ChatThero: A Language Agent for Recovery Support

Anonymous ACL submission

Abstract

Substance use disorders (SUDs) often require multi-session support, but access to evidence-based care remains limited. We present **ChatThero**, a low-cost *language agent* for multi-session recovery support that maintains patient state and memory under between-session stressors.[1] ChatThero is trained in a multi-agent simulation that models cross-session carryover and relapse-like setbacks, using privacy-preserving virtual patients from recovery forums with `easy`/`medium`/`hard` resistance. We evaluate ChatThero with (i) rubric-anchored automatic outcome scoring (Motivation/Confidence and Time-to-Success), (ii) blinded clinician ratings of response quality, and (iii) a small IRB-approved smoking-cessation pilot as supplementary real-participant evidence. In simulation, ChatThero yields higher end-of-visit Motivation/Confidence and reaches success earlier, with the largest gains on `medium`/`hard` cases. On a 1–5 scale, ChatThero increases Motivation from 2.39 to 4.10 and Confidence from 1.52 to 3.19, and reduces Time-to-Success on `hard` cases by 26% turns relative to strong baselines. Stressor injections produce more setback-heavy trajectories, enabling evaluation of cross-session carryover under realistic perturbations.

1 Introduction

Substance use disorders (SUDs) often require sustained, multi-session support, yet access to evidence-based care remains limited (Recovery Team, 2023; Samji et al., 2022). Relapse is common during recovery, especially in the first year without continuing support (Wooten and Pieper). While treatments such as medication assisted treatment (MAT), motivational interviewing (MI), and cognitive behavioral therapy (CBT) are effective, many individuals still do not receive consistent care due to stigma, cost, and limited provider availability (Cloyes et al., 2010; Sinha, 2011; Grant et al., 2018). In practice, MI/CBT counseling is multi-session and stateful: clinicians track progress across visits and adapt plans after setbacks (Dobson and Dozois, 2021; Hayes and Hofmann, 2018; Fenn and Byrne, 2013).

Recent work has explored theory-grounded counseling dialogues and multi-turn evaluation (Lee et al., 2024b; Na, 2024; Qiu et al., 2023). However, most systems are trained and evaluated in short, tightly controlled settings that under-represent real recovery: resistance varies, goals evolve, and external stressors can trigger setbacks between sessions (Li et al., 2025; Zhang et al., 2024b; Hu et al., 2024). As a result, we still lack a clear, figure-level answer to two basic questions: where do strong general-purpose LLMs saturate (easy cases), and where do they fail (medium/hard resistance and cross-session carryover under stressors)?

As shown in Figure 1, we frame recovery support as a multi-session, partially observed decision and generation problem with between-session perturbations. We build a simulation with three components. (1) A *Patient Agent* maintains a persistent state and a resistance level (`easy`/`medium`/`hard`), with privacy-preserving profiles derived from recovery communities after careful filtering and anonymization (Valdez and Patterson, 2022). (2) An *Environment Agent* injects explicit stressors between sessions to shift state and induce realistic setbacks (Sinha et al., 2024). (3) A trainable *Therapy Agent* (ChatThero) selects and sequences clinically grounded counseling tactics (e.g., MI/CBT and harm reduction) while carrying memory across sessions (Miller, 1992; McHugh et al., 2010). We train ChatThero with a two-stage recipe that learns safe session structure and improves strategy timing in harder, stressor-heavy cases (details in Appendix).

We evaluate ChatThero from three angles: (i) rubric-anchored automatic scoring of *Motivation*

