simulators. Specifically, for the Qwen3-32B model, we enable its thinking mode only when reporting self-confidence to mitigate overconfidence. The proxy user agent for persona acquisition is implemented using GPT-5-mini. We employ a widely adopted role-playing prompting (RPP) (Kong et al., 2024) as the simulation strategy. RPP includes a self-confirmation step, where the model acknowledges and internalizes the given persona before generating a response. Refer to Table 10 in Appendix A.1 for the RPP-style prompt. Specifically, the persona values are transformed into second-person expressions for response simulation. The inputs are anonymized to prevent answer leakage to the evaluated models, following Sang et al. (2022).

[3] https://huggingface.co/Qwen/Qwen3-32B

Persona information used for simulation	Simulation Quality (C_{choice})		Persona Efficiency		
	GPT-4.1	Qwen3-32B	#words ↓	#requests ↓	%acquired
No Persona (context only)	0.595	0.558	n/a	n/a	n/a
Query-agnostic Persona	0.642	0.575	533	n/a	n/a
Query-focused Persona (one-shot, GPT-4.1)	0.618	0.585	34	3.85	61.3%
(one-shot, Qwen3-32B)	0.635	0.595	33	4.99	62.5%
(progressive, GPT-4.1)	0.619	0.580	27	2.98	65.7%
(progressive, Qwen3-32B)	0.640	0.597	25	2.35	65.5%
(gold)	0.640	0.600	26	3.46	64.8%
RAG	0.617	0.562	n/a	n/a	n/a

Table 1: Results on PICQ-drama for simulating choice-making responses based on different persona information.

Personas setup We compare the following persona dimensions for simulation, whose values are derived from the dialogue history by GPT-5-mini.

No Persona generates responses using only the dialogue context and query.

Query-agnostic Persona uses persona values derived from the dialogue history without conditioning on the query, using a 1000-word limit. Refer to Table 9 in Appendix A.1 for the prompt.

Query-focused Persona (gold) uses up to five persona dimensions annotated by humans.

Query-focused Persona (one-shot) uses up to five persona dimensions predicted at once.

Query-focused Persona (progressive) uses persona dimensions predicted progressively, up to five (§ 4.1). The process terminates early if self-reported confidence exceeds 90 (on a 0–100 scale), as specified in the prompt in Table 8 (Appendix A.1).

We also evaluate **RAG** (retrieval-augmented generation) that uses the same dialogue history for retrieval. We retrieve the top-3 dialogue segments using Sentence-BERT (Reimers and Gurevych, 2019) embedding similarity with respect to the query, and condition the simulator on the retrieved segments, the dialogue context, and the query. Note that RAG is not universally applicable in our task since dialogue history may not always be available (*e.g.*, when persona values are provided by a human).

Metrics We evaluate the simulators in terms of simulation quality and persona efficiency. For quality, we report choice consistency. For efficiency, we report average persona word count, number of persona requests, and persona acquisition rate.

Choice Consistency (C_{choice}) measures the alignment between the generated and gold responses in terms of the expressed choice. While natural language inference (NLI) models are commonly used for sentence-level consistency, we found them insufficient for capturing fine-grained choice alignment. Therefore, we use GPT-4.1 to score consistency between the gold and generated responses, using the manually summarized choice situation provided with the PICQ-drama dataset (Su et al., 2026) as context. The consistency score is on a 2-point Likert scale: **0** (inconsistent) and **1** (consistent). The prompted GPT-4.1 is reported to achieve a Cohen's κ of 0.738 with two human annotators (the first author and a graduate student), whose inter-annotator agreement is $\kappa = 0.869$. The evaluation prompt is provided in Appendix A.1.

Average persona word Count (#words) indicates the average number of words used to describe persona information per simulation instance, reflecting the user's writing effort.

Average number of persona requests (#requests) indicates the average number of relevant persona dimensions queried to the proxy per simulation instance, reflecting the user's cognitive effort.

