Not All Persona Dimensions Evolve Alike: Selective Persona Maintenance under Temporal Heterogeneity

Anonymous ACL submission

Abstract

Maintaining coherent user state across histories is a central challenge in personalized language agents. Existing approaches largely assume homogeneous persona dynamics, either uniformly propagating persona state across interactions or inferring persona independently from each interaction. In practice, however, persona dimensions evolve at different temporal scales, creating a mismatch between persona dynamics and state maintenance.

In this paper, we investigate persona maintenance as a temporally heterogeneous inference problem. Specifically, we propose Structured Dynamic Persona Maintenance (SDPM), a uniform per-dimension inference framework that enables controlled analysis of how sequential state maintenance affects dimensions with different temporal persistence. We further introduce Persona-Stance, a longitudinal benchmark of temporally ordered Yelp review histories annotated with five-dimensional persona states.

Experiments reveal a consistent temporal asymmetry in persona modeling. Removing sequential updating collapses performance on slow variables such as price sensitivity, while stateless inference remains more effective for fast variables such as social context. Moreover, uniformly propagating persona state can degrade transient dimensions despite improving stable ones. Selective persona maintenance further improves downstream star-rating prediction. Together, these findings suggest that persona maintenance should be conditioned on temporal persistence at the dimension level.

1 Introduction

Personalized language agents increasingly rely on longitudinal interaction histories to infer and adapt to their users (Shanahan et al., 2023; Wang et al., 2024; Salemi et al., 2024). Existing approaches follow one of two paradigms: memory-augmented methods maintain persona as persistent state carried across interactions (Packer et al., 2024; Jiang et al., 2025; Feng et al., 2025), while stateless approaches infer persona independently from each interaction (Li et al., 2016; Zhang et al., 2018; Salemi et al., 2024). Both paradigms apply a single temporal maintenance strategy uniformly across all persona dimensions.

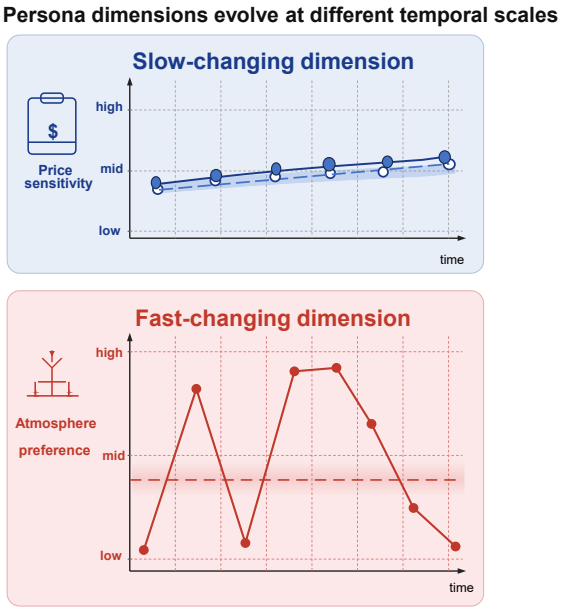

Figure 1: Temporal heterogeneity of persona dimensions motivates asymmetric maintenance.

However, persona dimensions are structurally heterogeneous in their temporal persistence. Some reflect stable, accumulative preferences that leave coherent traces across interactions; others capture transient contextual states that vary with situational factors independent of history. A uniform policy therefore systematically mismatches one class or the other: state propagation imports stale signal into fast-evolving dimensions, while stateless inference discards accumulated evidence from slow-evolving ones. We term this the *slow-*

fast asymmetry of persona dynamics.

To address this, we propose **Structured Dynamic Persona Maintenance (SDPM)**, a sequential inference framework that maintains persona as a structured discrete latent state updated incrementally. SDPM decomposes inference across dimensions via dimension-specialized adapters, allowing each dimension to integrate historical state and current interaction evidence through its own learned dynamics, before recomposing predictions into a unified structured representation.

Note that, since each dimension is parameterized by an independent adapter, selective persona maintenance reduces to enabling adapters for slow dimensions and disabling them for fast ones, requiring no architectural modification.

To evaluate persona maintenance under heterogeneous temporal persistence, we introduce **Persona-Stance**, a temporally ordered benchmark constructed from Yelp review histories with causal persona annotations across five dimensions. The benchmark enables controlled evaluation of how different temporal maintenance strategies affect both persona inference and downstream behavioral prediction.

Our contributions are as follows:

- We identify the ***slow-fast asymmetry*** as a structural property of longitudinal persona dynamics, showing that uniform maintenance strategies induce systematic errors across dimensions with different temporal persistence.
- We propose **SDPM**, a dimension-conditional inference framework that applies asymmetric temporal updates for slow- and fast-evolving persona attributes.
- We introduce **Persona-Stance**, a causally ordered benchmark for evaluating persona maintenance under heterogeneous temporal dynamics and its impact on downstream behavioral prediction.

2 Related Work

Stateless Persona Modeling Persona modeling has been a central thread in dialogue and personalized generation research. Early work grounded response consistency in explicit persona descriptions (Li et al., 2016; Zhang et al., 2018), with subsequent efforts extending this to implicit signals inferred from user text (Wu et al., 2021). As LLMs emerged, the community shifted toward prompting and role-playing (Shanahan et al., 2023; Wang et al., 2024; Malvicini et al., 2026), memory-augmented systems (Packer et al., 2024; Jiang et al., 2025), and activation-space steering (Feng et al., 2025; Liu et al., 2026) to sustain persona across interactions (Yang et al., 2026). Across these paradigms, however, persona is consistently treated as a monolithic representation, and the same temporal strategy tends to apply uniformly across all of its dimensions.

