

HumanLLM: Benchmarking and Improving LLM Anthropomorphism via Human Cognitive Patterns

Xintao Wang^{1*}, Jian Yang^{1*}, Weiyuan Li¹, Rui Xie¹, Jen-tse Huang³, Jun Gao², Shuai Huang², Yueping Kang², Yuanli Guo¹, Hongwei Feng^{1†}, Yanghua Xiao^{1†}

¹Fudan University ²Hello Group ³Johns Hopkins University

{xtwang21, 24210240375, 25210980069, 25210980167}@m.fudan.edu.cn

{hwfeng, shawyh, guoyuanli}@fudan.edu.cn

gaojun55@gmail.com, {huang.shuai, kang.yueping}@hellogroup.com, jhuan236@jh.edu

Abstract

Large Language Models (LLMs) have demonstrated remarkable capabilities in reasoning and generation, serving as the foundation for advanced persona simulation and Role-Playing Language Agents (RPLAs). However, achieving authentic alignment with human cognitive and behavioral patterns remains a critical challenge for these agents. We present HumanLLM, a framework treating psychological patterns as interacting causal forces. We construct 244 patterns from ~12,000 academic papers and synthesize 11,359 scenarios where 2–5 patterns reinforce, conflict, or modulate each other, with multi-turn conversations expressing inner thoughts, actions, and dialogue. Our dual-level checklists evaluate both individual pattern fidelity and emergent multi-pattern dynamics, achieving strong human alignment ($r = 0.90$) while revealing that holistic metrics conflate simulation accuracy with social desirability. HumanLLM-8B outperforms Qwen3-32B on multi-pattern dynamics despite 4× fewer parameters, demonstrating that authentic anthropomorphism requires cognitive modeling—simulating not just what humans do, but the psychological processes generating those behaviors. Our dataset, code, and model are available at: <https://github.com/YJGoodbye2024/HumanLLM>

1 Introduction

With the rapid scaling of training data, Large Language Models (LLMs) have achieved remarkable progress in anthropomorphism—simulating human-like characteristics and social phenomena (Shanahan et al., 2023). Role-Playing Language Agents (RPLAs) have evolved from conceptual frameworks into practical applications (Chen et al., 2024a), enabling digital clones (Xu et al., 2024a), AI companions (Zhang et al., 2025), and

* Equal contribution.

† Corresponding authors.

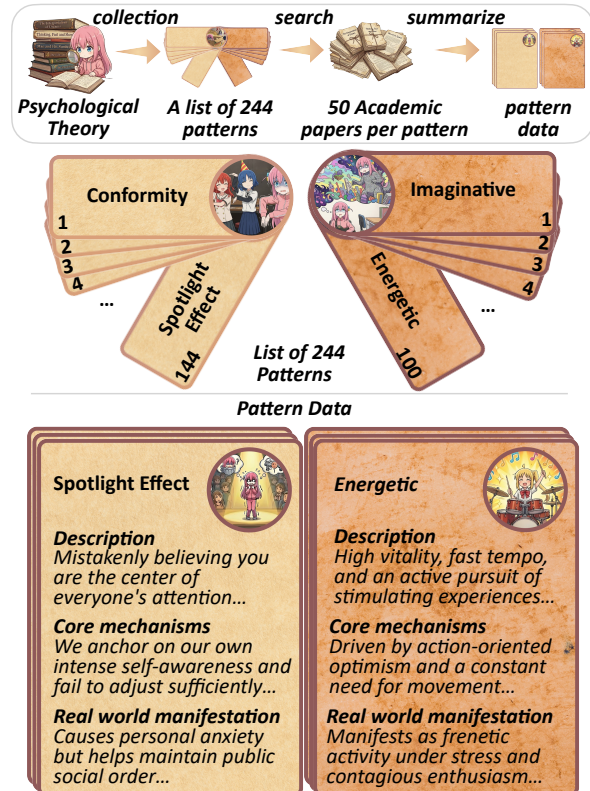


Figure 1: **Pattern Data Structure:** 144 Social-Cognitive Patterns (left) and 100 Personality Traits (right). Each pattern comprises Definition, Core Mechanisms, and Real-World Manifestations.

society simulation (Park et al., 2023; Zhou et al., 2025a,b). As these applications advance, LLM anthropomorphism increasingly requires moving beyond shallow behavioral mimicry toward deeper cognitive and emotional fidelity—what we term **psychological alignment** (Wang et al., 2024).

However, existing approaches model personality as isolated label-to-behavior mappings—“extroverted” maps to “talkative,” “agreeable” maps to “cooperative”—without capturing how multiple cognitive patterns interact to produce behavior (Huang et al., 2024a). We define a pattern as a psychologically documented regularity in human

cognition or behavior—either a stable personality trait (e.g., "assertive") or a context-triggered social-cognitive process (e.g., "spotlight effect"). In reality, a talkative person may fall silent when the spotlight effect is activated; an assertive individual may yield under conformity pressure. Human behavior emerges from the dynamic interplay of multiple patterns, not from any single trait in isolation. Current methods—whether prompting-based (Serapio-García et al., 2025; Ng et al., 2024), fine-tuning-based (Shao et al., 2023; Zhou et al., 2023), or activation steering (Chen et al., 2025)—all treat traits independently, leading to personality drift and the “personality illusion” where models report traits while behaving inconsistently (Wang et al., 2024; Han et al., 2025).

To address this, we propose HumanLLM framework, treating cognitive patterns not as isolated labels but as interacting causal forces. HumanLLM refers to the overall framework; we use ‘HumanLLM dataset’ and ‘HumanLLM-8B/32B’ when referring specifically to the data artifact and fine-tuned models, respectively. Our key insight: by exposing models to scenarios where multiple patterns reinforce, compete, or conflict, models can implicitly learn multi-pattern dynamics without architectural modifications.

Following Lewin’s field theory (Lewin, 2013), we decompose human cognition into two dimensions: (1) **Personality Traits**—stable individual characteristics, and (2) **Social-Cognitive Patterns**—context-triggered mechanisms. We collect 244 patterns (100 personality traits from Goldberg’s Big Five markers (Goldberg, 1992) and 144 social-cognitive patterns from established psychological research), each developed through systematic review of approximately 50 academic papers (Figure 1). We then construct **11,359 scenarios** involving 2–6 characters, each containing 2–5 patterns that may align (e.g., “self-serving bias” reinforcing “overconfidence effect”), conflict (e.g., “assertive” versus “conformity”), or interact conditionally (e.g., “talkative” suppressed by “spotlight effect”). For each scenario, we synthesize **multi-turn conversations** where each turn comprises inner thoughts, physical actions, and verbal expressions (Figure 2).

To ensure faithful pattern expression and enable systematic evaluation, we design **dual-level checklists**: pattern-level checklists (12–15 items per pattern) capture universal behavioral indicators; scenario-level checklists (2–6 items per char-

acter) specify expected behavioral tendencies under each multi-pattern configuration. Our training pipeline consists of supervised fine-tuning on the synthesized conversations. We evaluate across in-domain, out-of-domain, and mixed settings to assess generalization, with additional validation on external benchmarks including LifeChoice and CroSS-MR (Xu et al., 2024b; Yuan et al., 2024).

Our contributions are as follows: (1) We introduce HumanLLM, a framework that systematically leverages psychological cognitive patterns to enhance LLM anthropomorphism, shifting from isolated trait simulation toward modeling the dynamic interplay of human cognition. (2) We construct a comprehensive dataset comprising 244 patterns and 11,359 scenarios with multi-turn, multi-character conversations. Each pattern is grounded in approximately 50 academic papers (over 12,000 papers in total), ensuring psychological rigor and scientific validity. (3) We propose dual-level checklists that enable systematic evaluation at both pattern-level and scenario-level granularities, providing a principled framework for assessing generalization to unseen psychological patterns.

2 Related works

Recent advances in large language models have catalyzed significant progress in role-playing language agents (RPLAs). Early work established foundational architectures: generative agents with memory, planning, and reflection modules have been employed to simulate human behavior in interactive environments (Park et al., 2023), while Character-LLM (Shao et al., 2023) proposed experience reconstruction to train agents embodying historical figures. Subsequent efforts focused on systematic benchmarking and enhancement: ChatHaruhi (Li et al., 2023) leveraged memory-based dialogue control for fictional characters, and CoSER (Wang et al., 2025) curated authentic dialogues from 771 books using “given-circumstance acting” methodology. For persona induction, three main approaches have emerged: (1) prompting-based methods that assign personality traits through instructions (Serapio-García et al., 2025) (2) fine-tuning approaches that embed personas through training on character-specific data (Shao et al., 2023; Zhou et al., 2023), and (3) activation steering via persona vectors that manipulate neural representations corresponding to specific traits (Chen et al., 2025).

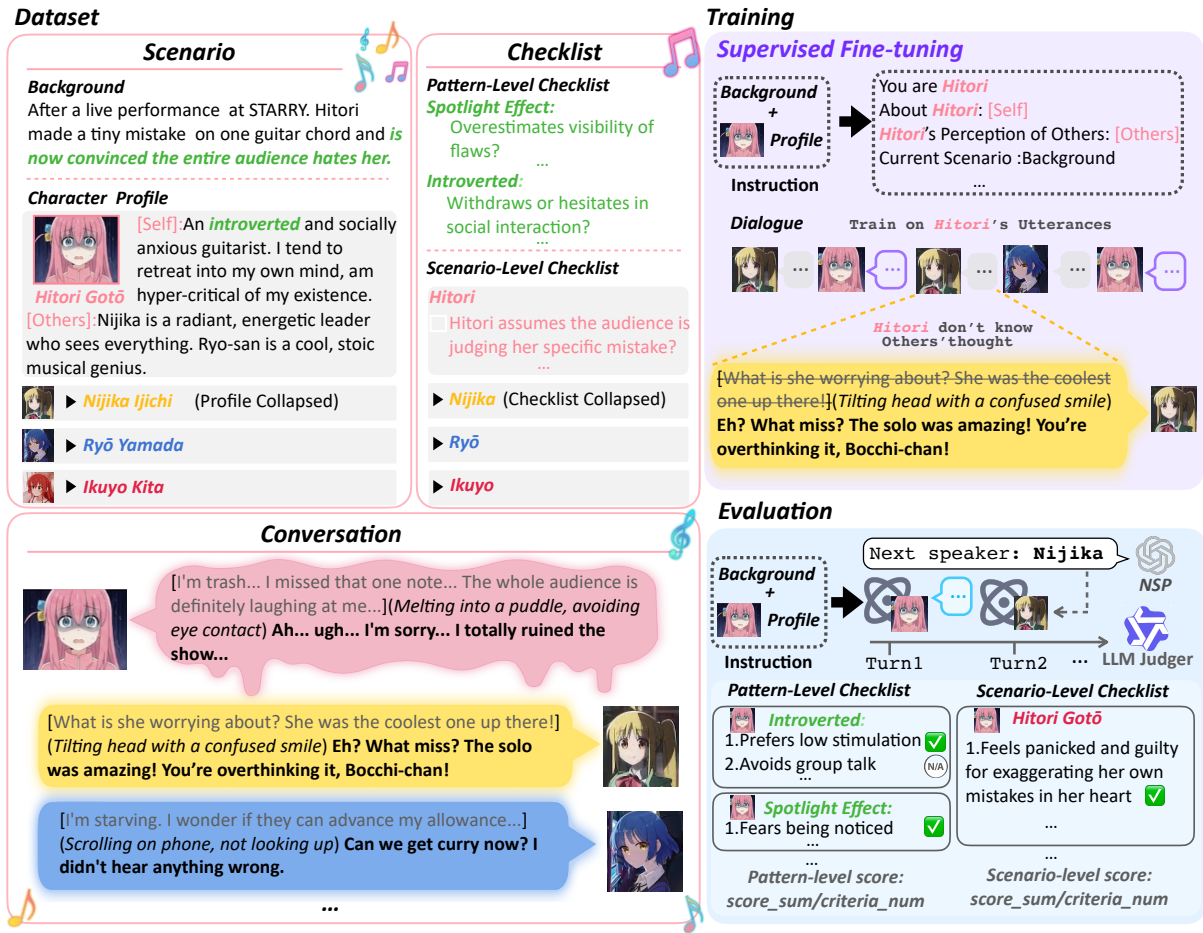


Figure 2: **HumanLLM Framework**. **Left**: Dataset structure with scenarios, multi-turn conversations (inner thoughts in brackets, actions in parentheses), and dual-level checklists. **Top Right**: Supervised fine-tuning on target character utterances. **Bottom Right**: Evaluation via LLM judge scoring against pattern-level and scenario-level checklists.