Persona acquisition rate (%acquired) is defined as the proportion of predicted persona dimensions successfully obtained.

Together, these metrics capture the trade-off between simulation quality and the persona efficiency, enabling a comprehensive evaluation of the query-focused individual simulation.

5.2 Results and Analysis

5.2.1 Main Results

Table 1 presents the main results of query-focused individual simulation on the PICQ-drama dataset across different persona acquisition strategies and

Persona information used for simulation	C_{choice}
No Persona (context only)	0.580
Query-focused Persona (one-shot, Qwen3-32B)	0.645
(progressive, Qwen3-32B)	0.650

Table 2: Human evaluation results on PICQ-drama for simulating choice-making responses conditioned on three selected settings using the GPT-4.1 simulator.

simulator models. We can see that incorporating persona information based on LLM-identified dimensions consistently improves performance over No Persona baseline. RAG improves over No Persona baselines and is comparable to GPT-4.1-based predictor and simulator settings, but under other configurations, it remains inferior to methods that explicitly acquire persona information.

Our progressive persona completion achieves comparable simulation quality to one-shot acquisition while requiring fewer persona words and requests. For example, with the GPT-4.1 simulator, our approach reaches C_{choice} of 0.640 using only 25 words and 2.35 requests on average. Both GPT-4.1 and Qwen3-32B predictors show that progressive completion reduces redundant acquisition by conditioning on previously verified persona information, highlighting its efficiency over one-shot strategies. This is further supported by human evaluation results, as shown in Table 2. While human-annotated persona dimensions yield the highest performance, our progressive approach rivals this upper bound with a similar persona cost, showing the feasibility of automated persona completion. In contrast, query-agnostic acquisition requires over 500 words on average for marginal gains, emphasizing the importance of query-focused persona collection.

Table 3 shows a representative example. After acquiring a critical persona dimension (*e.g.*, *he believes it's important to maintain professional dignity in public appearances*), the generated response becomes more consistent with the gold answer. Refer to Table 12 for another example of simulation.

Overall, our query-focused persona completion greatly reduces the amount of persona information for simulation, and progressive completion further improves efficiency while maintaining or even enhancing simulation quality. In what follows, we provide additional experiments on real conversation in Japanese (§ 5.2.2) and examine the impact of confidence-based stopping (§ 5.2.3) and the persona acquisition dynamics (§ 5.2.4).

Dialogue context:
Frasier: ...I just don't think I can go through with it, I'm afraid that I'm compromising my integrity as a psychiatrist.
Niles: Why do you feel like that?
Frasier: ...I need your guidance.
...
Frasier: Niles, will you just tell me, would you do it?

Gold Response:
Niles: Oh, certainly not, I'm a respected psychiatrist.

No Persona (Context only):
Niles: Of course I would do it. Look, Frasier, if you're going to be a psychiatrist, you need to understand people...

Query-focused Persona (progressive, Qwen3-32B):
Niles: No, Frasier, I wouldn't do it. Not if it makes you question your integrity...
Persona (acquired): he believes it's important to maintain professional dignity in public appearances

Table 3: Example of simulated choice-making responses in PICQ-drama using the Qwen3-32B simulator.

5.2.2 Results on Real Japanese Conversation

To evaluate the generalizability of our methods to real-world interactions, we conduct additional simulation experiments using the Nagoya University Conversation Corpus (NUCC) (Fujimura et al., 2012), which contains spontaneous conversations between native Japanese speakers of diverse ages and relationships, in contrast to the scripted PICQ-drama dataset. Following Su et al. (2026), we use GPT-4.1 to automatically extract persona-influenced choice questions (PICQs) and corresponding responses. We retain instances with more than 200 preceding utterances to ensure sufficient context for persona acquisition, resulting in approximately 400 candidates. After manual verification by the first and second authors, 145 high-quality instances are used for evaluation.