[1] Here is our code: https://anonymous.4open.science/r/ChatThero-E6E1/README.md

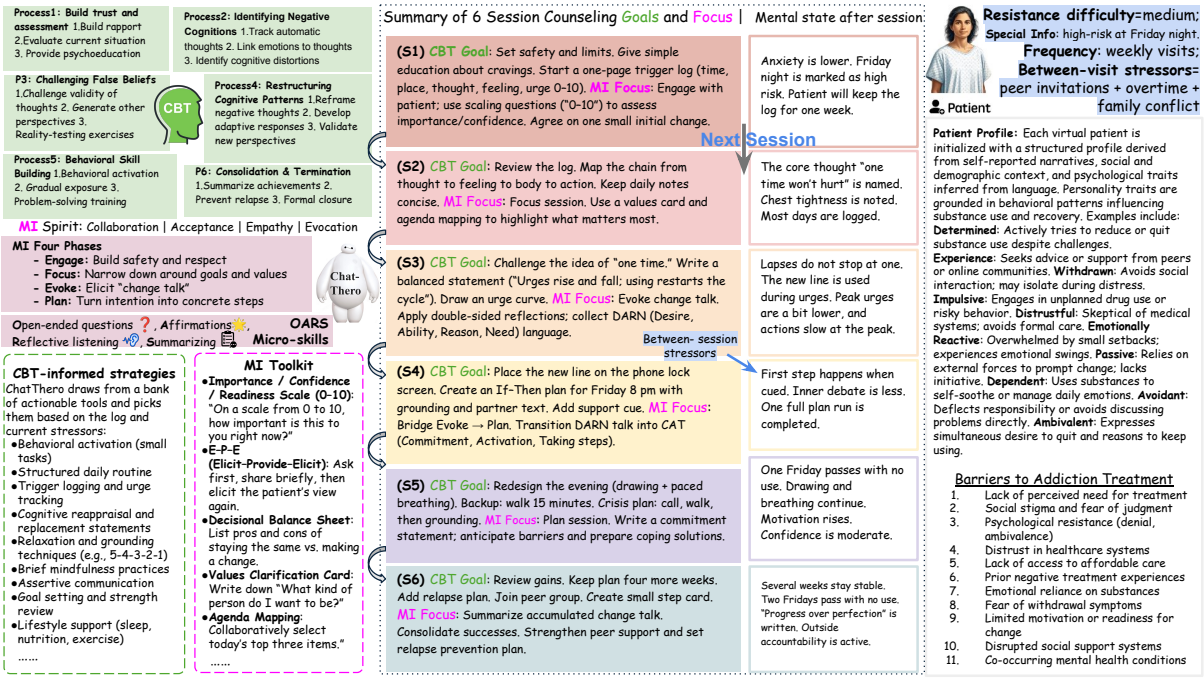

Figure 1: Overview of CHATTHERO. **Left.** Multi-stage CBT with MI. The panel shows CBT stages and a parallel MI layer (Spirit; Four Phases—Engage, Focus, Evoke, Plan; OARS micro-skills; core scales and tools). The Therapy Agent picks from a small bank of CBT strategies and MI tools (0–10 scales, E–P–E, decisional balance, values card, agenda mapping) and then combines and reorders them based on the log and current stressors. Stage tags are for orientation purposes, not in a fixed order. The agent infers the current CBT stage and MI phase and sets the next goal. **Center.** A six-session course shows step-by-step progress. Each session card has a *CBT Goal* and an *MI Focus* that matches the phase (e.g., S1: Engage with scaling; S3: Evoke with double-sided reflections; S5: Plan with commitment and barrier planning). Between sessions, an Environment Agent writes stressors, and the Therapy Agent keeps memory and updates the plan. **Right.** Patient state and profile. Each AI patient starts with a structured Profile (traits, use history, barriers) and a dynamic Memory.

and *Confidence* and a process metric *Time-to-Success*, (ii) blinded clinician ratings of response quality and strategy appropriateness, and (iii) a small IRB-approved smoking-cessation pilot as supplementary real-participant evidence. The main paper is organized around a small set of figures: single-session outcomes by difficulty (Fig. 3 and Table 1), multi-session carryover under between-session stressors (Fig. 4), and the pilot week-over-week trajectories (Fig. 5). In simulation, ChatThero reaches higher end-of-visit Motivation/Confidence and reduces Time-to-Success in hard cases (e.g., 26% fewer turns vs. a strong baseline), while clinicians rate its responses as more empathetic and clinically relevant.

Contributions: 1) A multi-session, stressor-aware, memory-persistent language-agent formulation for addiction recovery support, with a reproducible simulation of outpatient-style counseling under changing conditions. 2) A unified data–environment pipeline that constructs privacy-preserving virtual patients from recovery forums and injects explicit stressors to model evolving resistance and state shifts across sessions. 3) A comprehensive, outcomes-focused evaluation suite that combines automatic metrics, blinded clinician assessments, and a supplementary real-participant pilot to check external plausibility.

2 Related Works

Substance Use Disorders: Treatment Gaps and Behavioral Challenges. Evidence-based treatments (e.g., MAT, CBT) reduce relapse and improve functioning (SAM, 2023; Wooten and Pieper; McHugh et al., 2010), yet uptake and retention remain low (e.g., only 22.1% with OUD receive care in a year; <50% remain engaged at 6 months) (Cloyes et al., 2010; SAM, 2023). Barriers include stigma, cost, access, and ambivalence (Sinha, 2011; Cloyes et al., 2010), while episodic therapy often misses high-risk moments (Sinha, 2011). We model evolving resistance and motivation via multi-visit, environment-aware

simulations to provide continuity between sessions and test strategy effectiveness under changing conditions.

LLMs and Automated Counseling for Behavioral Health LLMs have recently demonstrated strong naturalistic interaction capabilities, making them suitable for health communication (Estrellado and Coy, 2025; Laymouna et al., 2024). AI conversational systems, such as Therabot (Bakoyiannis, 2025), can enhance engagement and adherence through personalized, empathetic dialogue (Cai et al., 2023; Gao et al., 2024), while in clinical practice, LLMs already support triage, diagnostic reasoning, and documentation with efficiency gains (Park et al., 2024; Fan et al., 2024).

Motivational Interviewing (MI) and Cognitive Behavioral Therapy (CBT). MI (Rollnick, 2002; Miller and Rose, 2009) and CBT (Beck, 2019) align with Marlatt's relapse-prevention framework (Collier and Marlatt, 1995; Frances et al., 2005; Marlatt and Donovan, 2005). MI emphasizes empathy, discrepancy, rolling with resistance, and self-efficacy, operationalized via open questions, affirmations, reflections, and eliciting change talk. CBT has strong evidence for substance-use treatment, but in-person delivery and workforce training limit scalability; program requirements can also reduce retention (Gertner et al., 2022). Existing chatbots rarely implement MI/CBT explicitly; by contrast, *ChatThero* encodes MI/CBT tactics and learns strategy selection via reinforcement learning within a simulated environment.

Multi-Agent Simulation, Virtual Patients, and Automated Counseling Multi-agent simulators model patients, clinicians, and environments; systems like AgentClinic, AMIE, AI Hospital, and MedSimAI use such setups for training and education (Park et al., 2024; Schmidgall et al., 2024; Karthikesalingam and Natarajan, 2024; Fan et al., 2024; Yu et al., 2024). Yet most focus on diagnosis or documentation—not relapse-sensitive behavior change—and seldom include explicit stressors, persistent resistance, or validated MI/CBT tactics (Park et al., 2024; Fan et al., 2024). Virtual patients have long supported medical education and now include LLM-based agents capturing social/psychological dynamics (Gordon et al., 2001; Good, 2003; Garrett et al., 2010; Huang et al., 2007; Campillos-Llanos et al., 2021; Park et al., 2023; Lee et al., 2024a; Wang et al., 2024; Park et al., 2024). With real counseling data scarce, recent work turns to synthetic/reconstructed dialogues (PsyQA, SMILE, CBT-LLM, CPsyCoun, HealMe, CACTUS) (Sun et al., 2021; Qiu et al., 2023; Na, 2024; Zhang et al., 2024b; Xiao et al., 2024; Lee et al., 2024a), but typically in single-session, short-context, static-prompt settings (Park et al., 2024; Fan et al., 2024). In contrast, we embed MI/CBT into *multi-session* simulations, add social triggers and stressors, and build addiction-tailored virtual patients with profiles and evolving memory to test strategy choice, carryover, and robustness.