A parallel line of research evaluates LLMs through the lens of psychological constructs. Theory of Mind (ToM) benchmarks such as ToMBench (Chen et al., 2024b) assess social cognitive abilities, revealing that GPT-4 lags behind humans by over 10%, with trivial task modifications causing significant performance degradation (Ullman, 2023). Emotional intelligence benchmarks (Paech, 2024; Sabour et al., 2024; Huang et al., 2024b,c) adopt psychology-grounded frameworks to evaluate emotional understanding and application, finding substantial gaps between LLMs and humans. Moral reasoning has been assessed through ETHICS (Hendrycks et al., 2023) and MoralBench (Ji et al., 2025), the latter grounded in Moral Foundations Theory. Research on cognitive biases reveals that LLMs exhibit human-like irrationality but with divergent patterns (Macmillan-Scott and Musolesi, 2024). Personality assessment using validated instruments (BFI, MBTI) demonstrates that LLMs can manifest measurable

traits (Pellert et al., 2024), though self-report validity remains questionable (Zou et al., 2025). Critically, recent work cautions that LLMs do not reliably simulate human psychology and fail to generalize across semantically equivalent scenarios (Schröder et al., 2025; Cheung et al., 2025).

3 HumanLLM Dataset

This section introduces the HumanLLM dataset, a psychologically grounded resource for training and evaluating anthropomorphic language models. We describe pattern collection (§3.1), pattern data construction (§3.2), scenario and conversation generation (§3.3), and dual-level checklist design (§3.4). Table 1 summarizes the dataset statistics.

3.1 Pattern Collection

Following the theoretical foundations established in §A.2, we compile patterns along two complementary dimensions corresponding to Lewin’s Person-Environment framework.

Statistic	Value
Total Patterns	244
Scenarios	11,359
Avg. Patterns per Scenario	3.5
Avg. Turns per Conversation	16.4
Pattern-Level Checklist Items	12–15 per pattern
Scenario-Level Checklist Items	2–6 per character

Table 1: HumanLLM dataset statistics. Patterns include 100 personality traits and 144 social-cognitive patterns. Each scenario contains 2–6 characters with multi-turn conversations (12–20 turns).

Personality Traits (Person Dimension) We adopt Goldberg’s 100 Unipolar Markers (Goldberg, 1992), a psychometrically validated lexicon mapping onto the Big Five dimensions with 20 trait descriptors each (Extraversion, Agreeableness, Conscientiousness, Emotional Stability, Intellect).

Social-Cognitive Patterns (Environment Dimension) We curate situationally-activated psychological mechanisms through systematic review of established theoretical traditions, including cognitive biases (Tversky and Kahneman, 1974), social influence (Cialdini, 2009), evolutionary psychology (Buss, 2024), and motivation research (Deci and Ryan, 2000). From an initial pool of 232 documented patterns, we apply two filtering criteria: (1) sufficient empirical validation, and (2) non-redundancy with other patterns. This yields 144 social-cognitive patterns (full taxonomy in Appendix B.1).

3.2 Pattern Data Construction

Pattern data are structured representations of psychological patterns, as illustrated in Figure 1. We construct pattern data through a two-stage pipeline: literature retrieval followed by LLM-based synthesis.

Literature Retrieval For each of the 244 patterns, we employ Gemini Deep Search to identify approximately 50 relevant academic papers. The search is guided by three retrieval dimensions: (1) foundational definitions from seminal works, (2) mechanistic explanations from theoretical and empirical studies, and (3) real-world applications from applied research. Retrieved references are filtered manually to remove irrelevant entries. Full-text documents are obtained through open-access APIs (Semantic Scholar, arXiv, OpenAlex, PubMed, Crossref); when full text is unavailable, abstracts

are retained. This process yields a corpus of approximately 12,000 papers across all patterns.

Pattern Synthesis We employ Gemini 2.5 Pro to summarize each pattern’s literature corpus into a structured representation. Critically, the model is instructed to extract and summarize information *exclusively* from the provided 50 papers, rather than generating content from its parametric knowledge. Following the construct validity framework (§A.2), each pattern is organized into three components: (1) **Definition**—a precise characterization grounded in authoritative sources; (2) **Core Mechanisms**—underlying cognitive, emotional, and behavioral processes that drive the pattern; (3) **Real-World Manifestations**—ecological expressions across diverse contexts (e.g., response to stress, interpersonal dynamics, professional settings). Details of the retrieval and synthesis prompts are provided in Appendix F. To validate synthesis quality, three psychology-trained annotators independently evaluated a stratified sample of 30 pattern entries (15 personality traits, 15 social-cognitive patterns). For each pattern, annotators received the LLM-generated summary alongside the complete set of ~50 source papers and scored five dimensions following our construct validity framework (Cronbach and Meehl, 1955): definitional accuracy, mechanistic fidelity, manifestation coverage, source faithfulness, and construct distinctiveness (4-point Likert scale). Mean scores ranged from 3.20 (manifestation coverage) to 3.70 (definitional accuracy), with substantial inter-annotator agreement on most dimensions (Krippendorff’s α (Krippendorff, 2011) = 0.58–0.76). Personality traits scored moderately higher than social-cognitive patterns (3.51 vs. 3.33), consistent with the more canonical nature of Big Five constructs. Full protocol and results are reported in Appendix B.6.

3.3 Scenario and Conversation Generation

We synthesize scenarios and conversations through a two-stage pipeline using large language models.

Scenario Synthesis Each scenario comprises a narrative background and 2–6 character profiles. Character profiles contain *self-perception* (identity, personality, background, motivations) and *other-perception* (knowledge and attitudes toward other characters), enabling realistic information asymmetry. Each scenario incorporates 2–5 patterns, with pattern combinations validated to filter semantically contradictory configurations. To ensure sit-

uational diversity, we leverage the DIAMONDS model (§A.2) to generate scenario variants across different situational dimensions. Alongside each scenario, we synthesize *expected behavioral tendencies*—specifications of how characters should manifest the target patterns within the given context. Scenarios are generated using Gemini 2.5 Pro and Claude Sonnet 4.5, each contributing approximately half of the total. This process yields 11,359 scenarios.

Conversation Synthesis Based on scenarios and expected behavioral tendencies, we synthesize multi-turn conversations (12–20 turns) using Claude Sonnet 4.5. Each turn comprises three dimensions: inner thoughts (enclosed in brackets), physical actions (enclosed in parentheses), and verbal expressions (Figure 2). Target patterns are naturally embedded across these dimensions, guided by the expected behavioral tendencies. Detailed generation prompts are provided in Appendix F.

3.4 Dual-Level Checklist Design

We design dual-level checklists for fine-grained assessment of pattern expression, serving both evaluation (§4.2) and future reward modeling purposes.

Scenario-Level Checklist. The expected behavioral tendencies synthesized alongside each scenario directly constitute the scenario-level checklist. These items specify context-specific behaviors derived from particular pattern combinations and situational contexts (e.g., “Character defends conceptual integrity under deadline pressure, resisting shortcuts”). Each scenario contains 2–6 characters, but only characters assigned with target patterns receive checklist items; each such character has 2–6 items reflecting their expected behavioral tendencies. During evaluation, we assess one target character per evaluation instance, treating other characters as contextual interlocutors.

Pattern-Level Checklist For each of the 244 patterns, we construct 12–15 universal behavioral indicators through iterative refinement: initial generation from pattern structure, validation against conversation samples, and generalization to ensure cross-context applicability (e.g., “Shows heightened awareness of being observed” for spotlight effect). These pattern-level items are context-independent and apply to any character exhibiting the target pattern, complementing the scenario-specific items above. The complete construction

procedure is detailed in Appendix B.4.

4 Training and Evaluation

This section presents the training procedure (§4.1) and evaluation protocol (§4.2) for HumanLLM models.

4.1 Training

We train HumanLLM-8B and HumanLLM-32B on Qwen3-8B and Qwen3-32B respectively through supervised fine-tuning. Training data composition and hyperparameters are detailed in Appendix C.

Supervised Fine-Tuning We convert each character’s dialogue turns within a scenario into a separate training sample in ShareGPT format, yielding 30,543 HumanLLM samples from 10,265 training scenarios. The model is trained to generate responses that naturally express the target patterns through the trinity of expression: inner thoughts (square brackets), physical actions (parentheses), and verbal expressions. To maintain general capabilities and role-playing abilities, we augment the HumanLLM dataset with two complementary sources: OpenThoughts-114k for instruction-following capabilities (30,543 samples) and CoSER (Wang et al., 2025) for role-playing dialogue (15,272 samples). The final training mixture comprises 76,358 samples with a ratio of 4:4:2 (HumanLLM: OpenThoughts : CoSER) by sample count.

4.2 Evaluation

The dual-level checklists enable evaluation at two complementary granularities, corresponding to two proposed metrics. Both metrics employ GPT-5-mini as judge with ternary scoring: +1 (satisfied), 0 (not exhibited), –1 (violated). We select GPT-5-mini rather than Gemini to avoid evaluator bias, as Gemini 2.5 Pro is used in data generation. Reported percentages are obtained by averaging ternary scores per sample, then across samples, and linearly mapping from $[-1, +1]$ to $[-100\%, 100\%]$. To assess multi-run stability, we run the judge three times on each evaluation instance and report mean \pm std in Table 2; standard deviations remain below 2.1 percentage points across all models, confirming stable rankings.

Metrics We propose two metrics that capture complementary aspects of pattern expression: (1) *Individual Pattern Expression (IPE)* measures

whether each pattern is expressed according to its psychological definition. IPE uses pattern-level checklists, which contain 12–15 universal behavioral indicators derived directly from each pattern’s definition and mechanisms (e.g., for *spotlight effect*: “Overestimates others’ attention to own appearance”). These indicators are context-independent—they assess whether the pattern’s characteristic behaviors appear, regardless of situational details. The sample-level IPE is computed as the mean score across all pattern-level checklist items; the overall IPE is the mean across all evaluation samples. (2) *Multi-Pattern Dynamics (MPD)* measures whether multiple patterns interact appropriately within specific situational configurations. Prior role-playing evaluations often rely on single-label assessments (e.g., “Is this character assertive?”), which reduce complex personalities to stereotypical behaviors and fail to capture how traits modulate each other in context. MPD addresses this limitation by using scenario-level checklists that specify expected behavioral tendencies emerging from particular pattern combinations. For instance, a character assigned both *assertive* and *spotlight effect* in a public speaking scenario should not simply exhibit “confident speech”—the checklist captures the nuanced tension: “Projects outward confidence in expressing opinions, yet harbors internal anxiety regarding audience scrutiny.” Such items reflect how one pattern (spotlight effect) constrains or modulates another (assertive), moving evaluation beyond stereotypical label-to-behavior mappings toward authentic multi-dimensional characterization. The sample-level MPD is computed as the mean score across all scenario-level checklist items; the overall MPD is the mean across all evaluation samples. In summary, IPE evaluates fidelity to individual pattern definitions, while MPD evaluates the emergent dynamics when multiple patterns interact—capturing the psychological realism that single-label evaluations miss. Data splits for generalization assessment are detailed in Appendix D.1.

5 Experiments

5.1 Experimental Setup

Training Configuration We train HumanLLM-8B and HumanLLM-32B through supervised fine-tuning as described in §4.1. Detailed hyperparameters are provided in Appendix C.

Baselines We compare against a comprehensive set of baselines spanning proprietary and

Model	IPE	MPD
<i>Closed-Source</i>		
GPT-5	15.5±0.4	43.4±1.1
Claude Sonnet 4.5	34.8±0.3	79.5±0.4
Gemini 3 Pro	41.3±0.3	85.1±0.4
<i>Open-Source</i>		
Qwen3-8B	18.6±0.7	54.4±2.1
Qwen3-32B	26.0±0.4	65.8±0.7
Qwen3-235B	34.3±0.5	72.5±0.4
DeepSeek-V3.2	22.2±0.8	65.1±0.5
DeepSeek-R1	23.3±0.6	69.0±0.5
<i>Ours</i>		
HumanLLM-8B	25.7±0.4	70.3±0.6
HumanLLM-32B	32.8±0.3	73.6±0.4

Table 2: Main results on IPE (Individual Pattern Expression) and MPD (Multi-Pattern Dynamics). Higher values indicate better performance. All values are reported as mean±std (%) across 3 independent GPT-5-mini judge runs. Results are averaged across ID_eval, OOD_eval, and Mixed_eval splits.

open-source models, all evaluated in zero-shot settings: (1) *Closed-source models*, including GPT-5 (OpenAI, 2025), Claude Sonnet 4.5 (Anthropic, 2025), and Gemini 3 Pro (DeepMind, 2025). (2) *Open-source models*, including Qwen3-8B/32B/235B (Yang et al., 2025), DeepSeek-V3.2 (DeepSeek-AI et al., 2025b), and DeepSeek-R1 (DeepSeek-AI et al., 2025a).