We use the same baselines and implementation settings for evaluation as on PICQ-drama, except that human-annotated persona dimensions are unavailable. To better preserve language-specific and cultural characteristics, we adopt Japanese prompts translated by GPT-4.1 for persona dimension prediction and value acquisition, which yield improved performance over English prompts. To account for cross-lingual differences in length measurement, we report average character count (#chars) instead of average word count. Since no human-annotated choice situation summaries are available, we use summaries generated by GPT-5.1 using the instructions for the human annotators used in Su et al. (2026) when evaluating choice consistency as described in § 5.1.

Persona information used for simulation	Simulation Quality (C_{choice})		Persona Efficiency		
	GPT-4.1	Qwen3-32B	#chars ↓	#requests ↓	%acquired
No Persona (context only)	0.600	0.503	n/a	n/a	n/a
Query-agnostic Persona	0.644	0.533	1581	n/a	n/a
Query-focused Persona (one-shot, GPT-4.1)	0.628	0.531	78	4.97	68.1%
(one-shot, Qwen3-32B)	0.628	0.527	78	5.00	68.8%
(progressive, GPT-4.1)	0.621	0.558	34	1.82	75.4%
(progressive, Qwen3-32B)	0.638	0.541	55	2.49	78.7%

Table 4: Results on NUCC for simulating choice-making responses based on different persona information.

Persona information used for simulation	Simulation Quality (C_{choice})		Persona Efficiency		
	GPT-4.1	Qwen3-32B	#words	#requests	%acquired
Query-focused Persona (progressive, GPT-4.1)	0.619	0.580	27	2.98	65.7%
w/o confidence-based stopping	0.616	0.575	40	5.00	72.3%
w/ one random persona value removed	0.600	0.555	20	n/a	n/a
Query-focused Persona (progressive, Qwen3-32B)	0.640	0.597	25	2.35	65.5%
w/o confidence-based stopping	0.630	0.591	46	5.00	73.0%
w/ one random persona value removed	0.617	0.585	17	n/a	n/a

Table 5: Ablation study on confidence-based stopping and perturbation experiments for progressive persona completion in PICQ-drama.

Table 4 presents the results on the NUCC. Overall, we observe trends consistent with those on the PICQ-drama dataset (Table 1). In particular, incorporating persona information consistently improves over the context-only baseline, and the progressive approach achieves comparable or better performance than one-shot acquisition with improved efficiency. These results suggest that our framework generalizes to real-world conversations and is not limited to English or scripted settings. Compared to the results on PICQ-drama, GPT-4.1 maintains similar performance, whereas Qwen3-32B exhibits a noticeable drop, likely reflecting its relatively weaker Japanese capability and the increased difficulty and noise of real-world conversations.

5.2.3 Impact of Confidence-Based Stopping

In this section, we evaluate the utility of confidence-based stopping for progressive persona completion. We focus on its ability to identify when sufficient persona information has been acquired while avoiding unnecessary persona requests.

Does Confidence-Based Stopping Enumerate Sufficient Personas? Table 5 compares our progressive persona completion with a variant without confidence-based stopping, where persona acquisition is forced to continue for a fixed number of iterations (five in our experiments). Across both GPT-4.1 and Qwen3-32B, confidence-based stopping achieves comparable or slightly better choice consistency while using substantially fewer persona requests and persona words.

For GPT-4.1, the confidence-based stopping maintains simulation quality while reducing the average number of persona requests from 5 to 2.98 and the average persona word count from 40 to 27. A similar pattern holds for Qwen3-32B, where stopping slightly improves choice consistency and reduces the number of persona requests by more than half (2.35 vs. 5). These results indicate that continuing acquisition beyond a certain point does not reliably improve simulation quality. Instead, confidence-based stopping effectively identifies the point of diminishing returns, avoiding redundant or uninformative requests.