Dynamic Persona Maintenance Dynamic user modeling has evolved from matrix factorization (Koren et al., 2009) and RNN-based session models (Hidasi et al., 2016) to Transformer architectures capturing long-range behavioral dependencies (Kang and McAuley, 2018). These methods track user state effectively but operate over structured behavioral logs rather than natural language, limiting transfer to LLM-driven settings. More recent work connects LLM personalization with longer interaction horizons (Park et al., 2023), yet cross-temporal consistency and inter-dimension structure have received limited attention (Song et al., 2021; Bae et al., 2022). Whether persona dimensions that evolve at different timescales are better served by stateful maintenance than by stateless re-inference, and how accumulated state shapes downstream behavioral prediction, remains underexplored and motivates our framework.

3 The PERSONA-STANCE Benchmark

Existing persona benchmarks primarily evaluate static profile consistency or persona-conditioned generation (Zhang et al., 2018; Salemi et al., 2024). In contrast, PERSONA-STANCE focuses on longitudinal persona maintenance: whether a model can track structured persona states as user interactions unfold over time, and whether these inferred states improve downstream behavioral prediction. The benchmark targets two goals: **(G1)** whether models can infer coherent multi-dimensional persona states under strict causal ordering, and whether structured sequential maintenance improves persona inference over prompt-only baselines; and **(G2)** whether inferred persona representations improve downstream behavioral prediction.

3.1 Data Construction

We construct PERSONA-STANCE from the Yelp Open Dataset (Yelp, 2023). For each user u, we

Table 1: Task definitions in PERSONA-STANCE.

Task	Input	Output
Persona inference	Review history $r_u^{1:t}$; prior state $\mathbf{p}_{t-1}$	$\mathbf{p}_t \in \mathcal{Y}^5$
Star-rating prediction	Recent reviews $r_u^{k+1:T-1}$; persona $\mathbf{p}$	$y_T \in \{1, \ldots, 5\}$

Table 2: Data split for persona inference.

Split	Users	Samples
Training	200	6,000
Test	134	4,020

build a temporally ordered review sequence $\mathcal{R}_u = \{r_u^1, r_u^2, \ldots, r_u^T\}$, where each review contains text, timestamp, and a star rating.

While the user count is moderate, each user contributes a temporally ordered trajectory of at least 15 reviews, yielding 6,000/4,020 train/test samples across 8,180 inference slots under a user-disjoint split, which is substantially larger and more structured than most existing persona benchmarks.

Persona dimensions. Each user state is represented as a five-dimensional persona vector over *price_value*, *atmosphere*, *convenience*, *social_context*, and *beverages*, each discretized into $\mathcal{Y} = \{\texttt{low}, \texttt{medium}, \texttt{high}\}$. These labels capture the degree of user attention toward each attribute rather than direct sentiment polarity or business quality. The dimensions are derived from Yelp business attributes and review-level lexical cues: structured metadata such as *Ambience*, *NoiseLevel*, and *RestaurantsDelivery*, together with keywords such as "price", "vibe", "parking", and "drinks", are mapped into higher-level persona dimensions.

Downstream data construction. For the sentiment prediction task, each sample is constructed from a user's review sequence at a fixed target step T, with the target review r_u^T withheld. The visible input consists of the review window $r_u^{k+1:T-1}$ and the business attributes of the target venue; the label is the star rating $y_T \in \{1, \ldots, 5\}$, binarized into high ($\geq$4) and low ($\leq$3) for evaluation. Strict causal ordering is preserved throughout: no review at or after step T is accessible during inference. Of the 218 users with evaluable persona states at t=15, 121 have gold star labels in the test split and form the aligned evaluation subset used for downstream comparison.

Temporal persona construction. Persona states are maintained under strict causal ordering. The initial state $\mathbf{p}_k$ is bootstrapped from the first k reviews, and each subsequent state is updated incrementally as each new review arrives:

$$\mathbf{p}_t = f(\mathbf{p}_{t-1}, r_u^t), \quad t \geq k+1. \tag{1}$$

No future observations are accessible during inference or evaluation.

Annotation. Persona labels are generated by Qwen2.5-32B-Instruct under a fixed JSON schema, following the same causal ordering as inference: $\mathbf{p}_k$ is produced from the first k reviews, and each subsequent state $\mathbf{p}_t$ is updated from $(\mathbf{p}_{t-1}, r_u^t)$. Rule-based weak supervision serves as a quality-control reference and is not used as evaluation gold. Full annotation procedures are detailed in Appendix B).

3.2 Tasks

As shown in Table 1, we define two evaluation tasks over PERSONA-STANCE, all enforcing strict causal ordering: the target review at step t is never accessible when predicting $\mathbf{p}_t$.

Multi-dimensional persona inference. This is the primary task. The model infers a five-dimensional persona state $\mathbf{p}_t \in \mathcal{Y}^5$ from chronologically ordered reviews in two phases. In the *initialization* phase, the model reads the first k reviews and produces $\mathbf{p}_k$. In the *sequential update* phase, it receives $\mathbf{p}_{t-1}$ and one new review r_u^t and outputs an updated $\mathbf{p}_t$.