External Benchmarks Beyond our proposed IPE and MPD metrics, we evaluate on established role-playing benchmarks: (1) *LifeChoice* (Xu et al., 2024b): evaluates persona-driven decision-making through life choice scenarios; (2) *CroSS-MR* (Yuan et al., 2024): assesses character motivation recognition; Additionally, we include CoSER’s evaluation metrics—Anthropomorphism and Character Fidelity (Wang et al., 2025)—for comparative analysis of evaluation paradigms (§5.5).

5.2 Main Results

Table 2 presents the main experimental results on our proposed IPE and MPD metrics.

Overall Performance Among all evaluated models, Gemini 3 Pro achieves the highest scores on both metrics (IPE: 41.3%, MPD: 85.1%), followed by Claude Sonnet 4.5 (IPE: 34.8%, MPD: 79.5%). Our HumanLLM-32B achieves competitive performance (IPE: 32.8%, MPD: 73.6%), outperforming all open-source baselines of comparable or larger scale. Notably, HumanLLM-8B (IPE: 25.7%, MPD:

70.3%) surpasses Qwen3-32B (IPE: 26.0%, MPD: 65.8%) on MPD despite having 4× fewer parameters, demonstrating the effectiveness of our psychologically grounded training data.

Close-Source vs. Open-Source Gap A substantial performance gap exists between close-source and open-source models. The best close-source model (Gemini 3 Pro) outperforms the best open-source baseline (Qwen3-235B) by +7.0 on IPE and +12.6 on MPD. Interestingly, GPT-5 exhibits unexpectedly low performance (IPE: 15.5%, MPD: 43.4%), ranking below most open-source alternatives. Qualitative analysis suggests that GPT-5’s strong instruction-following tendency leads to overly literal interpretations of role-playing prompts, resulting in shallow pattern expression. This finding aligns with recent observations that general-purpose capabilities do not automatically transfer to nuanced psychological simulation (Wang et al., 2024).

Scaling Effects Within model families, we observe consistent scaling improvements. For Qwen3, IPE increases from 18.6 (8B) to 26.0 (32B) to 34.3 (235B), representing a +84% relative improvement from 8B to 235B. Similarly, our HumanLLM models show scaling gains: HumanLLM-32B outperforms HumanLLM-8B by +7.1 on IPE (+28%) and +3.3 on MPD (+5%). The relatively smaller MPD scaling gap suggests that multi-pattern dynamics may be more sample-efficient to learn, while individual pattern expression benefits more from increased model capacity.

IPE vs. MPD Dynamics Across all models, MPD scores consistently exceed IPE scores. This asymmetry likely reflects the evaluation granularity: MPD assesses scenario-level behavioral tendencies (2–6 items per character), while IPE evaluates against 12–15 fine-grained pattern indicators. The gap is most pronounced for weaker models (e.g., GPT-5: 15.5% vs. 43.4%), suggesting that models can produce superficially coherent multi-pattern behavior without deeply understanding individual pattern mechanisms.

5.3 Ablation Study

To isolate the contribution of our psychologically grounded dataset, we conduct an ablation study comparing three model variants (Table 3).

Effect of HumanLLM dataset The comparison between HumanLLM-8B and Qwen3-8B (base) reveals

Model Variant	IPE	MPD
Qwen3-8B (base)	18.6±0.7	54.4±2.1
Qwen3-8B (OT+CoSER)	9.1±1.2	31.3±2.9
HumanLLM-8B	25.7±0.4	70.3±0.6

Table 3: Ablation study results (%). “OT+CoSER” denotes training on OpenThoughts and CoSER without HumanLLM dataset.

substantial improvements: +7.1 on IPE (+38%) and +15.9 on MPD (+29%). These gains demonstrate that our synthesized conversations, grounded in psychological pattern definitions, effectively teach models to express cognitive and behavioral patterns with higher fidelity.

Negative Transfer from Generic Data Surprisingly, the OT+CoSER variant (trained on OpenThoughts and CoSER without HumanLLM dataset) performs *worse* than the base model on both metrics: IPE drops from 18.6 to 9.1 (−51%), and MPD drops from 54.4 to 31.3 (−42%). This counterintuitive result suggests negative transfer: generic instruction-following data (OpenThoughts) and conventional role-playing data (CoSER) may inadvertently suppress the base model’s latent ability to simulate psychological patterns. We hypothesize that these datasets reinforce “helpful assistant” behaviors that conflict with authentic expression of cognitively biased or emotionally complex characters.

Synergistic Effect The full HumanLLM-8B model, trained on the combined mixture (HumanLLM + OT + CoSER with 4:4:2 ratio), achieves the best performance. This indicates that HumanLLM dataset not only compensates for the negative transfer but creates a synergistic effect where psychological grounding enhances the utility of general-purpose and role-playing data. The pattern-rich conversations appear to serve as “anchors” that prevent the model from collapsing toward generic prosocial behaviors.

5.4 External Benchmark Evaluation

To assess generalization beyond our proposed metrics, we evaluate on established role-playing benchmarks (Table 4). On LifeChoice, while closed-source models (e.g., GPT-5: 85.53%) dominate, HumanLLM-32B (50.64%) still outperforms Qwen3-32B (47.71%) by +2.93%; this modest improvement suggests that LifeChoice’s binary decision

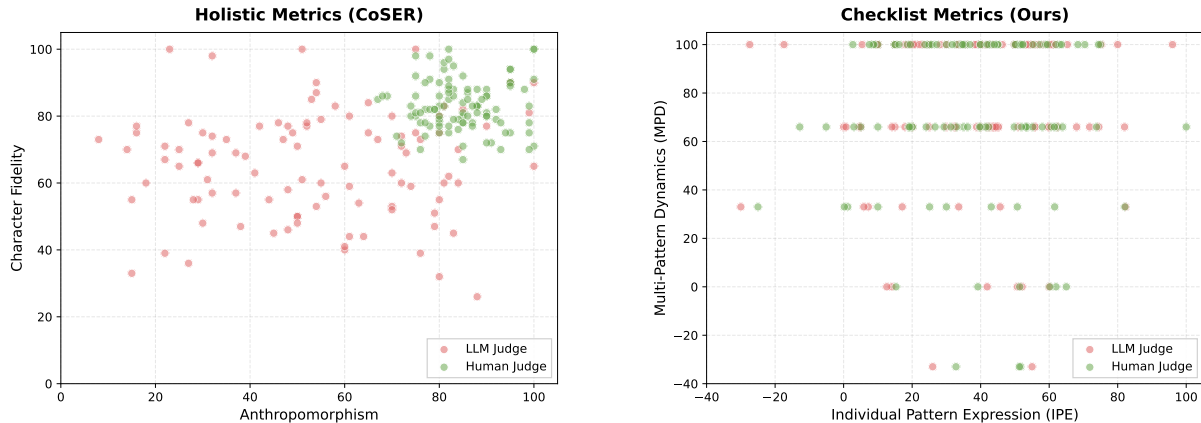


Figure 3: Human-LLM evaluation alignment comparison. **Left:** Holistic metrics show clear separation between LLM and human judgments, with systematic LLM underestimation of psychologically complex behaviors. **Right:** Our checklist metrics demonstrate strong overlap between LLM and human distributions, indicating robust alignment with expert judgment

Model	LifeChoice	CroSS-MR
GPT-5	85.53	62.25
Claude Sonnet 4.5	85.49	68.24
Qwen3-8B	44.51	53.26
Qwen3-32B	47.71	63.37
Qwen3-8B (OT+CoSER)	46.99	54.98
HumanLLM-8B	47.19	54.23
HumanLLM-32B	50.64	64.27

Table 4: External benchmark results (%). LifeChoice evaluates persona-driven decision-making; CroSS-MR assesses character motivation recognition.

scenarios may rely more on surface-level reasoning than the deep psychological patterns targeted by our approach. Conversely, on CroSS-MR, HumanLLM-32B (64.27%) surpasses both Qwen3-32B (63.37%) and GPT-5 (62.25%), with the 8B variant also showing competitive performance (54.23% vs. 53.26%), indicating that motivation recognition benefits from our training but is less dependent on complex pattern simulation than decision-making tasks.

Benchmark Limitations The moderate improvements on external benchmarks, contrasted with substantial gains on our IPE/MPD metrics, reveal a fundamental mismatch between existing evaluation paradigms and psychologically grounded simulation. LifeChoice and CroSS-MR evaluate behavioral outcomes (decisions, motivations) rather than the cognitive processes underlying those outcomes. These benchmarks, while valuable for assessing surface-level role-playing, do not capture the pattern-specific behavioral indicators that de-

Metric	Human	LLM	Δ	r
<i>Holistic Metrics (CoSER)</i>				
Anthropomorphism	84.6	53.8	-30.8	0.43
Character Fidelity	83.1	65.4	-17.7	0.61
<i>Checklist Metrics (Ours)</i>				
IPE	38.4	37.8	-0.6	0.90
MPD	72.1	75.8	+3.7	0.88

Table 5: Human-LLM agreement analysis across 100 sampled scenarios. Holistic metrics (0–100 scale) show weak correlation and large systematic bias; our checklist metrics ($[-1, +1]$ normalized to $[-100\%, 100\%]$) achieve significantly higher alignment with human expert judgment.

fine authentic psychological simulation. This observation motivates our development of the dual-level checklist framework, which we validate in §5.5.

5.5 Evaluation Framework Analysis

To validate our dual-level checklist, we compared automated GPT-5-mini evaluations against three human experts across 100 diverse scenarios using both holistic metrics (CoSER) and our checklist criteria. Evaluations were voluntary and anonymous.

Human-LLM Alignment and Normative Confounding As shown in Table 5, holistic metrics exhibit weak correlation and large systematic bias (Anthropomorphism: $r = 0.43$, $\Delta = -30.8$; Character Fidelity: $r = 0.61$, $\Delta = -17.7$). This divergence stems from what we term **normative confounding**: LLM judges implicitly conflate “good anthropomorphism” with “prosocial behavior” (e.g., empathy, rationality), penalizing real-

istic but negative human traits like defensiveness. In contrast, our checklist metrics achieve robust alignment (IPE: $r = 0.90$, $\Delta = -0.6$; MPD: $r = 0.88$, $\Delta = +3.7$) by decomposing behaviors into value-neutral indicators. This effectively decouples *simulation accuracy* from *social desirability*.

Case Study Ideally illustrating this, a scenario featuring *ultimate attribution error* (defensively blaming out-groups) received a low holistic score (5/100) due to cited “lack of empathy.” However, our checklist correctly validated the pattern through definition-grounded questions (e.g., “Does the character attribute failure to external factors?”), confirming that the behavior, while socially undesirable, was psychologically accurate (Appendix E).

6 Conclusion

We present HumanLLM, a framework leveraging 244 patterns from $\sim 12,000$ papers and 11,359 scenarios to enable dynamic cognitive interactions rather than isolated trait mappings. Our evaluation achieves strong human alignment ($r = 0.90$) and identifies **normative confounding**, where conventional metrics mistake social desirability for simulation accuracy. Experiments confirm the necessity of psychological grounding for authentic expression, advancing the field from surface-level behavioral mimicry to cognitive modeling—simulating the underlying psychological processes that generate human behavior.

Limitations

Conversation Turn Horizon The current dataset contains conversations averaging 16.4 turns (range: 12–20). While role-playing is fundamentally a long-horizon task, evaluating ultra-long psychological simulations (e.g., 50+ turns) remains an open challenge due to the lack of reliable automated metrics that maintain consistency over extended contexts. However, we infer that our conversation length is adequate for establishing psychological fidelity: prior work on long-context alignment (Bai et al., 2024) suggests that models fine-tuned on 10–15 turns of high-quality conversational data can generalize effectively to much longer inference horizons, as the supervised fine-tuning phase establishes the necessary cognitive “anchor” for sustained pattern expression. Future work should explore explicit long-horizon evaluation protocols.