Does Confidence-Based Stopping Enumerate Only Necessary Personas? To further assess whether the stopping signal reflects meaningful information rather than arbitrary confidence, we conduct a perturbation experiment in which a randomly selected acquired value of the predicted persona dimensions is removed after completion.

Table 5 shows the results, confirming a consistent drop in the simulation quality. This degradation indicates that the acquired persona information is indeed contributing to simulation quality, and that high-confidence states correspond to the presence of performance-critical persona information

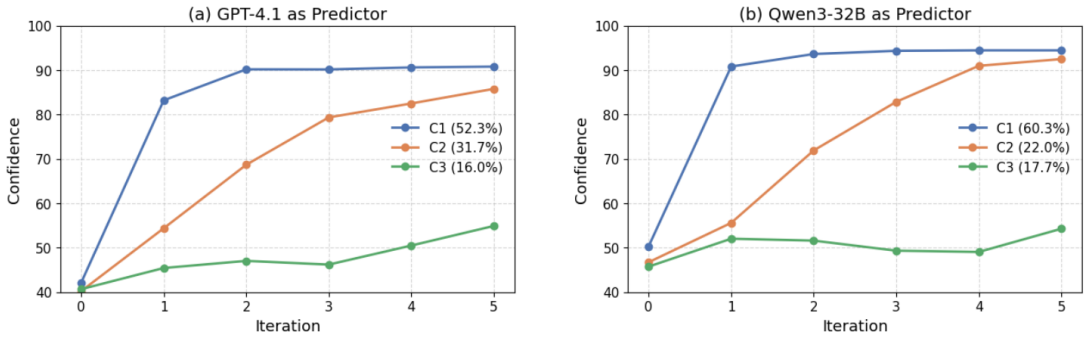


Figure 3: Progressive persona acquisition dynamics on PICQ-drama. Representative confidence trajectories for GPT- and Qwen-based predictors, showing early saturation, gradual improvement, and persistent uncertainty.

rather than incidental or redundant attributes. Overall, these results suggest that confidence provides a practical signal for determining persona sufficiency, enabling efficient and robust persona acquisition across different models and settings.

5.2.4 Persona Acquisition Dynamics

While we confirmed that confidence-based stopping is effective (§ 5.2.3), it does not explain why such a simple signal suffices. To better understand this, we analyze the dynamics of progressive persona acquisition through confidence trajectories. For each query, we record the model's self-reported confidence at each acquisition iteration, forming a trajectory vector. To ensure consistent dimensionality, trajectories are padded to a fixed length (with the final values repeated after termination). We then apply k-means clustering with $k = 3$ to these vectors to identify representative patterns.

Figure 3 shows the progressive persona acquisition dynamics of the resulting clusters for GPT- and Qwen-based predictors, which exhibit similar trends. The dominant pattern is rapid confidence saturation within one or two iterations, indicating that minimal additional persona information is sufficient. A second pattern shows gradual growth in confidence, corresponding to cases that benefit from progressive acquisition. A smaller cluster exhibits persistently low or unstable confidence, suggesting that the required persona dimensions are difficult to infer from available context. These patterns explain why confidence serves as an effective stopping signal: rapid saturation indicates sufficiency, while low or unstable confidence discourages further acquisition. They also suggest that different query types may benefit from adaptive stopping strategies.

6 Conclusions

We study query-focused individual simulation in cold-start settings, where persona information about the target individual is limited or unavailable. Existing user simulation approaches often assume access to rich persona profiles, which can be unrealistic in such scenarios. To address this, we formulate persona acquisition as a query-driven process that selectively elicits only the information necessary for each query. We explore two strategies to request missing but relevant persona dimensions: one-shot prediction and progressive prediction with confidence-based dynamic stopping. Experimental results on the PICQ-drama dataset and additional samples we extracted from NUCC show that our approach maintains or improves simulation quality while reducing persona acquisition cost compared to query-agnostic and one-shot strategies. These results highlight the value of treating persona acquisition as a dynamic, query-dependent process and provide a practical direction for efficient and scalable individual simulation.