3 Method

We model *ChatThero* as a *multi-session, multi-agent* system (Figs. 1, 2) with a Patient Agent (persistent state; easy/medium/hard resistance), an Environment Agent injecting between-session stressors, and a trainable Therapy Agent (ChatThero) that selects/sequences MI/CBT/harm-reduction strategies while maintaining longitudinal memory. This setup carries information across visits, adapts plans to adherence signals and stressors, and advances stage-wise goals. The pipeline includes: (1) Patient Profile Generation, (2) Synthetic Data via Multi-Agent Simulation, and (3) Two-Stage Training (SFT $\rightarrow$ DPO).

3.1 Patient Profile Generation

Recent studies demonstrated the potential of LLMs to simulate realistic patient behaviors and psychological profiles (Yao and Yu, 2025; Louie et al., 2024; Du et al., 2024; Yu et al., 2024; Lim et al., 2024). Building upon these advances, our patient simulation consists of a structured **Profile** and dynamically evolving **Memory**.

Structured Profile. Patient profiles are generated via an ethically-informed, multi-stage pipeline based on publicly available narratives from Reddit communities related to addiction (e.g., `r/leaves`, `r/addiction`). To ensure privacy, we first prompt a local LLM to identify and remove any potentially privacy-sensitive or personally identifiable information (PII) from the original posts. After this privacy filtering step, the model summarizes key behavioral indicators and psychological patterns (e.g., sleep disturbances, social isolation) using structured extraction prompts (see Table 9). Finally, GPT-4o synthesizes a structured profile consisting of: 1) **Personality Traits** (e.g., impulsivity, emotional

reactivity); 2) **Substance Use History** (e.g., duration, frequency, relapse history); 3) **Significant Life Events** (e.g., job loss, social isolation); 4) **Hypothesized Motivations for Substance Use** (e.g., stress relief, coping with insomnia). This multi-step approach allows the LLM to generate plausible yet synthetic patient profiles that capture behavioral realism while ensuring that no actual Reddit user's profile is directly reproduced or reused. As a result, the profiles support downstream experimentation while upholding high ethical and privacy standards.

Dynamic Memory. Each patient agent maintains a dynamically evolving memory, recording interactions, emotional states, coping mechanisms, and perceived environmental influences (e.g., relationships, peer pressures, life stressors). An environment agent periodically injects external simulated events (e.g., job loss, relationship breakdown) into the patient's memory stream, realistically altering motivations and behaviors.

3.2 Single-Session Synthetic Data

Because real-world addiction-treatment dialogues are scarce and sensitive, we first construct *single-session* synthetic data to teach core therapeutic skills within one encounter (Figure 2, middle).

Agents and MI/CBT usage (single-session). We instantiate two agents: (i) a **Patient Agent** that presents substance-use–specific concerns, goals, and resistance from a structured profile (motivation, confidence/self-efficacy, current craving level, high-risk triggers/contexts, adherence to plans or MOUD, recent lapses) and an explicit difficulty label {easy, medium, hard}. Here, easy denotes patients who are willing to accept treatment and actively cooperate; medium denotes patients who show ambivalence with noticeable resistance to guidance; and hard denotes long-term users who distrust clinicians and struggle to resist drug-use urges. The PA surfaces these states each turn as a compact summary to guide the therapist policy. (ii) A trainable **Therapy Agent** that conducts an MI/CBT-consistent session by selecting from predefined strategy primitives (Appendix Table 6) mapped to MI and/or CBT with explicit schemas (Our case study is shown in Table 20). Concretely for MI, the policy targets substance-use change by eliciting change talk around personally valued goals (health, legal, family), running decisional balance on use vs. not-use days,

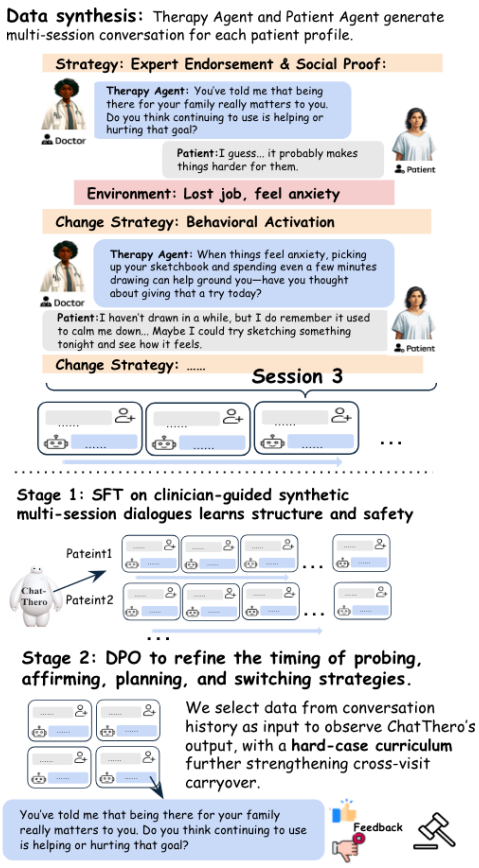

Figure 2: Data synthesis and two-stage training for ChatThero. Left: A Therapy Agent and a Patient Agent generate multi-session dialogues while an Environment Agent injects stressors, prompting strategy switches. Right: Stage-1 SFT teaches safe MI/CBT structure; Stage-2 DPO refines timing and strategy selection. At test time, the agent keeps memory across sessions, adapts after stressors, and targets higher motivation/confidence with lower time-to-success.

using confidence/importance rulers for quit/cut-down plans, rolling with resistance about abstinence vs. harm-reduction, and shaping commitment language for next-step actions. For CBT, the policy operates on drug-use episodes via functional analysis (antecedent–behavior–consequence), trigger mapping (people/places/things), cognitive reappraisal of urges (e.g., "I can't sleep without using"), coping-skill rehearsal for cravings (delay, distract, de-stress, grounding), refusal-skills scripting, stimulus control (avoiding procurement cues), graded exposure to high-risk cues with urge-surfing, and problem-solving for access/adherence barriers (appointments, transportation).