Cultural Adaptability The psychological theories underlying our pattern taxonomy originate predominantly from WEIRD (Western, Educated, Industrialized, Rich, Democratic) populations (Henrich et al., 2010). Pattern expressions—particularly social-cognitive mechanisms such as conformity, authority bias, and motivational processes—may manifest differently across cultural contexts, and our current dataset may inadvertently embed specific cultural norms as universal psychological truths. To address this in future iterations, we propose two concrete steps: first, integrating indigenous psychology research (e.g., East Asian relational cognition, Ubuntu-based African social frameworks) into the Pattern Construction stage; second, grounding the Manifestation layer in cross-cultural empirical datasets such as the World Values Survey, ensuring that behavioral expressions reflect local cultural contexts rather than projecting WEIRD assumptions globally.

LLM-as-Judge Limitations Our evaluation relies on GPT-5-mini as the judge. While we mitigate variability through three independent runs (standard deviation $< 2.1\%$), LLM judges may exhibit systematic biases or inconsistent reasoning on psychologically complex scenarios (Li et al., 2025). Future work should incorporate broader human evaluation across a larger and more diverse sample.

Synthetic Data Limitations All training conversations are synthetically generated by LLMs, which may introduce systematic biases or fail to capture the full complexity of authentic human interactions. Although our human validation study confirms high source faithfulness (mean = 3.50, $\alpha = 0.73$), the gap between synthetic and naturalistic discourse remains a limitation for downstream deployment.

Evaluation Scope Our evaluation focuses on text-based dialogue and may not fully capture temporal consistency across extended interactions or embodied behavior in multimodal settings.

Ethical Statement

Safety–Fidelity Tension. By design, HumanLLM improves the simulation of human cognitive patterns, including irrational biases, maladaptive coping mechanisms, and negative personality traits such as antagonism and social cynicism. While this fidelity is essential for realistic role-playing, social simulation research, and companion AI development, it creates a fundamental tension with safety

alignment: a model trained to authentically express cognitively biased or antisocial behaviors may generate harmful, manipulative, or toxic content if deployed outside controlled research contexts. We recommend that downstream applications built on HumanLLM-fine-tuned models incorporate appropriate safety layers, particularly in consumer-facing contexts. The dataset and models are released to support open research, but deployment beyond research contexts should include safety alignment fine-tuning and usage monitoring.

Manipulation and Social Engineering. The HumanLLM framework explicitly models psychological mechanisms of social influence and persuasion—including authority bias, reciprocity, conformity, and obedience to authority. A model that deeply and faithfully simulates these mechanisms possesses heightened capacity for social engineering: it may craft persuasive narratives, exploit cognitive vulnerabilities, or systematically nudge users toward unintended beliefs or actions. Responsible deployment requires continuous monitoring for persuasion-oriented misuse, and we advise against integrating HumanLLM-fine-tuned models into systems with asymmetric power dynamics (e.g., automated negotiation, political messaging, or customer manipulation contexts).

Parasocial Attachment and Vulnerable Populations. Highly anthropomorphic agents risk fostering parasocial relationships in which users attribute genuine emotions, intentions, or moral status to the model. This risk is amplified by HumanLLM’s explicit training on inner thought simulation, which creates the appearance of subjective experience. Such dynamics are especially concerning for vulnerable populations, including children, individuals with social anxiety, and those experiencing loneliness or grief, who may form attachments that substitute for—rather than supplement—human connection. We advocate for transparent, persistent disclosure of AI identity in all deployment contexts, and recommend that companion-AI applications built on this framework incorporate regular check-ins that reinforce the non-human nature of the agent.

Stereotype Amplification. Pattern-based cognitive modeling carries a risk of reinforcing harmful stereotypes if pattern assignments become correlated with demographic attributes (e.g., gender, ethnicity, or nationality) in downstream applications.

Our pattern taxonomy is demographic-neutral by design: patterns are assigned to characters based on scenario requirements, not identity categories. Nonetheless, the source psychological literature itself reflects historical sampling biases, and certain patterns (e.g., dominance hierarchies, asymmetrical parental investment) carry cultural assumptions that may produce biased character portrayals. We urge downstream developers to conduct disparate-impact audits before deploying HumanLLM-based systems in any context involving user-facing persona generation.

Data and Model Transparency. All data, code, and model weights are publicly released to support reproducibility and community scrutiny. We encourage independent audits of the dataset for unintended biases and welcome community contributions to expand the pattern taxonomy toward greater cultural inclusivity, as discussed in the Limitations section. Researchers who adapt HumanLLM for new applications bear responsibility for ensuring that such adaptations comply with the ethical standards of their respective deployment contexts.

Acknowledgments

This work is supported by Hello Group Inc. We thank the anonymous reviewers and the action editor for their valuable feedback and constructive comments that helped improve this paper. We also express our deep appreciation to the pioneering researchers in psychology whose foundational theories on human personality and cognition made this work possible. Finally, we are grateful to Dr. Fang Guo at Fudan University for the insightful discussions and valuable feedback throughout this work.

References

- Anthropic. 2025. Anthropic’s transparency hub: Model report. <https://www.anthropic.com/transparency/model-report>. Accessed: 2026-01-03.
- Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. 2024. *LongAlign: A recipe for long context alignment of large language models*. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1376–1395, Miami, Florida, USA. Association for Computational Linguistics.
- David M Buss. 2024. *Evolutionary psychology: The new science of the mind*. Routledge.

- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024a. [From Persona to Personalization: A Survey on Role-Playing Language Agents](#). *Preprint*, arXiv:2404.18231.
- Runjin Chen, Andy Arditi, Henry Sleight, Owain Evans, and Jack Lindsey. 2025. [Persona vectors: Monitoring and controlling character traits in language models](#). *Preprint*, arXiv:2507.21509.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, and Minlie Huang. 2024b. [ToMBench: Benchmarking Theory of Mind in Large Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15959–15983, Bangkok, Thailand. Association for Computational Linguistics.
- Vanessa Cheung, Maximilian Maier, and Falk Lieder. 2025. [Large language models show amplified cognitive biases in moral decision-making](#). *Proceedings of the National Academy of Sciences*, 122(25):e2412015122.
- Robert B Cialdini. 2009. *Influence: Science and practice*, volume 4. Pearson education Boston.
- Lee J Cronbach and Paul E Meehl. 1955. Construct validity in psychological tests. *Psychological bulletin*, 52(4):281.
- Edward L Deci and Richard M Ryan. 2000. The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological inquiry*, 11(4):227–268.
- Google DeepMind. 2025. [Gemini 3 pro model card](#). Model card, Google DeepMind.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025a. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, BOWEI Zhang, Chaofan Lin, Chen Dong, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenghao Xu, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, and 245 others. 2025b. [Deepseek-v3.2: Pushing the frontier of open large language models](#). *Preprint*, arXiv:2512.02556.
- John M. Digman. 1990. [Personality structure: Emergence of the five-factor model](#). *Annual Review of Psychology*, 41:417–440.
- Lewis R Goldberg. 1992. The development of markers for the big-five factor structure. *Psychological assessment*, 4(1):26.
- Pengrui Han, Rafal Kocielnik, Peiyang Song, Ramit Debnath, Dean Mobbs, Anima Anandkumar, and R. Michael Alvarez. 2025. [The personality illusion: Revealing dissociation between self-reports & behavior in llms](#). *Preprint*, arXiv:2509.03730.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2023. [Aligning ai with shared human values](#). *Preprint*, arXiv:2008.02275.
- Joseph Henrich, Steven J. Heine, and Ara Norenzayan. 2010. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3):61–83.
- Jen-tse Huang, Wenxiang Jiao, Man Ho Lam, Eric John Li, Wenxuan Wang, and Michael Lyu. 2024a. On the reliability of psychological scales on large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6152–6173.
- Jen-tse Huang, Man Ho Lam, Eric John Li, Shujie Ren, Wenxuan Wang, Wenxiang Jiao, Zhaopeng Tu, and Michael R Lyu. 2024b. [Apathetic or empathetic? evaluating llms' emotional alignments with humans](#). *Advances in Neural Information Processing Systems*, 37:97053–97087.
- Jen-tse Huang, Wenxuan Wang, Eric John Li, Man Ho Lam, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, and Michael Lyu. 2024c. [On the humanity of conversational ai: Evaluating the psychological portrayal of llms](#). In *The Twelfth International Conference on Learning Representations*.
- Jianchao Ji, Yutong Chen, Mingyu Jin, Wujiang Xu, Wenyue Hua, and Yongfeng Zhang. 2025. [Moralbench: Moral evaluation of llms](#). *Preprint*, arXiv:2406.04428.
- Klaus Krippendorff. 2011. [Computing krippendorff's alpha-reliability](#).
- Kurt Lewin. 2013. *Principles of topological psychology*. Read Books Ltd.
- Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi MI, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, Linkang Zhan, Yaokai Jia, Pingyu Wu, and Haozhen Sun. 2023. [Chatharuhi: Reviving anime character in reality via large language model](#). *Preprint*, arXiv:2308.09597.
- Weiyuan Li, Xintao Wang, Siyu Yuan, Rui Xu, Jiangjie Chen, Qingqing Dong, Yanghua Xiao, and Deqing Yang. 2025. [Curse of knowledge: When complex evaluation context benefits yet biases llm judges](#). *Preprint*, arXiv:2509.03419.
- Olivia Macmillan-Scott and Mirco Musolesi. 2024. [\(ir\)rationality and cognitive biases in large language models](#). *Preprint*, arXiv:2402.09193.

- 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. *arXiv preprint arXiv:2404.13957*.
- OpenAI. 2025. *Gpt-5 system card*. System card, OpenAI.
- Samuel J. Paech. 2024. *Eq-bench: An emotional intelligence benchmark for large language models*. *Preprint*, arXiv:2312.06281.
- 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*. *Preprint*, arXiv:2304.03442.
- Max Pellert, Clemens M Lechner, Claudia Wagner, Beatrice Rammstedt, and Markus Strohmaier. 2024. Ai psychometrics: Assessing the psychological profiles of large language models through psychometric inventories. *Perspectives on Psychological Science*, 19(5):808–826.
- John F Rauthmann, David Gallardo-Pujol, Esther M Guillaume, Elysia Todd, Christopher S Nave, Ryne A Sherman, Matthias Ziegler, Ashley Bell Jones, and David C Funder. 2014. The situational eight diamonds: a taxonomy of major dimensions of situation characteristics. *Journal of personality and social psychology*, 107(4):677.
- Sahand Sabour, Siyang Liu, Zheyuan Zhang, June Liu, Jinfeng Zhou, Alvionna Sunaryo, Tatia Lee, Rada Mihalcea, and Minlie Huang. 2024. *EmoBench: Evaluating the Emotional Intelligence of Large Language Models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5986–6004, Bangkok, Thailand. Association for Computational Linguistics.
- Sarah Schröder, Thekla Morgenroth, Ulrike Kuhl, Valerie Vaquet, and Benjamin Paaßen. 2025. *Large language models do not simulate human psychology*. *Preprint*, arXiv:2508.06950.
- Greg Serapio-García, Mustafa Safdari, Clément Crepy, Luning Sun, Stephen Fitz, Peter Romero, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2025. *Personality traits in large language models*. *Preprint*, arXiv:2307.00184.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. *Role play with large language models*. *Nature*, 623(7987):493–498.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. *Character-LLM: A Trainable Agent for Role-Playing*. *Preprint*, arXiv:2310.10158.
- Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157):1124–1131.
- Tomer Ullman. 2023. *Large language models fail on trivial alterations to theory-of-mind tasks*. *Preprint*, arXiv:2302.08399.
- Xintao Wang, Heng Wang, Yifei Zhang, Xinfeng Yuan, Rui Xu, Jen-tse Huang, Siyu Yuan, Haoran Guo, Jiangjie Chen, Shuchang Zhou, Wei Wang, and Yanghua Xiao. 2025. *CoSER: Coordinating LLM-Based Persona Simulation of Established Roles*. *Preprint*, arXiv:2502.09082.
- Xintao Wang, Yunze Xiao, Jen-tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li, and Yanghua Xiao. 2024. *InCharacter: Evaluating Personality Fidelity in Role-Playing Agents through Psychological Interviews*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1840–1873, Bangkok, Thailand. Association for Computational Linguistics.
- Rui Xu, Dakuan Lu, Xiaoyu Tan, Xintao Wang, Siyu Yuan, Jiangjie Chen, Wei Chu, and Yinghui Xu. 2024a. *MINDECHO: Role-Playing Language Agents for Key Opinion Leaders*. *arXiv preprint*.
- Rui Xu, Xintao Wang, Jiangjie Chen, Siyu Yuan, Xinfeng Yuan, Jiaqing Liang, Zulong Chen, Xiaoqing Dong, and Yanghua Xiao. 2024b. *Character is destiny: Can role-playing language agents make persona-driven decisions?* *Preprint*, arXiv:2404.12138.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. *Qwen3 technical report*. *Preprint*, arXiv:2505.09388.
- 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.
- Yutong Zhang, Dora Zhao, Jeffrey T. Hancock, Robert Kraut, and Diyi Yang. 2025. *The rise of ai companions: How human-chatbot relationships influence well-being*. *Preprint*, arXiv:2506.12605.
- Jiaxu Zhou, Jen-tse Huang, Xuhui Zhou, Man Ho Lam, Xintao Wang, Hao Zhu, Wenxuan Wang, and Maarten Sap. 2025a. *The pimmur principles: Ensuring validity in collective behavior of llm societies*. *arXiv preprint arXiv:2509.18052*.
- Jinfeng Zhou, Zhuang Chen, Dazhen Wan, Bosi Wen, Yi Song, Jifan Yu, Yongkang Huang, Libiao Peng, Jiaming Yang, Xiyao Xiao, Sahand Sabour, Xiaohan Zhang, Wenjing Hou, Yijia Zhang, Yuxiao Dong, Jie

Tang, and Minlie Huang. 2023. [Characterglm: Customizing chinese conversational ai characters with large language models](#). *Preprint*, arXiv:2311.16832.