Persona-based sentiment prediction. This task evaluates whether structured persona states improve behavioral prediction. Given a window of recent reviews and business attributes for the target venue, the model predicts the star rating $y_T \in \{1, \ldots, 5\}$ without access to r_u^T.

4 Methodology

4.1 Overview

We propose Structured Dynamic Persona Maintenance (SDPM), a sequential inference framework that addresses the slowfast asymmetry identified in ğ1. The core idea is to represent persona not as a monolithic embedding refreshed at each step, but as a structured discrete latent state maintained per dimension, so that each attribute can accumulate historical evidence at its own natural

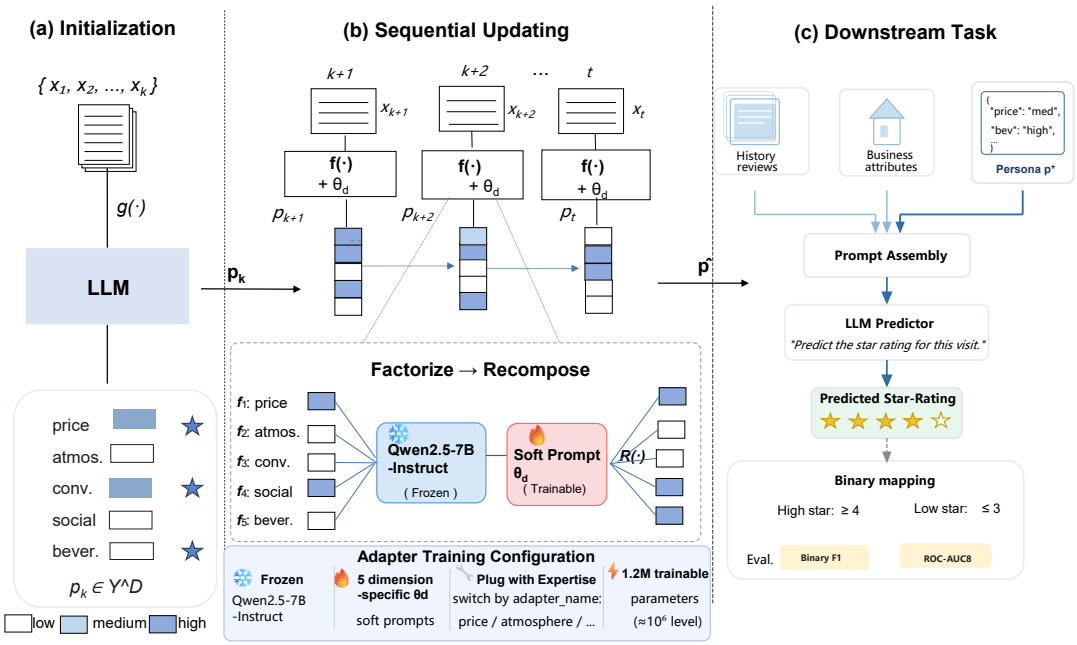

Figure 2: Overview of the proposed framework.

rate. As illustrated in Figure 2, SDPM consists of three components: Per-Dimension Inference, which updates each persona attribute through a dedicated dimension-specific adapter; Sequential State Maintenance, which propagates the structured state across interactions and resolves cross-dimensional conflicts; and Downstream Sentiment Prediction, which evaluates whether the maintained state improves forward-looking star-rating prediction.

4.2 Per-Dimension Persona Inference

SDPM represents each user's persona as a structured discrete state over $D{=}5$ dimensions:

$$\mathbf{p}_t \in \mathcal{Y}^D, \quad \mathcal{Y} = \{\texttt{low}, \texttt{medium}, \texttt{high}\}. \tag{2}$$

After $k{=}10$ chronologically ordered reviews, initialization produces

$$\mathbf{p}_k = g(x_1, \ldots, x_k), \tag{3}$$

where g is a structured LLM mapping under a fixed JSON schema (Li et al., 2016; Zhang et al., 2018). For $t > k$, each dimension d is updated via a dedicated adapter θ_d:

$$\hat{y}_{t,d} = f_d(\hat{\mathbf{p}}_{t-1},\, x_t;\, \theta_d), \quad d \in \mathcal{D}, \tag{4}$$

where θ_d is a dimension-specific soft prompt (Houlsby et al., 2019; Lester et al., 2021). Enabling or disabling θ_d per dimension directly implements selective maintenance, as motivated by the slow-fast asymmetry.

4.3 Sequential Persona Maintenance

For $t > k$, SDPM integrates per-dimension adapter outputs into a jointly consistent structured state:

$$\hat{\mathbf{p}}_t = \mathcal{R}\Big(f_1(\hat{\mathbf{p}}_{t-1}, x_t; \theta_1), \ldots, f_D(\hat{\mathbf{p}}_{t-1}, x_t; \theta_D)\Big). \tag{5}$$

At inference, $\hat{\mathbf{p}}_{t-1}$ is the model's own prediction from step $t{-}1$; gold personas are used only for evaluation. When all five dimensions are successfully parsed, the prior state is injected into the update prompt; otherwise an empty state $\{\}$ is substituted to prevent cascading parse errors. Recomposition $\mathcal{R}$ is an additional LLM pass that resolves inter-dimension conflicts across the five adapter outputs; disabling $\mathcal{R}$ reduces it to a union of independent parses. latex

4.4 Persona-based Sentiment Prediction

To assess whether maintained persona states benefit downstream prediction, we formulate a sentiment prediction task at step T. The predictor receives the visible review history and business attributes for the target venue:

$$\hat{y}_T = \mathrm{LLM}(h_T,\, b,\, z), \tag{6}$$

$$h_T = \mathrm{concat}(x_{k+1}, \ldots, x_{T-1}), \tag{7}$$

where the target review x_T is withheld, and the persona slot z takes one of four values:

$$z \in \{\, \varnothing,\ \hat{\mathbf{p}}_k,\ \hat{\mathbf{p}}_{T-1},\ \tilde{\mathbf{p}}_{T-1} \,\}. \tag{8}$$

Table 3: Overall aspect-level accuracy on Persona-Stance.