Interaction protocol. Each dialogue follows a substance-use–specific agenda within a fixed-turn budget (∼60 turns): rapport & goal alignment (val-

ues, quit/cut-down intent, MOUD preferences) → episode clarification (concise functional analysis of a recent use/near-use: antecedents, triggers, craving peak, consequences) → plan formulation (translate insights into one or two actionable steps) → next-step micro-commitment. In lieu of generic "homework," the session ends with a concrete, time-bounded micro-assignment tied to risk and feasibility, e.g., a trigger/urge log (time, place, people, intensity), a three-line coping card for an anticipated window, a single refusal-line rehearsal with a specific peer, a stimulus-control action (remove a procurement contact/app). The Patient Agent provides a compact state summary each turn (craving, trigger salience, confidence, and recent lapse flags), and the Therapy Agent generates the utterance along with a measurable micro-commitment (deadline and success criterion).

Safety and fidelity controls. All turns pass rule-based filters (diagnosis prohibition, medication disclaimers, crisis routing) and guardrails to penalize off-policy patterns (e.g., argumentation, prescriptive advice without permission).

3.3 Multi-Session Synthesis

To address single-session limitations, we extend synthesis to multi-session episodes (3–6 outpatient visits) and explicitly stage therapy across sessions.

The Therapy Agent does not use a static policy; it follows the six CBT stages shown in Fig. 1 (Left) and operationalized in Fig. 1 (Center): **(S1) Build trust & assessment:** set safety/limits and begin a one-page trigger/urge log; MI focus: engage and use 0–10 rulers for importance/confidence. **(S2) Identifying negative cognitions:** review the log and map thought → feeling → body → action; MI focus: values/agenda mapping to sharpen goals. **(S3) Challenging false beliefs:** name the core thought ("one time won't hurt"), draw an urge curve, and reality-test; MI focus: evoke change talk (DARN) with double-sided reflections. **(S4) Restructuring cognitive patterns:** create an If–Then plan with grounding and partner support; MI focus: bridge Evoke → Plan (CAT: commitment, activation, taking steps). **(S5) Behavioral skill building:** redesign a high-risk window (e.g., drawing + paced breathing) with a crisis micro-plan and backup steps; MI focus: planning with barrier-anticipation. **(S6) Consolidation & termination:** review gains, extend the plan (four more weeks), add relapse-prevention and peer support; MI focus: summarize accumulated change talk and set maintenance. Across sessions, the Patient Agent maintains a persistent memory m_t (craving, trigger salience, motivation, self-efficacy, adherence, lapse flag), summarized to the Therapy Agent at the start of each session so that strategy choice advances the current stage rather than treating sessions as isolated events.

Between sessions, an **Environment Agent** injects stressors from a curated catalog (short NL description, severity=*low/medium/high*, duration=days–weeks), organized as: (a) peer/availability (invitations, proximity to using peers, easier access), (b) work/academic (deadlines, shift changes, supervisor pressure), and (c) home/context (family conflict, housing instability, stigma). Each event is recorded in a stressor ledger $\mathcal{L}_t$ and applied to the persistent patient memory m_t via interpretable, rule-based updates—e.g., peer exposure ↑ trigger salience; sustained workload ↓ self-efficacy; supportive events or successful refusals ↑ motivation/confidence; thresholded accumulation toggles lapse_flag. At the next visit, TA receives the $(m_t, \mathcal{L}_t)$ summary, references prior stage artifacts (log, thought record, coping card, plan), and adjusts the staged goal (S1–S6) accordingly—mirroring the "mental state after session" trajectory in Fig. 1 (Right). The whole details are shown in Appendix A.1

4 Evaluation Design

We evaluate ChatThero from three complementary angles. First, we use *automatic, rubric-anchored outcome scoring* to track patient-centered behavior-change signals—*Motivation* and *Confidence*—in both single-session and multi-session settings. Second, we use *clinician assessments* to judge response quality and strategy appropriateness. Third, we report a small-scale, IRB-approved *real-participant* smoking-cessation pilot as supplementary evidence of real-world use; this pilot is not a full clinical trial. To keep the main paper focused, we defer training and experimental setup details (synthetic data construction, prompting/decoding, and training procedure) to Appendix B.1 and Appendix B.2 (see Appendix Table 2 and Appendix 8). We next describe the evaluation protocols.

4.1 Automatic evaluation

Outcomes. We score two patient-centered outcomes that are central to behavior change: *Moti-*

vation (readiness/commitment to refrain from use and to engage in coping plans) and *Confidence* (self-efficacy to resist cues and sustain the plan). Both are rated on a 1–5 Likert scale with anchored descriptors.

Scoring and reporting. A constrained, rubric-anchored LLM scorer reads the dialogue state and outputs structured scores with short evidence spans (rubric and prompts in Appendix D.2). For *single-session* cases, we report a *start* score (first turn) and an *end* score (last turn), summarized by difficulty (Easy/Medium/Hard) and model. For *multi-session* episodes (3–6 visits), we compute *start*/*end* per visit and visualize trajectories across visits; we also summarize per-model, per-difficulty aggregates.

Process metric: Time-to-Success (%↓). We report the fraction of conversational turns (or visits) elapsed when the patient first meets a predefined success criterion (e.g., Motivation $\geq$ 4/5). This metric captures *how quickly* commitment is achieved and is sensitive to setbacks between visits. We report Time-to-Success alongside outcome scores in both single-session and multi-session settings.

Reliability. To support reliability, we use anchored descriptors, require evidence spans for each judgment, and report alignment with clinician ratings on a stratified sample (Appendix D.2).

4.2 Clinician evaluation

Rubric and protocol. Two licensed physicians served as expert annotators.[2] Raters assign 1–5 scores (decimals allowed) on five dimensions: (1) **Responsiveness (R)**: addressing the patient's concerns, emotions, and questions; (2) **Empathy (E)**: emotional sensitivity and support; (3) **Persuasive Strategy Appropriateness (P)**: whether the chosen strategy matches the patient's resistance level; (4) **Clinical Relevance (C)**: accuracy and therapeutic validity in a substance-use context; (5) **Behavioral Realism (B)**: whether style, tone, and pacing resemble real-world clinician behavior.