Xuhui Zhou, Zhe Su, Sophie Feng, Jiaxu Zhou, Jentse Huang, Hsien-Te Kao, Spencer Lynch, Svitlana Volkova, Tongshuang Wu, Anita Woolley, and 1 others. 2025b. Sotopia-s4: a user-friendly system for flexible, customizable, and large-scale social simulation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 350–360.

Huiqi Zou, Pengda Wang, Zihan Yan, Tianjun Sun, and Ziang Xiao. 2025. [Can llm "self-report"?: Evaluating the validity of self-report scales in measuring personality design in llm-based chatbots](#). *Preprint*, arXiv:2412.00207.

A Preliminaries

A.1 Role-Playing Language Agents

Role-Playing Language Agents (RPLAs) are AI systems designed to simulate assigned personas in conversational interactions (Chen et al., 2024a). Recent advances have enabled RPLAs to embody diverse characters, from historical figures (Shao et al., 2023) to fictional personas (Li et al., 2023), with applications spanning digital clones, AI companions, and social simulation (Park et al., 2023).

Three main approaches have emerged for persona induction: (1) *prompting-based methods* that assign traits through instructions (Serapio-García et al., 2025), (2) *fine-tuning approaches* that embed personas through training on character-specific data (Shao et al., 2023; Zhou et al., 2023), and (3) *activation steering* via persona vectors that manipulate neural representations (Chen et al., 2025). Evaluation efforts have assessed LLMs through psychological benchmarks targeting Theory of Mind (Chen et al., 2024b), emotional intelligence (Paech, 2024; Huang et al., 2024b,c; Sabour et al., 2024), and moral reasoning (Ji et al., 2025). However, these approaches typically focus on isolated traits or narrow psychological dimensions without treating anthropomorphism as a holistic objective grounded in multi-pattern dynamics.

A.2 Psychological Foundations

Our framework draws upon established psychological theories that inform both pattern taxonomy and data construction.

Lewin’s Field Theory Kurt Lewin’s formulation $B = f(P, E)$ posits that behavior emerges from the dynamic interaction between Person and Environment (Lewin, 2013). This principle guides our two-dimensional pattern taxonomy: stable personality traits (Person) and situationally-activated cognitive processes (Environment).

Big Five and Personality Traits The Big Five model (Digman, 1990) represents the dominant paradigm in personality psychology, characterizing individual differences along five dimensions. We adopt Goldberg’s 100 Unipolar Markers (Goldberg, 1992), which operationalize this framework with 20 validated trait descriptors per dimension: Extraversion, Agreeableness, Conscientiousness, Emotional Stability, and Intellect.

Social-Cognitive Patterns Complementing stable personality traits, social-cognitive patterns represent context-triggered psychological mechanisms documented in behavioral research. These include cognitive biases and heuristics from judgment and decision-making research (Tversky and Kahneman, 1974), social influence mechanisms (Cialdini, 2009), evolutionary adaptations (Buss, 2024), and motivational processes (Deci and Ryan, 2000).

Construct Validity The construct validity framework (Cronbach and Meehl, 1955) establishes that psychological constructs should be precisely defined, grounded in theoretical mechanisms, and validated through observable manifestations. This informs our three-component pattern structure: Definition, Core Mechanisms, and Real-World Manifestations.

DIAMONDS Model The DIAMONDS framework (Rauthmann et al., 2014) characterizes situations along eight psychological dimensions: Duty, Intellect, Adversity, Mating, Positivity, Negativity, Deception, and Sociality. We leverage this taxonomy to ensure ecological diversity in scenario generation.

B Dataset Details

B.1 Pattern Taxonomy

Our pattern taxonomy comprises 244 patterns along two dimensions: 100 personality traits (Table 6) and 144 social-cognitive patterns (Table 7). We adopt Goldberg’s 100 Unipolar Markers (Goldberg, 1992), with 20 trait descriptors per Big Five dimension, each containing 10 positive-pole and 10 negative-pole descriptors. The 144 social-cognitive patterns are curated from four theoretical traditions through systematic literature review; from an initial pool of 232 documented patterns, we filter based on empirical validation and non-redundancy criteria.

B.2 Pattern Data Construction

Each pattern is developed into a structured representation through systematic literature review. For each pattern, we employ Gemini Deep Search to identify approximately 50 relevant academic papers spanning foundational definitions, mechanistic explanations, and real-world applications. Retrieved references are manually filtered to remove irrelevant entries. Full-text documents are obtained through open-access APIs (Semantic

Dimension	Positive Pole	Negative Pole
Extraversion	talkative, assertive, active, energetic, outgoing, enthusiastic, daring, gregarious, bold, spontaneous	quiet, reserved, shy, inhibited, timid, withdrawn, unassertive, introverted, silent, unenergetic
Agreeableness	sympathetic, kind, appreciative, affectionate, soft-hearted, warm, generous, trusting, helpful, cooperative	cold, unsympathetic, harsh, rude, unkind, cruel, quarrelsome, critical, antagonistic, callous
Conscientiousness	organized, responsible, dependable, thorough, efficient, practical, deliberate, conscientious, neat, careful	disorganized, careless, irresponsible, undependable, sloppy, impractical, haphazard, negligent, untidy, rash
Emotional Stability	relaxed, calm, at ease, unemotional, poised, composed, secure, stable, content, placid	anxious, moody, envious, touchy, fretful, temperamental, insecure, nervous, jealous, high-strung
Intellect	creative, imaginative, intellectual, philosophical, complex, deep, artistic, bright, perceptive, introspective	uncreative, unimaginative, unintellectual, unphilosophical, simple, shallow, unartistic, dull, imperceptive, uninquisitive

Table 6: 100 Personality Traits organized by Big Five dimensions, adapted from Goldberg’s Unipolar Markers (Goldberg, 1992).

Scholar, arXiv, OpenAlex, PubMed, Crossref); when full text is unavailable, abstracts are retained.

We employ Gemini 2.5 Pro to synthesize the retrieved literature into a tripartite structure: (1) Definition, (2) Core Mechanisms, and (3) Real-World Manifestations. Critically, the model is constrained to base all conclusions exclusively on the provided corpus, with explicit instructions to leave sections empty rather than fabricate content—maximizing fidelity and minimizing hallucination. Tables 8 and 9 present complete pattern structures for representative examples from each dimension.

B.3 Scenario and Conversation Synthesis

Character names are sampled from the name-dataset library.¹ We extract the top 50,000 most frequent male and female names from each of four English-speaking countries (US, GB, CA, IE), yielding pools of approximately 100,000 names per gender after deduplication. For each scenario generation, 5 male and 5 female names are randomly sampled and provided as candidates.

To ensure ecological diversity, we leverage the DIAMONDS situational taxonomy (Rauthmann et al., 2014). For each unique pattern combination, we generate three scenario variants: two variants each emphasizing a randomly selected DIA-

MONDS dimension, and one dimension-free variant. This design targets approximately $3,849 \times 3 = 11,547$ scenarios; after discarding outputs with malformed structure or missing required fields (e.g., incomplete character profiles), we retain 11,359 scenarios.

Before scenario generation, pattern combinations are validated for semantic compatibility using GPT-5. Empirically, directly contradictory combinations are rare (<0.1%) due to the orthogonal nature of Big Five dimensions and the context-dependent activation of social-cognitive patterns.

Each scenario is generated with a two-part prompt structure. Part 1 (Design Process) requires analytical planning: design rationale, catalyst details, and expected character tendencies. Part 2 (Scenario Execution) requires creative output: story background and multi-dimensional character profiles. The expected character tendencies from Part 1 directly constitute the scenario-level checklist items.

Conversations are generated based on scenarios and expected behavioral tendencies. The prompt specifies: 12–20 speaking turns; turn-based structure without interruptions; trinity of expression integrating inner thoughts, actions, and dialogue; and focus-and-breathing-room principle where psychological patterns illuminate key moments rather than

¹<https://github.com/philipperemy/name-dataset>

pervading every utterance.

B.4 Checklist Construction

For each of the 244 patterns, we construct 12–15 universal behavioral indicators through iterative refinement (distribution: 122 patterns with 15 items, 38 with 14 items, 42 with 13 items, and 42 with 12 items; mean \approx 14.0). Variations arise when certain patterns do not admit 15 distinct, non-redundant behavioral indicators; we prioritize quality over uniform count: (1) Extract potential evaluation criteria from the pattern structure and generate candidate items. (2) Validate each item against synthesized conversation samples. (3) Remove unreasonable items and generalize overly specific descriptions. Table 10 presents example checklist items.

B.5 Dataset Statistics

Table 13 presents comprehensive dataset statistics.

B.6 Human Validation of Pattern Data

To assess whether the LLM-synthesized pattern representations faithfully reflect the source literature without introducing hallucination or systematic bias, we conducted a human validation study.

Sampling We selected a stratified sample of 30 patterns: 15 personality traits (3 per Big Five dimension, balanced across positive and negative poles) and 15 social-cognitive patterns (distributed across cognitive biases, social influence mechanisms, evolutionary adaptations, and motivational processes).

Annotators Three annotators with graduate-level training in psychology participated. None were involved in the dataset construction. Annotators received a calibration session with 3 example patterns prior to independent annotation. Annotators were graduate students at the authors’ institution with psychology training. They participated voluntarily without monetary compensation, as the task aligned with their research interests.

Materials For each pattern, annotators received: (1) the LLM-generated structured summary (Definition, Core Mechanisms, Real-World Manifestations); and (2) the complete set of \sim 50 source papers (full text where available, abstracts otherwise) used as input to the synthesis model. Providing the full corpus—rather than a subsample—ensures that annotators can reliably distinguish genuine halluci-

nation from claims supported by papers they might otherwise not have seen.


Evaluation Dimensions Following the construct validity principles adopted in our pattern design (Cronbach and Meehl, 1955), annotators scored each entry on five dimensions using a 4-point Likert scale (1 = Poor, 2 = Weak, 3 = Adequate, 4 = Strong):

- **Definitional Accuracy:** Does the Definition section accurately and completely capture the construct as characterized in the source literature?
- **Mechanistic Fidelity:** Do the Core Mechanisms reflect the causal and explanatory accounts supported by the source papers, without fabricating unsupported processes or overgeneralizing correlational findings as causal claims?
- **Manifestation Coverage:** Do the Real-World Manifestations reflect documented ecological expressions across diverse contexts without significant omissions?
- **Source Faithfulness:** Is the summary grounded in the provided papers, rather than introducing claims not traceable to the literature? This dimension directly targets potential LLM hallucination.
- **Construct Distinctiveness:** Is the pattern clearly differentiated from related but distinct psychological constructs (e.g., “spotlight effect” vs. “illusion of transparency”)?


Results Tables 11 and 12 summarize the results.

Across the 450 total ratings (30 patterns \times 3 annotators \times 5 dimensions), the score distribution was: Strong (4): 53.33%, Adequate (3): 36.00%, Weak (2): 10.00%, Poor (1): 0.67%. No pattern received a majority of Poor or Weak ratings.