Method	Avg Aspect Acc. (%)↑
SDPM (Qwen2.5-7B + PEFT)	**49.00**
Qwen2.5-7B (prompt-only)	47.14
GPT-4o-mini (prompt-only)	45.17
LLaMA-3-8B (prompt-only)	34.75

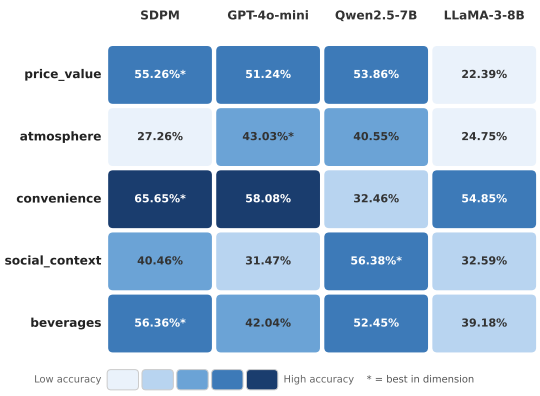

Figure 3: Per-dimension Accuracy Heatmap

Specifically, $\varnothing$ denotes no persona; $\hat{\mathbf{p}}_k$ is the SDPM state after initialization; $\hat{\mathbf{p}}_{T-1}$ is the fully propagated SDPM state; and $\tilde{\mathbf{p}}_{T-1}$ is a persona inferred from the same review window by a plain LLM without adapters or cross-step maintenance. All conditions share a fixed prompt template; only z varies.

5 Experiments

5.1 Setup

We evaluate on PERSONA-STANCE with a user-disjoint train/test split (6,000/4,020 samples). Reviews at t=1–10 initialize the persona state; $t \geq 11$ are reserved for sequential updating and evaluation. The upstream inference pipeline runs Qwen2.5-7B-Instruct with dimension-specialized PEFT modules. The downstream star-rating predictor uses deterministic decoding conditioned on t=11–14 review history and business attributes, excluding the target review at t=15. The aligned evaluation subset (n=121) retains users for which all four persona conditions produce parseable output. All reported results are averaged over 5 independent runs with different random seeds. Implementation details are in Appendix A.

5.2 Slow-Fast Asymmetry in Persona Inference

We compare SDPM against three stateless inference baselines: Qwen2.5-7B-Instruct, GPT-4o-mini, and LLaMA-3-8B-Instruct, each of which infers persona independently at each step without sequential state maintenance or dimension-specialized modules. Since SDPM shares the same backbone as the Qwen2.5-7B baseline, this comparison directly isolates the contribution of sequential state maintenance and dimension specialization from raw prompting capacity. Per-dimension accuracy results are shown in Figure 3.

SDPM consistently leads on slow dimensions, while stateless inference outperforms on fast dimensions. On slow dimensions, where preference signals surface as recurring, identifiable lexical cues, SDPM leads: `price_value` (55.26%), `convenience` (65.65%), and `beverages` (56.36%). On fast dimensions, however, stateless inference on the same backbone outperforms structured maintenance: `atmosphere` (40.55% vs. 27.26%) and `social_context` (56.38% vs. 40.46%), where signals are contextual, affective, and situationally variable. These results directly instantiate the slow-fast asymmetry: sequential state maintenance accumulates evidence where signals are consistent across interactions, and imports stale priors where they are not.

The aggregate gain masks the underlying asymmetry. SDPM achieves the highest average aspect accuracy (49.00%), ahead of Qwen2.5-7B (47.14%), GPT-4o-mini (45.17%), and LLaMA-3-8B (34.75%; Table 3). Strong performance on slow dimensions is partially offset by degraded performance on fast ones, and should not be interpreted as evidence for uniform benefit of sequential state maintenance.

The advantage of stateless inference on fast dimensions is robust across architectures. On `atmosphere`, GPT-4o-mini leads (40.55%), while on `social_context`, Qwen2.5-7B prompt-only leads (56.38%). The fact that two different stateless baselines outperform SDPM on different fast dimensions suggests that the performance gap reflects a structural property of those dimensions rather than a limitation of any particular model or inference approach.

Table 4: Ablation: aspect-level EM and persona-level EM (English slice, n=660 samples, G=132 groups). These figures are not directly comparable to Table 3 due to different protocols.

Variant	Removed	Aspect EM	Macro F1	Micro F1
Full SDPM	—	73.79%	15.15%	73.79%
w/o Recomposition	$\mathcal{R}$	73.18%	42.98%	73.18%
w/o Sequential updating	$f(\mathbf{p}_{t-1}, \cdot)$	71.97%	15.61%	72.11%
w/o `price_value` adapter	θ_{price}	70.61%	14.74%	70.61%
w/o `atmosphere` adapter	θ_{atm}	73.94%	14.82%	73.94%
w/o `convenience` adapter	θ_{conv}	70.76%	14.67%	70.76%
w/o `social_context` adapter	θ_{social}	76.52%	15.44%	76.52%
w/o `beverages` adapter	θ_{bev}	69.09%	15.15%	69.09%

5.3 Ablation: Causal Evidence for the Slow-Fast Asymmetry

We ablate three components of SDPM, consisting sequential state maintenance $f(\mathbf{p}_{t-1}, \cdot)$, structured recomposition $\mathcal{R}$, and each dimension-specific adapter θ_d. Results are reported in Table 4.