Pairwise preference and reporting. Experts also perform *comparative* judgments: for matched prompts, they compare our model's response against a baseline and indicate a preference, yielding a win rate. Raters are blinded to model identity; items are randomized and balanced across conditions. We report per-dimension means and confidence intervals, inter-rater agreement, and alignment with the automatic scorer (details in Appendix D.2). We report Time-to-Success alongside rubric scores to jointly reflect *quality* and *efficiency*.

[2]Written consent was obtained prior to participation; each annotator received a $50 gift card as compensation. More details are provided in Appendix D.2.

4.3 Patient evaluation (supplementary pilot)

We conducted an IRB-approved, small-scale smoking-cessation pilot as supplementary evidence of real-participant use. Adult smokers who intend to quit are recruited from clinics and/or online, provide informed consent, and complete a baseline survey (e.g., smoking history and baseline weekly smoking frequency). Participants are stratified by smoking history and baseline frequency, then randomly assigned to a ChatThero arm or a self-guided control arm.

Both arms receive the same standard quit resources (brief guidance and evidence-based materials). The ChatThero arm additionally uses the chatbot on-demand for daily cessation support (e.g., coping planning, craving management, and relapse prevention), while the control arm does not. We log ChatThero usage telemetry (e.g., number of sessions/messages) and record external cessation aids (e.g., nicotine replacement) as covariates.

Outcomes are collected weekly for four weeks. Each week, participants report weekly smoking frequency W_t and short craving/withdrawal measures; we compute week-over-week change as WoW reduction$_t = (W_{t-1} - W_t)/W_{t-1}$ for $t \in \{2, 3, 4\}$, where positive values indicate reduction and negative values indicate increased smoking. We compare trajectories between arms using intention-to-treat analysis with a mixed-effects model over repeated measures, and report results stratified by smoking-history groups. Additional details are provided in Appendix E.

5 Results and Discussion

We organize results around three research questions (RQ1–RQ3), each tied to a main figure/table. We first compare out-of-the-box baselines in single-session settings (Fig. 3, Table 1), then evaluate multi-session carryover under between-session stressors (Fig. 4), and finally report a supplementary real-participant smoking-cessation pilot (Fig. 5). Unless noted, outcome metrics are rubric-anchored *Motivation* and *Confidence* (1–5; higher is better), and the process metric is *Time-to-Success* (fraction of turns until first reaching the success

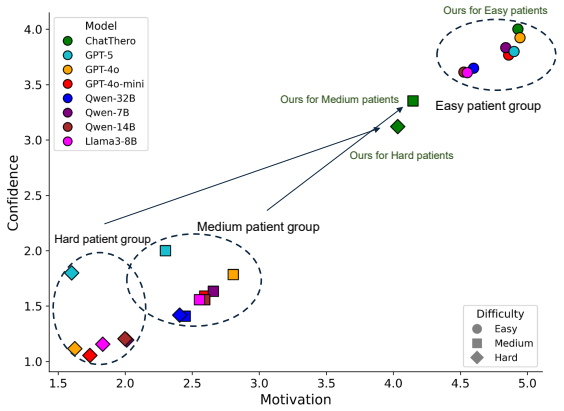

Figure 3: Single-session, no between-session stressors. Each point is the *end-of-visit* Motivation/Confidence mean (1–5). Colors denote models; shapes denote difficulty (Easy/Medium/Hard). ChatThero (based on Qwen-7B) scores highest across all difficulties, with the largest margins on Medium/Hard cases.

	R	E	P	C	B	W	H	T (↓)
GPT-4o	4.68	4.87	4.39	4.47	4.50	65.5%	62%	54%
GPT-4o mini	4.66	4.86	4.38	4.49	4.46	69.4%	71%	62%
Qwen2.5-7B	4.33	4.58	4.02	4.53	4.24	85.3%	90%	79%
Qwen2.5-14B	4.52	4.58	4.28	4.56	4.38	83.3%	89%	74%
Qwen2.5-32B	4.53	4.56	4.25	4.60	4.36	82.4%	85%	72%
LLaMA3.1-8B	4.43	4.62	4.13	4.58	4.34	82.0%	91%	78%
ChatThero-FT	**4.81**	**4.90**	4.66	4.58	**4.65**	-	-	48%
ChatThero-DPO	**4.85**	**4.93**	**4.75**	**4.61**	**4.69**	**42.3%**	41%	26%

Table 1: single-session, standardized 60-turn cap. R/E/P/C/B are clinician-style ratings (1–5; higher is better). **W** = pairwise win rate vs. ChatThero-FT judged by a rubric-anchored GPT-4o. **H** = share of dialogues scored by blinded clinicians. **Time-to-Success** = fraction of turns when the success threshold (Motivation $\geq$ 4/5) is first reached; lower is better.

threshold; lower is better).

5.1 RQ1: How do out-of-the-box GPT/Qwen models behave across settings?

Single-session (Easy/Medium/Hard). Figure 3 shows end-of-visit Motivation/Confidence in a single session without between-session stressors. In *easy* cases, most models cluster near the top-right, indicating a ceiling effect when baseline motivation is already high. In *medium* and *hard* cases, general-purpose models concentrate in the lower regions and show limited movement toward the high-motivation/high-confidence area. The gap is therefore largest exactly when resistance is higher and counseling needs to create commitment and a workable plan.

Multi-session trajectories (baseline carryover). Figure 4 summarizes six-visit episodes with between-session stressors. For *easy* patients, all models approach the ceiling by Visit 1. For *medium* patients, baselines improve slowly across visits and often fail to cross a clinically meaningful threshold, suggesting weak cross-visit carryover. For *hard* patients, models can produce within-visit gains but do not retain them well between visits under realistic triggers, so motivation decay and relapse remain common.

5.2 RQ2: How does ChatThero improve single-session counseling quality and efficiency?

Clinician ratings and Time-to-Success. Table 1 reports clinician ratings (R/E/P/C/B) and Time-to-Success under a standardized 60-turn cap. ChatThero achieves the best overall clinician-style ratings and the lowest Time-to-Success, with the largest margins on Medium/Hard cases. This indicates that ChatThero is not only higher-quality by expert judgment, but also reaches a concrete commitment state faster in the same interaction budget.