Future work includes extending the method to more complex interaction settings (*e.g.*, multi-turn dialogues), improving the robustness of persona inference under noisy or conflicting inputs, and exploring more principled strategies for balancing acquisition cost and simulation quality. Another promising direction is to validate the approach in real-world applications where persona information must be obtained from human users in the loop.

Acknowledgement

This work was supported by Institute for Digital Observatory, the University of Tokyo. We thank Yuhan Zhou for his help in evaluating the quality of simulated responses for the PICQ-drama dataset.

Limitations

We use an LLM-based proxy to provide persona information and self-reported confidence. While this design promotes the reproducibility of experiments, it does not account for variability in human behaviors, making human-in-the-loop evaluation an important next step. Additionally, confidence-based stopping relies on self-reported confidence, which may be imperfectly calibrated and model-dependent. Although effective in our experiments, more robust uncertainty estimation or confidence calibration methods could improve reliability. Our approach relies on predefined persona categories and linguistic patterns, which may limit the expressiveness of persona representations and may overlook more complex or implicit individual traits. However, this design helps maintain structural consistency while allowing some linguistic variability.

References

Yun-Shiuan Chuang, Krirk Nirunwiroj, Zach Studdiford, Agam Goyal, Vincent V. Frigo, Sijia Yang, Dhavan V. Shah, Junjie Hu, and Timothy T. Rogers. 2024. Beyond demographics: Aligning role-playing LLM-based agents using human belief networks. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14010–14026, Miami, Florida, USA. Association for Computational Linguistics.

Zhuoyun Du, Lujie Zheng, Renjun Hu, Yuyang Xu, Xiawei Li, Ying Sun, Wei Chen, Jian Wu, Haolei Cai, and Haochao Ying. 2025. LLMs can simulate standardized patients via agent coevolution. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17278–17306, Vienna, Austria. Association for Computational Linguistics.

Itsuko Fujimura, Shoju Chiba, and Mieko Ohso. 2012. Lexical and grammatical features of spoken and written japanese in contrast: Exploring a lexical profiling approach to comparing spoken and written corpora. In *Proceedings of the VIIth GSCP International Conference. Speech and Corpora*, pages 393–398.

Silin Gao, Beatriz Borges, Soyoung Oh, Deniz Bayazit, Saya Kanno, Hiromi Wakaki, Yuki Mitsufuji, and Antoine Bosselut. 2023. PeaCoK: Persona commonsense knowledge for consistent and engaging narratives. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6569–6591, Toronto, Canada. Association for Computational Linguistics.

Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2025. Scaling synthetic data creation with 1,000,000,000 personas. *Preprint*, arXiv:2406.20094.

Tiancheng Hu and Nigel Collier. 2024. Quantifying the persona effect in LLM simulations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10289–10307, Bangkok, Thailand. Association for Computational Linguistics.

Zhengyu Hu, Jianxun Lian, Zheyuan Xiao, Max Xiong, Yuxuan Lei, Tianfu Wang, Kaize Ding, Ziang Xiao, Nicholas Jing Yuan, and Xing Xie. 2025. Population-aligned persona generation for LLM-based social simulation. *Preprint*, arXiv:2509.10127.

Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. 2024. Better zero-shot reasoning with role-play prompting. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4099–4113, Mexico City, Mexico. Association for Computational Linguistics.

Ang Li, Haozhe Chen, Hongseok Namkoong, and Tianyi Peng. 2025. LLM generated persona is a promise with a catch. In *First Workshop on Social Simulation with LLMs*.

Man Tik Ng, Hui Tung Tse, Jen tse Huang, Jingjing Li, Wenxuan Wang, and Michael R. Lyu. 2024. How well can LLMs echo us? evaluating AI chatbots' role-play ability with ECHO. *Preprint*, arXiv:2404.13957.