Sequential state maintenance selectively benefits slow dimensions. Removing sequential state maintenance reduces overall aspect EM from 73.79% to 71.97%, but this aggregate conceals a dimension-level collapse. `price_value` drops from 74.24% to 29.55%, a reduction of 44.7 pp, while `atmosphere` and `social_context` shift by less than 3 pp. It is evident that slow dimensions depend critically on cross-step state propagation: price preferences are expressed through recurring, comparable signals across reviews, and sequential state maintenance converts this distributed evidence into a stable estimate that single-step stateless inference cannot reliably recover. Fast dimensions, whose signals vary with situational context, show no comparable dependence on the propagated state.

Structured maintenance actively degrades fast dimensions. The adapter ablation reveals the symmetric cost. Removing adapters for slow dimensions consistently degrades performance: `beverages` (-4.7%), `price_value` (-3.2%), and `convenience` (-3.0%). In contrast, removing the `social_context` adapter raises overall accuracy by 2.73%, with the dimension's own accuracy improving from 57.58% to 72.73%; removing the `atmosphere` adapter yields a further gain of 0.15%, each evaluated independently against full SDPM. A dimension-specialized adapter trained on a fast dimension overfits to spurious co-occurrences, since the underlying signal is inconsistent across reviews, and the unspecialized backbone avoids this bias. These results directly prescribe a selective maintenance policy: structured maintenance benefits slow dimensions and actively degrades fast ones, and the optimal configuration applies sequential state maintenance only where the underlying signal warrants it. We examine the mechanism underlying this identifiability gap in §6.4.

Structured recomposition enforces cross-dimension coherence. Removing $\mathcal{R}$ leaves per-aspect EM nearly unchanged (73.18%) but reduces persona-level EM by 0.75 pp, the expected behavior of a conflict-resolution operator that harmonizes independently produced dimension predictions into a jointly consistent structured state. Sequential state maintenance and $\mathcal{R}$ therefore target distinct failure modes and are complementary rather than redundant.

We evaluate four conditions on the aligned subset: Qwen2.5-7B without persona, Qwen2.5-7B augmented with SDPM-$\mathbf{p}_{10}$, with SDPM-$\mathbf{p}_{14}$, and with single-pass $\tilde{\mathbf{p}}_{14}$. All conditions share the same downstream predictor and prompt template, differing only in the persona slot z; performance differences therefore reflect persona representation quality rather than predictor capacity. Results are reported in Table 5.

It is evident that persona state consistently improves over no-persona across all conditions, establishing user preference representation as a useful conditioning signal for downstream sentiment prediction. Moreover, SDPM-$\mathbf{p}_{14}$ outperforms SDPM-$\mathbf{p}_{10}$ (71.91% vs. 71.19% F1), indicating that continued sequential state maintenance beyond initialization accumulates signal not yet captured by the early persona estimate. Given the evaluation cohort, fine-grained differences between conditions should be treated as indicative rather than conclusive. The ablation in §5.3 provides the

Table 5: Results of persona-based sentiment prediction. All conditions share the same Qwen2.5-7B downstream predictor; only the persona slot z varies.

Condition	F1 (%)	ROC-AUC (%)	Acc (%)	Prec (%)	Rec (%)
Qwen2.5-7B (no persona)	69.36	47.76	56.20	63.16	76.92
Qwen2.5-7B + SDPM-$\mathbf{p}_{10}$	71.19	48.52	57.85	63.64	80.77
Qwen2.5-7B + SDPM-$\mathbf{p}_{14}$	**71.91**	49.17	**58.68**	64.00	**82.05**
Qwen2.5-7B + $\tilde{\mathbf{p}}_{14}$ (single-pass)	71.59	**49.69**	**58.68**	**64.29**	80.77

$\tilde{\mathbf{p}}_{14}$ is a structured persona state inferred from the same review window as SDPM-$\mathbf{p}_{14}$ in a single forward pass, without dimension-specific adapters or cross-step state maintenance.

stronger and more direct evidence: removing fast-dimension adapters individually improves persona inference by 2.73% (`social_context`) and 0.15% (`atmosphere`), confirming that a selective maintenance policy yields higher-quality persona representations than uniform sequential maintenance.

6 Analysis and Discussion

6.1 Why Sequential State Maintenance Benefits Slow Dimensions

Sequential state maintenance encodes prior inferences as an explicit structured state $\mathbf{p}_{t-1}$, allowing incremental revision as new evidence arrives rather than re-deriving predictions from the full concatenated history at each step. This matters most for slow dimensions, whose preference signals are sparse within any single review but consistent across many: sequential propagation converts distributed evidence into a stable estimate that single-pass stateless inference underweights. The ablation confirms this directly: `price_value` drops from 74.24% to 29.55% upon removing sequential state maintenance, while `atmosphere` changes by less than 3%, where the relevant cues do not accumulate in the same way.