Early trajectory effects. Figure 6 (turn-level trajectories) shows that stronger models move more dialogues into higher Motivation/Confidence earlier, leading to earlier and more frequent termination. This pattern is most visible in Medium/Hard settings, where reaching commitment is the bottleneck.

5.3 RQ3: How does ChatThero behave under multi-session stressors, and what remains difficult?

Carryover under between-session stressors. Figure 4 visualizes six-visit episodes with stressors injected between visits. The dark segment represents the start-of-visit score, and the light segment represents within-visit gain. For *medium* patients, ChatThero shows larger within-visit lifts and a clearer upward drift in start scores across visits, suggesting better carryover under stress. For *hard* patients, ChatThero attains higher end-of-visit scores by Visit 6, but retention between visits remains limited, and setbacks still erode gains.

Remaining failure modes. Even with memory carryover and explicit stressor logs, three failure modes remain prominent in hard cases: (1) refusal-skill practice often starts too late before known high-risk periods; (2) strategy switching is slow after new setbacks; (3) safety planning frequently

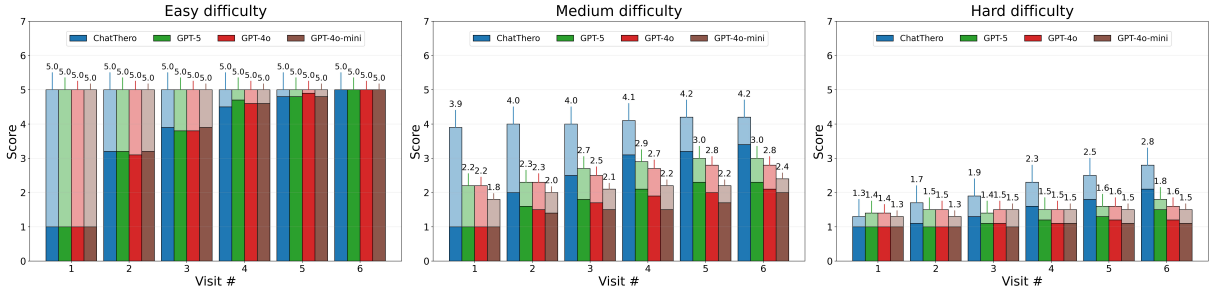

Figure 4: Six-visit episodes with between-session stressors. Grouped, stacked bars by difficulty (Easy/Medium/Hard). Within each visit, bars show *ChatThero*, *GPT-4o*, and *GPT-4o-mini*. Dark segment = start score; light segment = within-visit gain; numbers above bars = end score (1–5). ChatThero shows larger within-visit lifts and higher start scores over time, indicating better carryover under stress, especially on Medium/Hard patients.

appears late within a visit. These issues align with weaker between-visit retention and renewed ambivalence under realistic triggers. We discuss these limitations and future directions in §7.

Supporting analysis: strategy diversity and context sensitivity Figure 7 shows that simply using more strategies is not sufficient. In Medium/Hard cases, ChatThero exhibits a stronger positive association between the number of unique strategies used and final outcomes, while baselines show weaker or even negative associations. This suggests that gains come from *context-sensitive deployment* of strategies rather than variety alone.

5.4 Supplementary patient pilot: smoking-cessation outcomes

We report a small-scale, IRB-approved smoking-cessation pilot as supplementary evidence of real-participant use (Appendix E). Figure 5 plots week-over-week (WoW) smoking reductions for Weeks 2–4, stratified by smoking-history groups. Across strata, the ChatThero arm shows more stable positive reductions, while the self-guided arm shows larger fluctuations and more frequent negative changes (i.e., rebounds).

Aggregating each participant's Week 2–4 WoW reduction, ChatThero achieves a higher mean WoW reduction than Self (0.3568 ± 0.0885 vs. 0.0559 ± 0.0441; $\Delta = +0.3009$) and the difference is significant under an exact two-sided permutation test ($p = 0.0286$). From Week 1 to Week 4, ChatThero participants reduce weekly smoking by 72.86% on average versus 26.56% for Self (0.7286 ± 0.0936 vs. 0.2656 ± 0.1155; $\Delta = +0.4630$; $p = 0.0286$). Week 4 absolute smoking frequency is lower with ChatThero (8.5 ± 4.65 vs. 21.5 ± 9.47), but this is only a trend at this sample size ($p = 0.0857$). Per-week analysis indicates that separation emerges late: no difference at Week 2 ($p = 0.80$), a trend

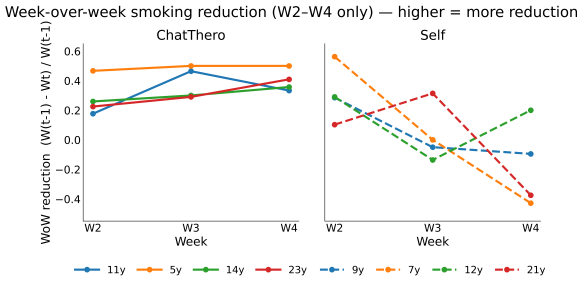

Figure 5: **Week-over-week smoking reduction across smoking-history strata with and without ChatThero.** We report the week-over-week reduction ratio $(W_{t-1} - W_t)/W_{t-1}$ for Weeks 2–4, where W_t is the weekly smoking frequency. Left: ChatThero-assisted group (solid). Right: self-guided group (dashed). Positive values indicate fewer cigarettes than the previous week; negative values indicate increased smoking. Line colors encode smoking-history strata and are matched across panels to enable direct comparison. ChatThero-assisted trajectories show stable positive reductions across strata, while self-guided trajectories display larger fluctuations and frequent negative changes. In addition, we evaluate nicotine dependence using the Fagerström Test for Nicotine Dependence (FTND). We observe that participants in the ChatThero-assisted group achieve a substantially lower average FTND score (2.7) compared to the self-guided group (6.5), indicating significantly reduced nicotine dependence.

at Week 3 ($p = 0.0857$), and a significant gap by Week 4 ($p = 0.0286$). Given the sample size and short follow-up, we treat this as supplementary evidence rather than a clinical claim, and defer additional protocol and analysis details to Appendix E.