OpenAI. 2024. GPT-4 technical report. *Preprint*, arXiv:2303.08774.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA. Association for Computing Machinery.

Joon Sung Park, Carolyn Q. Zou, Aaron Shaw, Benjamin Mako Hill, Carrie Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S. Bernstein. 2024. Generative agent simulations of 1,000 people. *Preprint*, arXiv:2411.10109.

QwenTeam. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Yiting Ran, Xintao Wang, Rui Xu, Xinfeng Yuan, Jiaqing Liang, Yanghua Xiao, and Deqing Yang. 2024. Capturing minds, not just words: Enhancing role-playing language models with personality-indicative data. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14566–14576, Miami, Florida, USA. Association for Computational Linguistics.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on*

Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Yisi Sang, Xiangyang Mou, Mo Yu, Shunyu Yao, Jing Li, and Jeffrey Stanton. 2022. TVShowGuess: Character comprehension in stories as speaker guessing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4267–4287, Seattle, United States. Association for Computational Linguistics.

Shoetsu Sato, Naoki Yoshinaga, Masashi Toyoda, and Masaru Kitsuregawa. 2017. Modeling situations in neural chat bots. In *Proceedings of ACL 2017, Student Research Workshop*, pages 120–127, Vancouver, Canada. Association for Computational Linguistics.

Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-LLM: A trainable agent for role-playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13153–13187, Singapore. Association for Computational Linguistics.

Weiwen Su, Yuhan Zhou, Zihan Wang, Naoki Yoshinaga, and Masashi Toyoda. 2026. Is he extroverted? identifying missing relevant personas for faithful user simulation. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 131–149, Rabat, Morocco. Association for Computational Linguistics.

Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.

Noah Wang, Z.y. Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Jie Fu, and Junran Peng. 2024. RoleLLM: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14743–14777, Bangkok, Thailand. Association for Computational Linguistics.

Tao Wu, Jingyuan Chen, Wang Lin, Mengze Li, Yumeng Zhu, Ang Li, Kun Kuang, and Fei Wu. 2025. Embracing imperfection: Simulating students with diverse cognitive levels using LLM-based agents. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9887–9908, Vienna, Austria. Association for Computational Linguistics.

Miao Xiong, Zhiyuan Hu, Xinyang Lu, YIFEI LI, Jie Fu, Junxian He, and Bryan Hooi. 2024. Can LLMs express their uncertainty? an empirical evaluation of confidence elicitation in LLMs. In *International Conference on Learning Representations*, volume 2024, pages 23650–23678.

Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, and Deqing Yang. 2024. Evaluating character understanding of large language models via character profiling from fictional works. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8015–8036, Miami, Florida, USA. Association for Computational Linguistics.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Yunyao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. 2025. GA-S^3: Comprehensive social network simulation with group agents. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8950–8970, Vienna, Austria. Association for Computational Linguistics.

A Appendix

A.1 Prompts

This section lists the prompts used in the experiments; the prompt for predicting the one-shot missing relevant dimensions is from (Su et al., 2026). Table 6 shows the prompt for progressively predicting one relevant persona dimension with auxiliary probes. Table 7 shows the prompt for acquiring personas for given dimensions. Table 8 shows the prompt for generating a self-reported confidence score for estimating persona sufficiency. Table 9 shows the prompt for query-agnostic persona acquisition. Table 10 shows the RPP style (Kong et al., 2024) prompt for simulating a response of a target individual. Table 11 shows the prompt for automatically evaluating the choice-consistency between the generated responses and the gold responses, which is also referred to by human evaluators.