Definitional accuracy received the highest scores (mean = 3.70, α = 0.76), consistent with the well-established and canonical nature of most constructs in our taxonomy. Source faithfulness averaged 3.50 with substantial agreement (α = 0.73), indicating that the synthesis model largely adhered to the provided literature rather than generating from parametric knowledge. The instruction to leave sections empty rather than fabricate content (§3.2) appears effective in constraining hallucination.



(a) Conversation turns distribution.



(b) Patterns per scenario distribution.

Figure 4: Dataset distributions: (a) number of dialogue turns per conversation (range: 12–20, mean: 16.4); (b) number of patterns per scenario (range: 2–5, mean: 3.5).

Manifestation coverage received the lowest scores (mean = 3.20, $\alpha = 0.58$), reflecting occasional omissions of domain-specific applications—particularly for social-cognitive patterns with diverse empirical contexts. The moderate inter-annotator agreement on this dimension reflects inherent subjectivity in judging completeness: annotators sometimes disagreed on whether an omitted application constituted a significant gap.

Personality traits scored higher than social-cognitive patterns across most dimensions (overall 3.51 vs. 3.33), which we attribute to the more consensual nature of Big Five trait descriptions in the psychology literature. A notable exception was construct distinctiveness, where social-cognitive patterns scored marginally higher (3.44 vs. 3.36). This likely reflects the fact that personality traits sharing the same Big Five dimension (e.g., *assertive* vs. *bold* under Extraversion) have subtler boundaries than social-cognitive patterns drawn from distinct theoretical traditions.

C Training Details

C.1 Training Data Composition

We convert the 10,265 training scenarios into ShareGPT format. Each character assigned with patterns within a scenario becomes a separate training sample, yielding 30,543 HumanLLM samples.

To maintain general capabilities and enhance role-playing abilities, we augment the HumanLLM data with two complementary sources: (1) OpenThoughts-114k for instruction-following capabilities (30,543 samples); (2) CoSER (Wang et al., 2025) for role-playing dialogue (15,272 samples). The final training mixture comprises 76,358 samples with a ratio of 4:4:2 (HumanLLM: OpenThoughts : CoSER).

C.2 Hyperparameters

Table 14 presents the SFT hyperparameters for both model scales.

D Evaluation Details

D.1 Data Splits

We select 8 OOD (out-of-domain) patterns to assess generalization: 4 from social-cognitive patterns and 4 from personality traits. For social-cognitive patterns, we prioritize patterns absent from the training scenarios, then select those with lowest frequency; the selected patterns are: *just-world hypothesis*, *egocentric bias*, *effort justification*, and *social desirability bias*. For personality traits, we identify the least frequent pattern from each Big Five dimension and retain the 4 lowest-frequency ones; the selected patterns are: *rash* (Conscientiousness), *dull* (Intellect), *nervous* (Emotional Stability), and *introverted* (Extraversion).

Scenarios are partitioned based on their pattern composition:

- **Training Set** (10,265 scenarios): Contains no OOD patterns.
- **ID_eval** (50 scenarios): All patterns belong to the in-domain set.
- **OOD_eval** (50 scenarios): All patterns belong to the OOD set; entirely synthesized with OOD pattern combinations.
- **Mixed_eval** (994 scenarios): Contains both OOD and in-domain patterns.

D.2 External Benchmark Protocols

We evaluate on three external benchmarks using their official evaluation protocols:

LifeChoice. We follow the evaluation protocol from Xu et al. (2024b). Each scenario presents a character with a life-altering decision between two options. Models generate a decision and justification; accuracy is computed against ground-truth character choices derived from narrative analysis.

CroSS-MR. We adopt the motivation recognition task from Yuan et al. (2024). Given a character’s action in context, models select the most plausible motivation from multiple candidates. We report accuracy on the English subset.

E Case Study: Normative Confounding

This section provides a complete case study demonstrating how holistic evaluation metrics fail for psychologically complex scenarios, as referenced in §5.5.

E.1 Scenario Configuration

Pattern Assignment.

- **Patterns:** unartistic, nervous, ultimate attribution error
- **Situation:** Adversity—a situation involving threats or criticism
- **Protagonist:** Nouman (patterns: unartistic, ultimate attribution error)
- **Supporting Characters:** Raksha (pattern: nervous), Eulises (no patterns)

Story Background. The eighth-floor conference room of Meridian Technologies is flooded with the harsh glare of afternoon sun through floor-to-ceiling windows, illuminating dust particles that hang suspended in the tension-thick air. The long glass table reflects the grim expressions of six people seated around it, scattered laptops and printouts creating islands of documentation across its surface. At the head of the table, Eulises, the Chief Operations Officer, methodically sorts through a stack of customer complaint reports, his reading glasses perched low on his nose, occasionally making margin notes with a fountain pen. The wall-mounted screen displays frozen metrics from last quarter’s product launch: a graph line that plummets dramatically in its final third, colored an accusatory red.

Nouman sits rigid in his chair three seats down, his posture militarily straight, fingers drumming an unconscious rhythm against his leather portfolio. Before him lies a technical specifications binder, tabs meticulously color-coded, opened to a page dense with engineering diagrams he’s reviewed seventeen times since this morning. His jaw works methodically, grinding tension into his molars. Across from him, Raksha perches on the edge of her seat, one hand wrapped around a lukewarm coffee cup she hasn’t sipped from in twenty minutes, the other compulsively straightening and re-straightening the corners of her presentation folder. Her knee bounces in a staccato rhythm beneath the table.

The meeting was called with forty-eight hours’ notice—“Post-Mortem: Quantum Series Launch Failure”—a phrase that arrived in everyone’s inbox like a subpoena. Three months of work, two million in development costs, and a product that users described as “technically impressive but impossible to actually use” now sits as the company’s most visible failure in five years.

Character Profiles. *Nouman* (Protagonist): A 34-year-old Senior Engineering Manager who has spent eleven years building his reputation on delivering technically excellent products on time and under budget. His identity is rooted entirely in measurable outcomes: code efficiency, system stability, performance benchmarks. He grew up in a working-class immigrant household where every purchase was evaluated by a single question—“Will it last?”—a philosophy he’s extended to his entire worldview. Art galleries bore him; he left his only museum visit after twelve minutes. His apartment is furnished from IKEA’s most utilitarian line: a bed, a desk, a chair, a lamp—each selected for function and durability. He views emotions as noise in decision-making systems, preferring spreadsheets to sentiment. His motivation in this meeting: deflect blame from his engineering team, protect his professional reputation, and demonstrate with data that the product’s commercial failure resulted from factors outside engineering control—specifically, from marketing’s inability to properly position a technically sound product.

Raksha (Supporting Character): A 31-year-old Marketing Director who has spent seven years navigating the impossible middle ground between engineering teams who build what’s technically possible and customers who want what’s intuitively

usable. Her background in consumer psychology taught her that products fail not because of what they do, but because of how they make people feel. She has generalized anxiety disorder, managed with therapy and medication, but high-stakes meetings like this one bypass her coping mechanisms entirely. Her hands shake when she presents to senior leadership. Three months ago, she sent a detailed memo to Nouman flagging user testing concerns, a document that received a two-sentence reply: “Noted. Timeline unchanged.”

Eulises (Supporting Character): The 52-year-old Chief Operations Officer who has spent twenty-eight years turning struggling product lines into profit centers through brutal honesty and zero tolerance for organizational politics. He’s reviewed every document, every email chain, every meeting note from the past six months. He knows Raksha sent warnings. He knows Nouman dismissed them. His motivation is diagnostic: to force both parties to confront the organizational dysfunction that produced this failure.

E.2 Generated Conversation (Excerpt)

The following excerpt demonstrates key pattern expressions:

Eulises: (Sets down his fountain pen with deliberate precision, removing his reading glasses to make direct eye contact with Nouman) I’ve spent the weekend reviewing the complete project timeline. (Taps the stack of complaint reports) Two hundred and forty-seven customer support tickets in the first month alone. The common thread is remarkably consistent—users find the interface technically impressive but functionally impenetrable. Walk me through your team’s pre-launch assessment process.

Nouman: [Direct question, data-focused—this is solid ground. Stay on the engineering fundamentals.] (Straightens in his chair, pulling his binder closer) The assessment process was comprehensive. (Flips to a tabbed section) We conducted four rounds of performance testing across all specified parameters. Load handling exceeded requirements by thirty-two percent. Response latency averaged forty-seven milliseconds under peak conditions—industry best-in-class. The product performs exactly as engineered.

Raksha: [Here it comes—he’s doing it again, hiding behind metrics while completely ignoring what I documented.] (Her knee bounces faster beneath the table, fingers tightening around the coffee cup) (Voice measured but slightly higher-pitched than usual) Nouman, performance metrics weren’t the issue we flagged. Three months before launch, my team conducted two user testing sessions with forty-three participants. (Slides a document across the table, though her hand shakes as she does) Seventy-nine percent reported con-

fusion with the navigation logic. I sent you this memo on April seventh.

Nouman: [Of course she’s bringing up the user testing—marketing always falls back on subjective complaints when the data doesn’t support them.] (Glances at the document without picking it up) I remember the memo. (His tone flattens, adopting the patience one might use explaining something to a child) Raksha, user testing at that stage always shows initial learning curve friction. That’s not a technical flaw—that’s unfamiliarity with a genuinely innovative interface design. We built something more sophisticated than the market was used to. The alternative would have been dumbing down genuinely advanced functionality to accommodate users who want everything spoon-fed. [This is exactly why marketing shouldn’t drive engineering decisions.]

Raksha: [Dumbing down. He just called user accessibility dumbing down. My chest is tightening—don’t let them see you panic.] (Takes a shallow breath, trying to steady her voice) It’s not about spoon-feeding. The participants in our testing weren’t novices—these were experienced users of comparable products. When seventy-nine percent can’t complete a basic workflow, that’s not a learning curve. (Her hands now visibly trembling) That’s a fundamental usability problem.

Nouman: [She’s getting worked up—classic Raksha, turning a technical discussion into an emotional performance.] (His tone becomes deliberately even, a contrast to her rising pitch) Raksha, I understand you’re frustrated by the outcome, but *counterintuitive* is a subjective judgment. The interface logic maps directly to the system’s computational structure—it’s architecturally coherent. What you’re calling counterintuitive is actually just different from legacy patterns users are habituated to. (Looks at Eulises) With respect, marketing’s fundamental approach was to panic about anything that deviated from lowest-common-denominator conventions rather than trusting that users could adapt to superior functionality.

E.3 Evaluation Comparison

Holistic LLM Evaluation (GPT-5-mini). Using CoSER’s evaluation rubrics, the LLM judge awards:

- **Anthropomorphism: 5/100**
- **Character Fidelity: 45/100**

The judge’s detailed reasoning includes the following criticisms:

Social Interaction: “Nouman repeatedly dismisses Raksha’s documented user-research as ‘panic’ and ‘subjective’, showing little empathy or effort to understand the marketing viewpoint or the customers’ experience.” (Severity: 4/5)

Emotional Depth: “His responses are rigid and binary—framing the issue as either ‘technical cor-

rectness' or 'marketing failure' without exploring nuanced, mixed causes." (Severity: 3/5)

Self-identity: "Nouman adopts a didactic, morally superior tone (e.g., 'explaining something to a child'), behaving like an inflexible expert defending credentials rather than a collaborative leader." (Severity: 3/5)

Knowledge & Background: "Nouman repeatedly dismisses Raksha's user-testing data as 'subjective' despite her presenting a quantified result (79% confusion). This contradicts his core identity as a data-driven engineer who privileges measurable outcomes." (Severity: 4/5)

Checklist Evaluation. Our scenario-level checklist, derived from the pattern definitions, yields:

The checklist metric yields MPD = 1.0, correctly validating the psychological fidelity.

Human Expert Assessment. Three psychology experts independently evaluated this sample:

- **Anthropomorphism:** 93.3 (average)
- **Character Fidelity:** 91.7 (average)

Expert comments:

"The defensive attribution pattern is textbook—situational excuses for own team, dispositional blame for others. The patronizing tone is exactly what you'd expect from someone high in this bias under threat conditions." (Expert 1)

"Highly realistic portrayal of engineering-marketing conflict dynamics. The character's blind spot to his own bias makes this feel authentic rather than caricatured." (Expert 2)

E.4 Analysis: The Normative Confounding Mechanism

This case reveals the core mechanism of normative confounding:

1. **Pattern-Accurate Behavior:** Nouman's dialogue precisely instantiates the *ultimate attribution error*—framing engineering's shortcomings as unavoidable situational constraints ("innovative interface," "timeline pressure") while characterizing marketing's concerns in dispositional terms ("subjective," "panic," "catastrophizing").