Structured recomposition $\mathcal{R}$ addresses a distinct problem: when five dimension predictions are produced by independent adapters, the resulting persona state may be internally inconsistent. $\mathcal{R}$ resolves these conflicts at the persona level, which is why its effect appears in persona-level exact match rather than per-dimension accuracy. Sequential state maintenance and $\mathcal{R}$ are therefore complementary operators targeting distinct failure modes.

6.2 The Slow-Fast Asymmetry Generalizes Across Domains

The slow-fast asymmetry observed on Yelp may reflect domain-specific properties of restaurant reviews rather than a general structural property of persona inference. To test this, we conduct a controlled experiment on the PersonaHub `knowledge` split (Ge et al., 2024), a completely different dataset with different personas, synthesized text, and no temporal structure. Inference runs without SDPM, adapters, or any temporal chain, isolating whether dimension-level heterogeneity persists independently of the framework. We sample n=200 instances and define five dimensions in parallel with the Yelp schema.

It is evident that the same asymmetric pattern reproduces. Inference accuracy ranges from 54.0% (`topic_specificity`) to 13.5% (`learning_curiosity`), with a macro average of 35.2% (Table 6). Dimensions whose signals surface through explicit, recurring textual cues perform substantially better than dimensions whose labels reflect latent motivation or identity states. Specifically, `learning_curiosity` yields the lowest accuracy because the gold label reflects the persona's stated motivation, while the synthesized text adopts an expert-sharing register that provides almost no signal about underlying curiosity level; the model consequently defaults to predicting `high` in 165 of 200 cases. This mirrors the `atmosphere` and `social_context` pattern on Yelp, where fast dimensions share the same failure mode: diffuse textual evidence and skewed label distributions. These results confirm that the slow-fast asymmetry is a general property of persona inference from natural language, not an artifact of the restaurant review domain or the SDPM framework.

6.3 The Role of Persona as a Conditioning Variable

The downstream results reveal a consistent pattern: any persona condition improves over no-persona, but the gap between the static $\mathbf{p}_{10}$ and the dynamically updated $\mathbf{p}_{14}$ is small. This pattern has a structural explanation distinct from questions about method quality.

Table 6: PersonaHub `knowledge` dimension identifiability: lexical baseline vs. GPT-4o-mini on synthesized text (n=200 per dimension; accuracies in %). Not comparable to Yelp main-task numbers.

Dimension	Inferred semantics	Lex.	LLM	Δ	Signal
`topic_specificity`	Narrow vs. broad topic (high/low)	43.5	**54.0**	+10.5	●
`expertise_domain`	Expertise domain (6 categories)	37.0	**42.5**	+5.5	●
`professional_level`†	Expert / practitioner / enthusiast	**86.5**	35.5	−51.0	●
`primary_stance`	Educator / researcher / enthusiast / practitioner	**60.5**	30.5	−30.0	●
`learning_curiosity`	Learning-oriented vs. specialized output	**37.0**	13.5	−23.5	×
Macro avg.	—	52.9	35.2	—	—

● = moderate text–label alignment; × = weak alignment. †Lexical baseline exploits "As a…" openings; LLM over-predicts `expert` on `professional_level` (124 vs. 6 gold). JSON parse rate = 100%.

Persona functions as a slow-moving summary variable. Most stable preferences, including price sensitivity, beverage choices, and accessibility concerns, are established early in a user's review history and change slowly thereafter. $\mathbf{p}_{10}$, initialized from the first ten reviews, already captures the bulk of this stable structure; updates from t=11 to 14 refine the estimate at the margins. Furthermore, the downstream prompt already includes the t=11–14 review texts directly, so the information that drives those updates is partially present in the input regardless of persona condition. These two factors together bound the observable benefit of dynamic maintenance in this evaluation setup.

The short-horizon prediction experiment (Appendix D) makes this concrete: predicting $\mathbf{p}_{t+1}$ from $\mathbf{p}_t$ alone barely exceeds a persistence baseline (49.55% vs. 49.85% macro accuracy), confirming that persona state transitions carry no autoregressive structure exploitable without new observations. Sequential state maintenance is therefore valuable for consolidating evidence from incoming text, not for extrapolating state dynamics in the absence of new signal.

6.4 Identifiability as the Mechanism Behind the Asymmetry

The asymmetry in adapter contributions is consistent in direction across all five dimensions: averaged across slow dimensions, removing the adapter degrades performance by 3.63%; averaged across fast dimensions, removing the adapter improves performance by 1.44%. This systematic reversal in sign confirms that dimension type determines whether structured maintenance is beneficial, rather than any incidental property of individual dimensions.

The slow-fast distinction maps onto a more fundamental property of each dimension: whether its preference signals surface reliably in text, which we term *dimension identifiability*. Slow dimensions (`price_value`, `convenience`, `beverages`) have high identifiability; their signals appear as recurring, lexically distinct cues that accumulate stably across reviews, making them amenable to structured state maintenance and dimension-specialized adapters. Fast dimensions (`atmosphere`, `social_context`) have low identifiability; their signals are expressed implicitly and inconsistently, and in the ablation the specialized adapter for these dimensions actively hurts accuracy by overfitting to spurious co-occurrences.

Identifiability also clarifies what dimension-specialized adapters can and cannot do. Adapters are effective when the task reduces to learning a reliable mapping from identifiable surface patterns to discrete labels, effectively a structured extraction problem. They are not effective when the label reflects latent attributes that do not surface predictably in text, since there is no consistent signal to amplify. Some weakly identifiable dimensions ma y require richer supervision or domain-specific annotation, independent of the sequential state maintenance framework.