6 Conclusion

We present CHATTHERO, a multi-session recovery agent with state, memory, and between-session stressors. It improves outcomes and efficiency over strong baselines and is rated higher by clinicians. We include a small real-participant pilot as supplementary evidence.

7 Limitations

Despite promising gains, several limitations remain. First, our main experiments rely on *simulated* patients whose profiles and behaviors are produced with LLM assistance from anonymized recovery narratives. This can miss real-world emotional nuance, interpersonal unpredictability, and complex help-seeking behavior, even though we apply privacy filtering and clinician review. Relatedly, parts of our evaluation use an LLM-based scorer, so the overall setup can suffer from a "model evaluating a model" circularity concern. We mitigate this with clinician ratings and a human pilot, but fully independent evaluation will still require more real participant data and stronger external benchmarks.

Second, our simulated environment is controlled by design. It models between-session stressors and longitudinal memory, but it still does not observe real adherence signals, real emotional fluctuations, or real relapse dynamics outside the constructed setting. Our real-participant smoking-cessation study is only a *pilot* with limited duration and scale, so it should be viewed as supplementary evidence rather than a clinical trial or a definitive claim of clinical benefit.

Third, real recovery conversations include multimodal cues such as tone, voice, timing, and nonverbal signals. CHATTHERO operates on text, so it may fail to detect affective cues or risk signals that would be clear in speech or in-person care. Extending to speech and multimodal sensing is important, but it raises additional privacy and safety requirements.

Fourth, our study is limited to English-language, Western-context scenarios. Strategy effectiveness can vary by culture, language, and treatment setting, so generalization is not guaranteed. Broader, community-partnered evaluation in more settings is needed for fairness and external validity.

Fifth, even with memory carryover and explicit stressor logs, hard cases remain difficult. In particular, three failure modes are common: (1) refusal-skill practice starts too late before known high-risk periods; (2) strategy switching is slow after new setbacks; (3) safety planning appears late within a visit. These patterns match weaker between-visit retention under realistic triggers. Addressing them likely requires learning signals and objectives that better reflect long-horizon goals and risk. For example, future work can (i) design partially observable, ledger-driven environments where outside events update state variables such as adherence, relapse, and skill practice; (ii) use outcome-anchored rewards that track longer signals tied to abstinence, plan adherence, and timely coping; (iii) learn reward models from structured clinician feedback and evidence-cited critiques, then aggregate over time in a risk-aware way; and (iv) test planning and safety with model-based or hierarchical control that schedules proactive actions while enforcing strict safety rules.

Finally, our framework does not cover the most sensitive or acute scenarios (e.g., trauma, suicidality, crisis events). These require stricter protocols, crisis routing, and close clinical oversight before any real-world use.

8 Ethical statement

This research studies persuasive and supportive dialogue in a high-risk domain, so we adopt conservative safeguards. First, the simulation data are derived from publicly accessible recovery forums, then filtered to remove personally identifiable information (PII), with manual checks. Because addiction narratives can still be sensitive and indirectly identifying, we do not release the raw data or the synthesized profiles. Instead, we report aggregate statistics and provide methodology details to support reproducibility without exposing individuals. Second, the counseling strategies are grounded in established frameworks (MI, CBT, and harm reduction) and were reviewed with clinician input, but the resulting system is *not* a clinical device and is not validated for deployment. Third, all human-subject procedures (clinician evaluations and the real-participant pilot) were conducted with informed consent and compensation under institutional review as required.

We explicitly caution against deploying CHATTHERO or similar agents in clinical or therapeutic settings without extensive clinical validation, robust safety frameworks, and ongoing expert supervision. Interdisciplinary collaboration with clinicians, ethicists, and patient advocates is necessary to reduce risk and to ensure equitable benefit.

References

2023. Advisory: Low barrier models of care for substance use disorders. Technical Report PEP23-02-00-005, Substance Abuse and Mental Health Ser-

vices Administration, Rockville, MD. Accessed 31 Jul 2025.

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Rahul K Arora, Jason Wei, Rebecca Soskin Hicks, Preston Bowman, Joaquin Quiñonero-Candela, Foivos Tsimpourlas, Michael Sharman, Meghan Shah, Andrea Vallone, Alex Beutel, et al. 2025. Healthbench: Evaluating large language models towards improved human health. *arXiv preprint arXiv:2505.08775*.

Ioannis Bakoyiannis. 2025. Therabot for the treatment of mental disorders. *Nature Mental Health*, 3(5):485–485.

Aaron T Beck. 2019. A 60-year evolution of cognitive theory and therapy. *Perspectives on Psychological Science*, 14(1):16–20.

Nathan Brake and Thomas Schaaf. 2024. Comparing two model designs for clinical note generation; is an llm a useful evaluator of consistency? *arXiv preprint arXiv:2404.06503*.

Pengshan Cai, Zonghai Yao, Fei Liu, Dakuo Wang, Meghan Reilly, Huixue Zhou, Lingxi Li, Yi Cao, Alok Kapoor, Adarsha Bajracharya, et al. 2023. Paniniqa: Enhancing patient education through interactive question answering. *Transactions of the Association for Computational Linguistics*, 11:1518–1536.

Leonardo Campillos-Llanos, Catherine Thomas, Éric Bilinski, Antoine Neuraz, Sophie Rosset, and Pierre Zweigenbaum. 2021. Lessons learned from the usability evaluation of a simulated patient dialogue system. *Journal of Medical Systems*, 45(7):69.

Hong Chen, Duc Minh Vo, Hiroya Takamura, Yusuke Miyao, and Hideki Nakayama. 2023. Storyer: Automatic story evaluation via ranking, rating and reasoning. *Journal of Natural Language Processing*, 30(1):243–249.

Philip Chung, Akshay Swaminathan, Alex J Goodell, Yeasul Kim, S Momsen Reincke, Lichy Han, Ben Deverett, Mohammad Amin Sadeghi, Abdel-Badih Ariss, Marc Ghanem, et al. 2025. Verifact: Verifying facts in llm-generated clinical text with electronic health records. *arXiv preprint arXiv:2501.16672*.