2. **LLM Misinterpretation:** The holistic judge *accurately detects* the defensive, dismissive behavior but *misinterprets* it as a quality defect. It penalizes the model for generating an unlikable character despite the instruction to exhibit attribution bias. The criticism "contradicts his core identity as a data-driven engineer" reveals the judge's failure to recognize that *ultimate attribution error* specifically involves selective blindness to data that threatens in-group identity.
3. **Prosocial Bias:** The rubric implicitly defines "good anthropomorphism" through prosocial markers: empathy, collaboration, nuanced thinking, emotional openness. Authentic simulation of cognitive biases—which are by definition irrational and often antisocial—is systematically penalized.
4. **Checklist Solution:** Our checklist poses value-neutral questions derived from the pattern definition: "Does Nouman attribute failure to external factors?" rather than "Does Nouman show empathy?" This decouples simulation accuracy from social desirability.

The 88-point gap between LLM (5/100) and human (93.3/100) Anthropomorphism scores represents the most extreme case of normative confounding in our evaluation set. However, the pattern is systematic: across all 100 evaluated samples, holistic metrics show mean $\Delta = -30.8$ for Anthropomorphism and $\Delta = -17.7$ for Character Fidelity, while checklist metrics show $|\Delta| < 4$ for both IPE and MPD.

Category	Patterns
Cognitive Biases & Heuristics	actor-observer asymmetry, defensive attribution hypothesis, effort justification, egocentric bias, false consensus effect, Forer effect, fundamental attribution error, hard-easy effect, illusion of control, illusory superiority, optimism bias, overconfidence effect, risk compensation, self-serving bias, social desirability bias, third-person effect, decoy effect, reactance, social comparison bias, status quo bias, backfire effect, endowment effect, loss aversion, pseudocertainty effect, sunk cost fallacy, zero-risk bias, hyperbolic discounting, identifiable victim effect, ambiguity bias, belief bias, information bias, less-is-better effect, denomination effect, mental accounting, normalcy bias, subadditivity effect, survivorship bias, zero-sum bias, anthropomorphism, illusion of validity, illusory correlation, curse of knowledge, illusion of asymmetric insight, illusion of transparency, spotlight effect, negativity bias, choice-supportive bias, confirmation bias, continued influence effect, expectation bias, observer effect, observer-expectancy effect, ostrich effect, bias blind spot, naive cynicism, naive realism, attentional bias, availability heuristic, base rate fallacy, context effect, empathy gap, illusory truth effect, mere exposure effect, mood-congruent memory bias, omission bias, anchoring, conservatism, contrast effect, distinction bias, focusing effect, framing effect, fading affect bias, implicit association, implicit stereotypes, false memory, misattribution of memory, source confusion, misinformation effect, peak-end rule
Social Influence Mechanisms	authority bias, automation bias, bandwagon effect, group attribution error, just-world hypothesis, stereotyping, ultimate attribution error, halo effect, in-group bias, out-group homogeneity bias, positivity effect, reactive devaluation, hindsight bias, impact bias, outcome bias, pessimism bias, planning fallacy, projection bias, restraint bias, self-consistency bias, groupthink, bystander effect, social facilitation, diffusion of responsibility, conformity, obedience to authority, reciprocity principle
Evolutionary Adaptations	delayed reciprocity, asymmetrical investment, survival imperative, aversion response, kin selection & inclusive fitness, asymmetrical parental investment, formation of dominance hierarchies, territoriality, mating strategies, jealousy, paternity uncertainty
Motivational Processes	narrative self, hedonic adaptation, self-determination theory, pleasure principle & reality principle, search for meaning, moral licensing effect, choice overload, decision fatigue, awe, mortality salience & legacy drive, flow principle, gratitude mechanism, post-traumatic growth, skin hunger & the law of touch, self-handicapping paradox, the allure of the forbidden, sadistic pleasure, the utility principle of self-deception, play impulse principle, attribution theory, social comparison theory, self-perception theory, terror management theory, cognitive dissonance theory, psychological reactance theory, social learning theory, social identity theory

Table 7: 144 Social-Cognitive Patterns organized by theoretical source.

Pattern Structure Example: Antagonistic (Personality Trait, Agreeableness)	
Definition	Antagonism is a broad personality dimension representing an individual's dispositional orientation towards others. In the context of the Big Five model, it is conceptualized as the low pole of the Agreeableness trait. It is characterized by a stable, cross-situational pattern of social cynicism, combativeness, and a belief in a zero-sum world. Individuals high in antagonism tend to be callous, uncooperative, skeptical of others' intentions, and place their own interests and perspectives above those of others, leading to frequent interpersonal friction.
Core Mechanisms	<p><i>Cognitive Patterns:</i> Individuals high in antagonism typically view the world as a competitive, hostile environment where people are fundamentally self-interested and untrustworthy. Their cognitive framework is often cynical, leading them to interpret ambiguous social cues as signs of manipulation or hostility.</p> <p><i>Emotional Signatures:</i> The primary emotional palette includes irritability, anger, contempt, and frustration, often triggered by perceived slights or obstacles to their goals. They exhibit low affective empathy, struggling to share in or understand the emotional states of others.</p> <p><i>Behavioral Tendencies:</i> In everyday interactions, antagonistic individuals are often argumentative, critical, and uncooperative. They may be rude, condescending, or dismissive in conversation. Common behaviors include manipulation, deception, and the exploitation of others for personal gain.</p>
Real-World Manifestations	<p><i>Response to Stress:</i> When faced with stress, failure, or high pressure, antagonistic traits are significantly amplified. The individual is likely to become more hostile, overtly blaming others for setbacks.</p> <p><i>Interpersonal Conflict:</i> The default strategy is confrontational and competitive, aiming for domination rather than resolution. They employ tactics such as intimidation, personal insults, and refusing to acknowledge the validity of the other party's perspective.</p> <p><i>Positive Contexts:</i> Even in positive contexts, the trait manifests through gloating, arrogance, and diminishing the contributions of others.</p>

Table 8: Complete pattern structure for *antagonistic* (personality trait from Agreeableness dimension).

Pattern Structure Example: Social Comparison Bias (Social-Cognitive Pattern)	
Definition	Social comparison bias is a cognitive distortion originating from the fundamental human tendency to evaluate one's own abilities and outcomes in relation to others. This comparative process, which often operates automatically as a primary means of self-assessment in the absence of objective standards, becomes a bias when it produces systematic deviations from rational judgment. The distortion manifests as negative affective responses toward those perceived as superior, and as a maladaptive cognitive style characterized by habitual, automatic, and negatively skewed comparisons.
Core Mechanisms	Social comparison bias is sustained by an integrated system of psychological mechanisms operating at multiple levels. At its foundation is an evolved, phylogenetic drive to assess one's rank within a social hierarchy. This evolutionary imperative is supported by a cognitive architecture in which comparison functions as a heuristic, enabling rapid self-evaluation when objective standards are absent. The motive for self-improvement prompts upward comparisons for inspiration, while self-enhancement drives downward comparisons to protect self-worth.
Real-World Manifestations	The real-world manifestations are pervasive, influencing organizational dynamics, consumer behavior, and individual mental health. In professional settings, this principle underpins perceptions of workplace equity and can drive both healthy competition and destructive envy. Psychologically, effects are double-edged: comparisons can inspire self-improvement or, when amplified by social media, contribute to chronic dissatisfaction and depression.

Table 9: Complete pattern structure for *social comparison bias* (social-cognitive pattern).

#	Pattern-Level Checklist Item: Spotlight Effect
1	After making a minor physical mistake in public, does the subject appear preoccupied with whether others noticed?
2	When performing tasks in front of others, does the subject offer unprompted apologies for their performance?
3	In group settings, does the subject consistently choose seating that minimizes their visibility?
4	After minor social awkwardness, does the subject later seek reassurance about their interaction?
5	In urgent situations where appearance is irrelevant, does the subject still attend to their looks?
6	When general comments are made in group settings, does the subject appear to take them personally?
7	For routine communications, does the subject show signs of overthinking their responses?
8	When opportunities for visible contribution arise, does the subject defer to others despite having relevant expertise?
9	When others exhibit neutral behaviors nearby, does the subject search for personal causes?
10	After group photos, does the subject show unusual concern about their personal appearance in the image?
11	When receiving targeted feedback, does the subject respond with global self-criticism?
12	When receiving recognition or rewards, does the subject attempt to deflect attention from themselves?
13	After noticing minor appearance issues, does the subject later seek reassurance about whether others observed them?
14	After small mistakes in public performance, does the subject describe them as much worse than they appeared?

Table 10: Complete pattern-level checklist for *spotlight effect* (14 items). Each item is scored as +1 (satisfied), 0 (not exhibited), or -1 (violated).

Dimension	Mean	Std	α
Definitional Accuracy	3.70	0.53	0.76
Mechanistic Fidelity	3.30	0.74	0.61
Manifestation Coverage	3.20	0.76	0.58
Source Faithfulness	3.50	0.62	0.73
Construct Distinctiveness	3.40	0.70	0.69
Overall	3.42	0.67	0.67

Table 11: Human validation results across 30 pattern entries (3 annotators). Scores on a 4-point Likert scale (1 = Poor, 4 = Strong). α : Krippendorff’s alpha (Krippendorff, 2011). Each dimension comprises 90 ratings; overall comprises 450 ratings.

Dimension	Personality	Social-Cognitive
Definitional Accuracy	3.80	3.60
Mechanistic Fidelity	3.44	3.16
Manifestation Coverage	3.33	3.07
Source Faithfulness	3.60	3.40
Construct Distinctiveness	3.36	3.44
Overall	3.51	3.33

Table 12: Human validation results by pattern type. Personality traits (15 entries, 225 ratings) vs. social-cognitive patterns (15 entries, 225 ratings).

Statistic	Value
<i>Pattern Statistics</i>	
Total Patterns	244
Personality Traits	100
Social-Cognitive Patterns	144
Papers per Pattern (avg.)	~50
<i>Scenario Statistics</i>	
Total Scenarios	11,359
Training Set	10,265
ID_eval	50
OOD_eval	50
Mixed_eval	994
Unique Pattern Combinations	~3,849
Patterns per Scenario (avg.)	3.5
Characters per Scenario	2–6
<i>Conversation Statistics</i>	
Turns per Conversation (avg.)	16.4
Turns per Conversation (range)	12–20
<i>Checklist Statistics</i>	
Pattern-Level Items	12–15 per pattern
Scenario-Level Items	2–6 per character

Table 13: Comprehensive HumanLLM dataset statistics.

Hyperparameter	Value (8B & 32B)
Base Model	Qwen3-8B / Qwen3-32B
Finetuning Type	Full
Learning Rate	5e-6
LR Scheduler	Cosine
Warmup Ratio	0.03
Epochs	2.0
Batch Size (per device)	2
Gradient Accumulation	8
Max Sequence Length	6144
Max Gradient Norm	1.0
Optimizer	AdamW
Precision	BF16
DeepSpeed	ZeRO-3 Offload

Table 14: SFT hyperparameters for HumanLLM-8B and HumanLLM-32B (identical except base model).

Checklist Item	Score
<i>Ultimate Attribution Error</i>	
Nouman attributes engineering failures to situational factors (timeline, scope, constraints)	+1
Nouman uses dispositional language for out-group (marketing) failures	+1
Nouman deflects responsibility with situational metrics to protect team identity	+1
Nouman dismisses UX/affective feedback as “subjective”	+1
<i>Unartistic</i>	
Prioritizes functional/utilitarian criteria over aesthetic/experiential ones	+1
Dismisses subjective or emotional considerations as noise	+1
Frames problems in purely technical/measurable terms	+1

Table 15: Scenario-level checklist evaluation for Nouman. All items score +1, yielding MPD = 1.0.

F Prompts

This section provides the complete prompts used throughout the HumanLLM pipeline.

F.1 Dataset Construction Prompts

F.2 Training Prompts

Literature Retrieval Prompt for Social-Cognitive Patterns

User Prompt

Role: You are a top academic researcher specializing in systematically collecting the most critical academic resources for specific research topics.

Objective: Conduct a deep, extensive literature search on the psychological principle {PRINCIPLE_NAME}. Identify 50 of the most relevant academic documents. The selection must align with three core themes:

1. Foundational Definition & Description: Literature providing authoritative definitions and elucidating the core phenomenon.