7 Conclusion

We identify the slow-fast asymmetry in longitudinal persona dynamics and propose SDPM, a dimension-conditional framework for sequential persona maintenance. Experiments on PersonaStance show that sequential updating benefits slow-evolving dimensions, while fast-evolving dimensions are better handled through local inference. These findings suggest persona maintenance should be conditioned on temporal persistence rather than applied uniformly across dimensions.

Limitations

This work is limited by computational constraints, experiments are conducted on a single GPU, restricting the scale of both training and evaluation. The downstream sentiment prediction experiments are evaluated on a single dataset (PERSONA-STANCE) with a small aligned subset. Future work should extend this evaluation to PersonaHub and other longitudinal corpora to validate the generalizability of selective maintenance across domains. Additionally, all experiments use Qwen2.5-7B-Instruct as the backbone; extending SDPM to larger models and different task settings, such as dialogue and recommendation, remains an open direction.

References

Sanghwan Bae, Donghyun Kwak, Soyoung Kang, Min Young Lee, Sungdong Kim, Yuin Jeong, Hyeri Kim, Sang-Woo Lee, Woomyoung Park, and Nako Sung. 2022. Keep me updated! memory management in long-term conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3769–3787, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

X. Feng, L. Zhao, W. Zhong, Y. Huang, Y. Gu, L. Kong, and B. Qin. 2025. Persona: Dynamic and compositional inference-time personality control via activation vector algebra. In *International Conference on Learning Representations (ICLR)*.

Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2024. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*.

Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. *Preprint*, arXiv:1511.06939.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. *Preprint*, arXiv:1902.00751.

Bowen Jiang, Yuan Yuan, Maohao Shen, Zhuoqun Hao, Zhangchen Xu, Zichen Chen, Ziyi Liu, Anvesh Rao Vijjini, Jiashu He, Hanchao Yu, Radha Poovendran, Gregory Wornell, Lyle Ungar, Dan Roth, Sihao Chen, and Camillo Jose Taylor. 2025. Personamem-v2: Towards personalized intelligence via learning implicit user personas and agentic memory. *Preprint*, arXiv:2512.06688.

Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 197–206. IEEE.

Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37.

Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *Preprint*, arXiv:2104.08691.

Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and William B Dolan. 2016. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003.

Yuxin Liu, Mingye Zhu, Siyuan Liu, Bo Hu, and Lei Zhang. 2026. Enhancing persona following at decoding time via dynamic importance estimation for role-playing agents. *arXiv preprint arXiv:2603.01438*.

Stephanie Anneris Malvicini, Ewelina Gajewska, Arda Derbent, Katarzyna Budzynska, Jarosław A. Chudziak, and Maria Vanina Martinez. 2026. A natural language agentic approach to study affective polarization. *arXiv preprint arXiv:2603.02711*.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. 2024. Memgpt: Towards llms as operating systems. *Preprint*, arXiv:2310.08560.

Joon Sung Park, Joseph C. O'Brien, Carrie J. 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)*, UIST '23, New York, NY, USA. Association for Computing Machinery.

Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. LaMP: When large language models meet personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392.

Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role-play with large language models. *Preprint*, arXiv:2305.16367.

Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, and Ting Liu. 2021. BoB: BERT over BERT for training persona-based dialogue models from limited personalized data. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–177. Association for Computational Linguistics.

Noah Wang, Zy Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, and 1 others. 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.

Chen Henry Wu, Yinhe Zheng, Xiaoxi Mao, and Minlie Huang. 2021. Transferable persona-grounded dialogues via grounded minimal edits. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2368–2382.

Jing Yang, Moritz Hechtbauer, Elisabeth Khalilov, Evelyn Luise Brinkmann, Vera Schmitt, and Nils Feldhus. 2026. Persona prompting as a lens on LLM social reasoning. *arXiv preprint arXiv:2601.20757*.

Yelp. 2023. Yelp open dataset. https://www.yelp.com/dataset. Accessed: 2026-05-26.

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.

Acknowledgments

A Implementation Details

Dimension-specialized PEFT. Five independent soft-prompt modules are trained on the shared Qwen2.5-7B-Instruct backbone, one per aspect dimension, each with the same number of virtual tokens. At inference time, the active module is switched before each aspect prediction. The *−Prompt Tuning* ablation omits all modules and relies solely on a text-only dimension prompt, serving as the no-specialization baseline.

Temporal state injection. Update-phase prompts contain a placeholder that is filled at runtime by serializing $\mathbf{p}_{t-1}$ into JSON. The previous state is injected only when all five dimensions at $t-1$ were successfully parsed; if any dimension is missing, an empty object is injected instead to prevent partial-state contamination. The *w/o Sequential Updating* ablation suppresses this injection entirely.

Structured recomposition. After all five aspect predictions for a group are produced independently, $\mathcal{R}$ passes them jointly to the backbone, which resolves inter-dimension conflicts and fills missing slots under the fixed schema, yielding $\hat{\mathbf{p}}_t$. When $\mathcal{R}$ is disabled, the persona-level prediction is formed by taking the union of the five independently parsed labels without conflict resolution.

Metrics. For persona inference, we report average aspect-level accuracy and per-dimension accuracy across the five persona dimensions. For ablation analysis, we additionally report aspect-level exact match (EM), which requires the predicted label to exactly match the gold label for each dimension independently, and persona-level EM, which requires all five dimensions to match simultaneously for a given user at a given timestep. Macro F1 is reported where noted; due to class imbalance, it should be interpreted alongside EM rather than in isolation. For downstream sentiment prediction, we report binary F1 and ROC-AUC for the classification task (high star $\geq$4 vs. low star $\leq$3), together with accuracy, precision, and recall.