A.2 Example

Table 12 presents another simulation example from the PICQ-drama dataset, including gold responses, context-only generations, and responses generated with progressively acquired persona information by Qwen3-32B.

```
Given dialogue context, question utterance, known persona of {maker}, and banned persona dimensions,
your task is to:
Identify the single most critical missing persona dimension of {maker} that will influence the choice
of the options provided in the question.
  1. Prioritize the missing persona dimension based on factors that:
    a. Act as necessary conditions, core constraints, or essential enablers (e.g., affordability
for a large purchase, prerequisite required skills).
    b. Represent strong motivations or driving forces behind the choice (e.g., key personal goals,
deeply held beliefs, very strong preferences).
    c. Are critical factors when its value falls within a certain range (e.g., "spice tolerance"
when choosing a Sichuan (spicy) vs. Japanese restaurant).
  2. You should first choose a category and choose one specific linguistic pattern associated
with the category to describe the missing persona dimension. The identified persona dimension must
strictly follow the selected linguistic pattern.
    a. personality:
      (1) whether s/he is ADJ (ADJ is an adjective describing a personality trait); for example,
        - whether s/he is introverted
        - whether s/he is adventurous
    b. beliefs (personal values, moral principles, and views on social norms):
      (1) whether s/he believes it's ADJ to VP (ADJ is an adjective to comment a behavior, VP is a
verb phrase); for example,
        - whether s/he believes it's important to save money
        - whether s/he believes it's wrong to lie to others
      (2) whether s/he believes propN should VP (propN is a target, VP is a verb phrase); for example,
        - whether s/he believes children should have less screen time
        - whether s/he believes the government should invest more in public transport
    c. tastes:
      (1) whether s/he (dis)likes VP (VP is a verb phrase); for example,
        - whether s/he likes traveling
        - whether s/he dislikes waking up early
      (2) whether s/he (dis)likes N (N is a noun); for example,
        - whether s/he likes spicy food
        - whether s/he dislikes crowded places
    d. relationships:
      (1) whether s/he is N of propN (N is a noun representing a human relationship, propN is a name
of a person); for example,
        - whether s/he is a close friend of Alex
        - whether s/he is the sibling of Sarah
      (2) whether s/he is ADJ + P + propN (ADJ represents an adjective or a past participle used
adjectivally, describes the subject's view, attitude, feeling, or judgment regarding the person
propN, P is prepositional); for example,
        - whether s/he is annoyed with Maria
        - whether s/he is loyal to their team
    e. attributes:
      (1) whether his/her ATTR is X (ATTR is a noun describing an attribute of a person (e.g., gender,
occupation, age, height, weight, income), X is a value or range); for example,
        - whether his/her physical stamina is suitable for a long hike
        - whether his/her disposable income is suitable for luxury purchases
    f. goals (short-term or long-term goals):
      (1) whether s/he aims to VP (VP is a verb phrase describing the goal); for example,
        - whether s/he aims to get a promotion
        - whether s/he aims to learn a new language
    g. experience (write 'experience' only if a specific past event directly influences the choice):
      (1) whether s/he has V (V is a past participle phrase); for example,
        - whether s/he has been to that restaurant before
        - whether s/he has had a bad experience with online shopping
  3. For this persona dimension, propose up to two auxiliary probes that:
    - are observable behaviors that indirectly reflect the same underlying trait
    - are likely to be observed given the dialogue context and known persona
    - do NOT introduce any additional persona dimensions
Persona dimensions in [Banned persona dimension] must NOT be selected again.
You should think step by step as follows:
  1. Summarize the choice situation
  2. Review known persona and banned dimensions
  3. Identify the most critical missing persona dimension and propose auxiliary probes
Output format (dictionary):
{"situation_summary": "...", "dimension_category": "...", "dimension_content": "...",
"auxiliary_probes": ["...", "..."]}
```