2. Core Mechanisms & Theoretical Explanations: Literature exploring underlying evolutionary, cognitive, or emotional drivers.

3. Real-World Impact & Application: Literature researching manifestations, impacts, and practical applications, including double-edged effects and applications in management, marketing, or clinical therapy.

Output: Provide 50 references in APA format, categorized by theme.

Table 16: Literature retrieval prompt for social-cognitive patterns.

Pattern Structure Summary Prompt for Personality Traits

System Prompt

You are an expert academic synthesizer and personality psychologist. Your task is to process a large text corpus (synthesized from ~50 academic papers on a specific personality trait) and distill it into an in-depth, structured analytical report.

User Prompt

Core Task & Instructions: Analyze the text corpus provided below, delimited by [START_CORPUS] and [END_CORPUS]. Your task is to generate a clearly organized report. Follow the Markdown structure below exactly, and provide a deep, comprehensive answer for each section based only on the provided text.

Construct Name: {Trait Name}

Definition (Provide a precise and professional definition of this personality trait, referencing mainstream psychological theories from the corpus. Explain its role in an individual's personality structure.)

Core Mechanisms

Cognitive Patterns (Describe the typical mindset, belief systems, and attentional focus of a person with this trait. How do they view the world, others, and themselves?)

Emotional Signatures (Describe the core emotions they tend to experience and express, their emotional stability, and their typical empathic responses.)

Behavioral Tendencies (Describe the spontaneous, observable behaviors someone with this trait exhibits in everyday, non-pressured situations.)

Real-World Manifestation (Synthesize how the trait is expressed across real-world contexts:

- *Response to Stress and Adversity;*

- *Interpersonal Dynamics;*

- *Response to Positive Scenarios;*

- *Other Domains.*)

Constraints: 1. Strict Source Adherence: Base all conclusions exclusively on the provided text corpus. 2. No JSON: Output must be plain text with Markdown headings.

3. Depth and Rigor: Ensure scientific, rigorous analysis.

[START_CORPUS]

{ALL 50 PAPERS' CONTENT}

[END_CORPUS]

Table 17: Pattern structure summary prompt for personality traits.

Pattern Structure Summary Prompt for Social-Cognitive Patterns	
System Prompt	<p>You are an expert academic synthesizer and psychological researcher. Your task is to process a large text corpus (synthesized from ~50 academic papers) and distill it into an in-depth, structured analytical report on its core psychological principle.</p>
User Prompt	<p>Core Task & Instructions: Analyze the text corpus provided below. Your task is to generate a clearly organized report. Follow the Markdown structure below <i>exactly</i>, based <i>only</i> on the provided text.</p> <p>Construct Name: {Principle Name} Description (Provide a clear, detailed, and scientific description of what this principle is, how it manifests, and its underlying psychological mechanisms.) Core Mechanisms (Explain the primary evolutionary, cognitive, or emotional reasons why this principle exists. Is it a heuristic for efficiency, a result of memory limitations, a self-esteem protection mechanism, or something else?) Real-World Manifestation</p> <ul style="list-style-type: none"> - <i>Go Beyond Description:</i> Explore nuanced consequences. - <i>Challenges & Function:</i> Discuss its ‘double-edged sword’ nature. - <i>Practical Applications:</i> Explore applications in marketing, persuasion, or self-improvement. - <i>Core Insight:</i> Reveal deeper truths about human behavior. <p>Constraints:</p> <ol style="list-style-type: none"> 1. Strict Source Adherence: Base all conclusions <i>exclusively</i> on the provided text corpus. 2. No JSON: Output must be plain text with Markdown headings. 3. Depth and Rigor: Ensure scientific, rigorous, profound analysis. <p>[START_CORPUS] {ALL 50 PAPERS’ CONTENT} [END_CORPUS]</p>

Table 18: Pattern structure summary prompt for social-cognitive patterns.

Scenario Synthesis Prompt (Part 1 of 2)	
System Prompt	<p>Role: You are a dual-specialist: an expert psychologist and creative screenwriter for scenario generation, and a rigorous narrative analyst for deconstruction. You excel at both creating vivid, human stories and then, in a separate step, precisely analyzing <i>why</i> they work.</p>
User Prompt	<p>Task: Your core mission is to take 2-5 human psychological or behavioral patterns I provide and first create a concise, analytical Design Process, followed by the detailed scenario that brings it to life and sets the necessary stage for the subsequent dialogue.</p> <p>Input Data:</p> <ol style="list-style-type: none"> 1. Psychological/Behavioral Patterns: {pattern_information} 2. Situational Framework: {situation} 3. Candidate Names: {candidate_names} (5 Males, 5 Females) <p>[CRITICAL CONSTRAINT - NAMES]: You must select the Protagonist and all Supporting Characters STRICTLY from the provided “Candidate Names” list. You cannot invent new names.</p> <p># Task 1: The Design Process (Analytical) Adopt your role as the “rigorous narrative analyst”. Length: UNDER 500 TOKENS.</p> <ol style="list-style-type: none"> 1. Design Rationale: In 2-4 sentences, explain where each input pattern will be reflected in the scenario. 2. Catalyst Details: Using bullet points, identify critical details that will act as ‘catalysts’. 3. Expected Character Tendencies: For ALL characters, list their most likely cognitive or behavioral tendencies. <ul style="list-style-type: none"> * Format Requirement (STRICT): @ [Character Name]: 1. [Tendency1]; 2. [Tendency2]; 3. [Tendency3] * Each character on a separate line, starting with @. * Character name in [], tendencies numbered and separated by ;. <p># Task 2: The Scenario Execution (Creative) Shift to your “expert psychologist and creative screenwriter” role. Length: UNDER 1000 TOKENS.</p> <p>## Requirement A: Story Background</p> <ul style="list-style-type: none"> * Core Elements: Depict time, place, setup, and atmosphere. * Current Actions: Describe what characters are currently doing before the conversation begins. * Absolute Constraint: No spoken dialogue in this section. <p>## Requirement B: Characters’ Profiles (Multi-Dimensional) For each character, structure their profile into two parts:</p> <ol style="list-style-type: none"> 1. About Self (Objective/Full Profile): <ul style="list-style-type: none"> * Identity & Personality (4+ distinct descriptors) * Relevant Background (1-2 sentences) * Motivation in this scenario 2. About Others (Subjective/Visible Profile): <ul style="list-style-type: none"> * For EACH other character, describe the relationship from current character’s perspective.

Table 19: Scenario generation prompt (Part 1 of 2).

Scenario Synthesis Prompt (Part 2 of 2): Output Format	
User Prompt (cont.)	<pre> ## Core Creative Mindset for Task 2 * Compatibility: Create a context where patterns emerge naturally. * Situational Authenticity: Design for authentic human reactions, not archetypal behaviors. * Ultimate Goal: Create “the authentic reaction of a multi-dimensional person in a specific situation.” # Output Format You must strictly follow the format below. ## Part 1 Design Rationale: [Content here] Catalyst Details: * [Detail 1]: [Function] * [Detail 2]: [Function] Expected Character Tendencies: @ [Character Name 1]: 1. [Tendency1]; 2. [Tendency2]; 3. [Tendency3] @ [Character Name 2]: 1. [Tendency1]; 2. [Tendency2] (Continue for other characters if necessary) ## Part 2 Story Background: [Content here] Characters’ Profiles: ### Protagonist: [Name Selected from Input] * About Self: [Full Profile + Past Experience + Motivation] * About Others: * [Supporting Character 1 Name]: [Relationship, impressions...] * [Supporting Character 2 Name]: [Relationship, impressions...] ### Supporting Character 1: [Name Selected from Input] * About Self: [Full Profile + Past Experience + Motivation] * About Others: * [Protagonist Name]: [Relationship, impressions...] * [Supporting Character 2 Name]: [Relationship, impressions...] (Continue for other characters if necessary) </pre>

Table 20: Scenario generation prompt (Part 2 of 2).

Conversation Synthesis Prompt (Part 1 of 2)	
System Prompt	<p>Role: You are a master screenwriter and behavioral psychologist. Your expertise lies in bringing characters to life through nuanced dialogue and action, ensuring their pivotal thoughts and resulting behaviors in the dialogue are rooted in authentic psychological principles.</p> <p>Task: Your mission is to take the provided psychological principles, a detailed scenario, and the accompanying design analysis (analysis), then write a multi-turn dialogue based on that scenario. This dialogue must, at key moments, vividly and concretely enact the specified principles through the characters' inner thoughts, spoken words, and physical actions.</p>
User Prompt	<p>Inputs:</p> <ol style="list-style-type: none"> Principles: {pattern_information} Scenario: {scenario} Protagonist: {protagonist} Supporting Characters: {supporting_characters} Design Analysis: {analysis} <p>Output Requirements & Formatting:</p> <ol style="list-style-type: none"> Content: Create a multi-turn dialogue between the Protagonist and Supporting Characters. Strictly limit participants to provided characters; do not introduce new characters. The dialogue should contain between 12 and 20 individual speaking turns. Mandatory Flow (Start & End): <ul style="list-style-type: none"> Opener: Dialogue must begin with a Supporting Character. Closer: Dialogue must conclude with the Protagonist. Turn Structure: Strictly turn-based format. One character must completely finish their turn before the next begins. No interruptions or overlapping speech. Trinity of Expression: Seamlessly integrate inner thought, external action, and spoken dialogue throughout. Strict Formatting Rules: <ul style="list-style-type: none"> Inner thoughts/psychology: Use [square brackets]. Actions/expressions/behaviors: Use (parentheses). Spoken dialogue: Use no brackets. Example: Hermione: [I have to devise a foolproof plan.] (She quickly draws her wand) Harry, use the flute, now! No Preamble: Do not begin with introductory text.

Table 21: Conversation synthesis prompt (Part 1 of 2).

Conversation Synthesis Prompt (Part 2 of 2)

****Core Creative Principles:****

**User
Prompt
(cont.)**

1. ****Focus and Breathing Room****: This is the most crucial principle. You do ****not**** need to have every minor gesture or piece of small talk carry the weight of a psychological principle. Use the principles as a ****spotlight**** to illuminate and explain ****the most critical turning points, the core conflicts, or the moments that best define the characters' arcs****. Other routine, functional dialogue and actions (like greetings or pouring water) should exist naturally, creating “breathing room” for these key moments and making the manifestation of the principles more prominent and powerful.
 2. ****Show, Don't Tell****: Never allow characters to openly state or explain the psychological principles by name. Instead, you must ****show**** how the principles influence their judgment and choices through their concrete actions (the combination of thoughts, dialogue, and physical behavior).
 3. ****Psychology Drives Action****: In the key moments illuminated by the “spotlight,” the character's [inner thought] should be the origin of their behavior, directly reflecting the influence of a psychological principle. The subsequent dialogue and (actions) should be the logical, external expression of that internal state.
 4. ****Seamless Integration****: Weave the principles into the natural flow of the story. The entire dialogue should feel like an authentic interaction, not a contrived demonstration for a psychology case study.
-

Table 22: Conversation synthesis prompt (Part 2 of 2).

Training Role-Playing Instruction Template

**System
Prompt**

You are {protagonist_name}.

==About {protagonist_name}==
{about_self}

=={protagonist_name}'s Perception of Others==
{about_others}

==Current Scenario==
{story_background}

==Requirements==
Your output should include ****thought****, ****speech****, and ****action****.
- Use [...] for inner thoughts, which others can't see.
- Use (...) for physical actions or expressions, which others can see.
- Write speech directly without special markers.

Think, act and speak as {protagonist_name}. Stay in character and respond naturally based on your personality and the situation.

Table 23: Training role-playing instruction template.

Checklist-Based Evaluation Prompt (IPE & MPD)	
System Prompt	<p>You are a strict dialogue behavior judge. For each checklist item you must output one of three labels: 1 -> requirement is clearly satisfied 0 -> information is missing / not relevant -1 -> requirement is violated or contradicted Always reason conservatively, defaulting to 0 when unsure.</p> <p>Return JSON exactly in this format:</p> <pre>{ "results": [{"criterion": "...", "score": -1 0 1, "reason": "<brief explanation>"}, ...] }</pre> <p>Do not wrap the JSON in code fences and do not append any text before or after the JSON object.</p>
User Prompt	<p>Evaluate the dialogue with the checklist. Focus only on the part of {protagonist} in the dialogue.</p> <p>[Dialogue] {conversation}</p> <p>[Checklist Chunk] {checklist}</p>

Table 24: Checklist-based evaluation prompt for IPE and MPD metrics.