B Persona Label Construction

Persona gold labels follow a two-stage protocol (§3.1). All main and ablation results are scored against **Stage 2** outputs unless noted otherwise.

Stage 1: Rule-based weak supervision (auxiliary). We derive auxiliary ordinal signals from (i) structured business attributes associated with each review and (ii) lexical mention indicators in review text, following the weak-supervision paradigm of combining multiple noisy labeling functions. Signals are discretized into $\{\texttt{low}, \texttt{medium}, \texttt{high}\}$ per dimension. Stage 1 labels support quality control and descriptive analysis only; they are *not* used as the primary evaluation reference in §5.

Stage 1 discretization. For boolean business attributes, we use $\texttt{true_frequency} = \texttt{true_count}/\texttt{total_count}$. For *price_value*, ordered price tiers $\{1, 2, 3, 4\}$ map to `low`, `medium`, and `high` ($1 \rightarrow$`low`, $2 \rightarrow$`medium`, $3, 4 \rightarrow$`high`); out-of-range values use $<$ $2 \rightarrow$`low`, else `high`. For frequency-based dimensions, relative concern $\texttt{relative_concern} = \texttt{numeric_freq}/\texttt{baseline}$ yields `high` if > 1.1, `low` if < 0.9, and `medium` otherwise; missing baselines fall back to fixed thresholds or binary splits.

Stage 2: LLM structured pseudo-labeling (primary gold). Under a fixed JSON schema, **Qwen2.5-32B-Instruct** generates structured persona states from the dynamic-persona prompts: $\mathbf{p}_{10}$ is inferred from reviews $r_u^1, \ldots, r_u^{10}$; for $t \geq 11$,

$\mathbf{p}_t$ is produced by conditioning on $\mathbf{p}_{t-1}$ and the review at t under the same prompt template. Responses are parsed into the closed label space; unparseable outputs are excluded from metric denominators where applicable. Reported accuracy and F1 therefore measure agreement with this pseudo-labeling protocol, not human adjudication.

Label semantics. Each dimension encodes *preference salience*—how strongly a user attends to an attribute in their writing—rather than sentiment toward a business. Both price-seeking and price-avoidant language can indicate high salience on *price_value*. Prompt templates and the label-filling script (`05_fill_dynamic_persona_labels.py`) are included in the supplementary materials.

C Evaluation Protocols

Main experiment. Average aspect accuracy is computed as the fraction of aspect slots with a valid gold label that are predicted correctly:

$$\text{Avg Aspect Acc} = \frac{\text{\# correct aspect slots}}{\text{\# aspect slots with gold}}.$$

SDPM is evaluated over the full sequential trajectory (8,180 slots total); prompt-only baselines use a single-aspect prompt protocol (4,020 slots total). Per-dimension denominators differ accordingly and are not interchangeable across methods. An earlier version of this computation included slots with empty gold labels; all figures in the paper use the corrected denominator.

Ablation slice. The ablation runs on an English slice of 660 single-aspect samples (132 five-dimension groups). Aspect-level EM scores each sample by exact label match. Persona-level EM scores each group by requiring all five dimensions to match simultaneously; for Full SDPM this is evaluated on the output of $\mathcal{R}$, and for *w/o Recomposition* on the union of independent predictions. An output is *parseable* if it decodes into a valid schema-conforming structure; EM $\Rightarrow$ parseable by definition, so EM_rate $\leq$ parseable_rate always holds. The ablation EM figures ($\approx$70–77%) are not comparable to the main-experiment average aspect accuracy ($\approx$49%), as the two protocols differ in slice composition, trajectory length, and inference setup.

Downstream aligned subset. The aligned subset (n=121) retains only users for which all four persona conditions produce parseable output on the same review. Persona states come from the SDPM or Pure-LLM pipeline and do not represent a gold upper bound. No significance tests are reported; the $\approx$0.003 F1 gap between SDPM-$\mathbf{p}_{14}$ and Pure-LLM-$\tilde{\mathbf{p}}_{14}$, together with the reversed ROC-AUC ordering, should be interpreted under the sampling uncertainty appropriate to this scale.

D Short-Horizon Persona Forecasting

To assess the auto-regressive predictability of persona dynamics, we evaluate a model that predicts $\mathbf{p}_{t+1}$ from $\mathbf{p}_t$ alone, without access to any new review text.

Table 7: Short-horizon persona prediction. *Forecast*: model input is $\mathbf{p}_t$ only; *Oracle*: gold $\mathbf{p}_t$ is provided as input; *Persistence*: predict $\mathbf{p}_{t+1} = \mathbf{p}_t$.

Metric	Forecast	Oracle	Persistence
Macro Acc. (%)	49.55	95.37	49.85

The forecast model nearly matches the persistence baseline (49.55% vs. 49.85%), while the oracle ceiling (95.37%) confirms that the label space is discriminative. The near-zero gap between forecast and persistence indicates that persona transitions carry no auto-regressive structure that can be exploited without new observations: predicting $\mathbf{p}_{t+1} = \mathbf{p}_t$ is essentially optimal in the absence of incoming review text. This result motivates the observation-driven design of SDPM's sequential updating, and places a clear bound on what state-only extrapolation can achieve.