Table 6: Prompt for identifying the most critical missing persona dimension.

```
You are {seeker} recalling your past conversations with {character}.
You are given:
    - Conversation history between you and {character}, supporting you recall your interactions.
    - One target persona dimension (e.g., "whether he is shy") representing a question about {character}.
    - Two auxiliary probes that describe observable behaviors that may indirectly inform this persona dimension (e.g., "whether he tends to give short or minimal responses").
Your task is to recall, based on the conversation history and your memory of {character}, whether you can form a clear impression about the target persona dimension of {character}.
About auxiliary probes:
    - Use the auxiliary probes only as guidance to help you reflect on relevant evidence from the conversations.
    - Do NOT treat auxiliary probes as separate persona dimensions.
    - Do NOT introduce any persona traits beyond the target dimension.
For the target persona dimension:
    - If your memory or the conversation history provides sufficient evidence to form a clear impression, write the inferred value of that persona dimension in a declarative form (e.g., "he is shy").
    - If you cannot confidently infer the persona based on that, output [Unknown].
You should think step by step as follows:
    1. Keep the persona dimension and corresponding auxiliary probes in mind as guiding requests.
    2. Selectively recall the parts of your memory or the conversation history that are relevant to that dimension.
    3. Infer the value of that persona dimension.
Please strictly follow the exact output format below:
<inferred persona value or [Unknown]>
```

Table 7: Prompt for acquiring persona value for a target dimension.

```
Given dialogue context, question utterance, and known persona of {maker}, your task is to:
Assess whether the current known persona is sufficient to predict {maker}'s choice with high certainty.
You should:
    1. Examine how the known persona supports or contradicts specific options of the choice
    2. Determine whether acquiring additional persona dimensions would likely change the outcome
    3. Assign a confidence score (0–100) reflecting persona sufficiency:
      (90–100): Sufficient; additional persona is unlikely to change the choice
      (70–89): Mostly sufficient, with minor uncertainty
      (50–69): Weakly sufficient; additional persona may change the outcome
      (<50): Insufficient; further persona information is needed
You should think step by step.
Output only the confidence score (0–100) without any explanation.
Example: 45
```

Table 8: Prompt for generating self-reported confidence.

```
You are {seeker} recalling your past conversations with {character}.
You are given:
    - Conversation history between you and {character}, supporting you recall your interactions.
Your task is to recall, based on the conversation history and your memory of {character}, and write down a summary of {character}'s persona
The summary must be within {word limit} words
```

Table 9: Prompt for query-agnostically acquiring persona information.

```
# Stage 1: Role-Setting Prompt
From now on, you are called {maker}.
Here is your persona description: {known_persona}.

# Stage 2: Response Generation Prompt
Now you will be given a dialogue context in which you have participated.
You should generate a response (i.e., make a choice) to the last utterance (question).
Dialogue context:
{context}
Output: <your response>
```

Table 10: RPP style (Kong et al., 2024) two-stage prompting for query-focused response simulation: lines starting with '#' are comments (and are not included in the actual prompt).

```
Given a summary of a choice-making situation (describing what the person needs to choose and why),
and two specific choices:
Your task is to score the consistency between the two choices on a scale from 0 to 1:
    0: The two choices are inconsistent.
    1: The two choices are fully consistent.
Note: Focus only on the meaning and content of the two choices.
Do not consider differences in wording, phrasing, or style unless they affect the actual meaning.
```

Table 11: Prompt for evaluating choice consistency.

Dialogue Context:
Rachel: Mindy.
Mindy: Hey, you.
Rachel: Hey, you.... So, what's up?
Mindy: Um.. we should really be sitting for this.
Rachel: Sure we should... So.
Mindy: ...Okay, I'm just gonna ask you this once, and I want a straight answer.
Rachel: Okay.
Mindy: Will you be my maid of honour?

Gold Response:
Rachel: Of course!

No Persona (Context only):
Rachel: (quietly) ...I need to think.

Query-focused Persona (progressive, Qwen3-32B):
Rachel: (pauses, considering) Of course I will.
Persona (acquired): she believes it's important to support close friends in major life events

Table 12: Another example of simulated choice-making responses in PICQ-drama using the Qwen3-32B simulator.