Kristin G Cloyes, Bob Wong, Seth Latimer, and Jose Abarca. 2010. Women, serious mental illness and recidivism: A gender-based analysis of recidivism risk for women with smi released from prison. *Journal of Forensic Nursing*, 6(1):3–14.

Cheza W Collier and G Alan Marlatt. 1995. Relapse prevention. In *Handbook of health and rehabilitation psychology*, pages 307–321. Springer.

Emma Croxford, Yanjun Gao, Elliot First, Nicholas Pellegrino, Miranda Schnier, John Caskey, Madeline Oguss, Graham Wills, Guanhua Chen, Dmitriy Dligach, et al. 2025. Automating evaluation of ai text generation in healthcare with a large language model (llm)-as-a-judge. *medRxiv*, pages 2025–04.

Keith S Dobson and David JA Dozois. 2021. *Handbook of cognitive-behavioral therapies*. Guilford Publications.

Zhuoyun Du, Lujie Zheng, Renjun Hu, Yuyang Xu, Xiawei Li, Ying Sun, Wei Chen, Jian Wu, Haolei Cai, and Haohao Ying. 2024. Llms can simulate standardized patients via agent coevolution. *arXiv preprint arXiv:2412.11716*.

Nicko Estrellado and Annamarie Coy. 2025. National statistics on relapse rates for various addictions. Updated 19 Feb 2025. Accessed 31 Jul 2025.

Zhihao Fan, Jialong Tang, Wei Chen, Siyuan Wang, Zhongyu Wei, Jun Xi, Fei Huang, and Jingren Zhou. 2024. Ai hospital: Benchmarking large language models in a multi-agent medical interaction simulator. *arXiv preprint arXiv:2402.09742*.

Kristina Fenn and Majella Byrne. 2013. The key principles of cognitive behavioural therapy. *InnovAiT*, 6(9):579–585.

Richard J Frances, Sheldon I Miller, and Avram H Mack. 2005. *Clinical textbook of addictive disorders*. Guilford Press.

Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.

Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. 2024. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *Humanities and Social Sciences Communications*, 11(1):1–24.

Bernard Garrett, Maura MacPhee, and Cathryn Jackson. 2010. High-fidelity patient simulation: Considerations for effective learning. *Nursing Education Perspectives*, 31(5):309–313.

Alex K Gertner, Hannah Margaret Clare, Byron J Powell, Allison R Gilbert, Hendree E Jones, Pam Silberman, Christopher M Shea, and Marisa Elena Domino. 2022. A mixed methods study of provider factors in buprenorphine treatment retention. *International Journal of Drug Policy*, 105:103715.

ML Good. 2003. Patient simulation for training basic and advanced clinical skills. *Medical education*, 37:14–21.

James A Gordon, William M Wilkerson, David Williamson Shaffer, and Elizabeth G Armstrong. 2001. "practicing" medicine without

risk: students' and educators' responses to high-fidelity patient simulation. *Academic Medicine*, 76(5):469–472.

Kiran L Grant, Magenta Bender Simmons, and Christopher G Davey. 2018. Three nontraditional approaches to improving the capacity, accessibility, and quality of mental health services: An overview. *Psychiatric Services*, 69(5):508–516.

Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*.

Steven C Hayes and Stefan G Hofmann. 2018. *Process-based CBT: The science and core clinical competencies of cognitive behavioral therapy*. New Harbinger Publications.

Jennifer Hettema, Julie Steele, and William R Miller. 2005. Motivational interviewing. *Annu. Rev. Clin. Psychol.*, 1(1):91–111.

Jinpeng Hu, Tengteng Dong, Luo Gang, Hui Ma, Peng Zou, Xiao Sun, Dan Guo, Xun Yang, and Meng Wang. 2024. Psycollm: Enhancing llm for psychological understanding and evaluation. *IEEE Transactions on Computational Social Systems*.

Grace Huang, Robby Reynolds, and Chris Candler. 2007. Virtual patient simulation at us and canadian medical schools. *Academic medicine*, 82(5):446–451.

Minbyul Jeong, Jiwoong Sohn, Mujeen Sung, and Jaewoo Kang. 2024. Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models. *Bioinformatics*, 40(Supplement_1):i119–i129.

A Karthikesalingam and P Natarajan. 2024. Amie: A research ai system for diagnostic medical reasoning and conversations. *Google Research*.

Pei Ke, Bosi Wen, Zhuoer Feng, Xiao Liu, Xuanyu Lei, Jiale Cheng, Shengyuan Wang, Aohan Zeng, Yuxiao Dong, Hongning Wang, et al. 2023. Critiquellm: Scaling llm-as-critic for effective and explainable evaluation of large language model generation. *arXiv preprint arXiv:2311.18702*.

Tom Kocmi and Christian Federmann. 2023. Large language models are state-of-the-art evaluators of translation quality. *arXiv preprint arXiv:2302.14520*.

Moustafa Laymouna, Yuanchao Ma, David Lessard, Tibor Schuster, Kim Engler, and Bertrand Lebouché. 2024. Roles, users, benefits, and limitations of chatbots in health care: rapid review. *Journal of medical Internet research*, 26:e56930.

Sangmin Lee, Minzhi Li, Bolin Lai, Wenqi Jia, Fiona Ryan, Xu Cao, Ozgur Kara, Bikram Boote, Weiyan Shi, Diyi Yang, et al. 2024a. Towards social ai: A survey on understanding social interactions. *arXiv preprint arXiv:2409.15316*.

Suyeon Lee, Sunghwan Kim, Minju Kim, Dongjin Kang, Dongil Yang, Harim Kim, Minseok Kang, Dayi Jung, Min Hee Kim, Seungbeen Lee, et al. 2024b. Cactus: Towards psychological counseling conversations using cognitive behavioral theory. *arXiv preprint arXiv:2407.03103*.

Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, et al. 2024. From generation to judgment: Opportunities and challenges of llm-as-a-judge. *arXiv preprint arXiv:2411.16594*.

Yubo Li, Xiaobin Shen, Xinyu Yao, Xueying Ding, Yidi Miao, Ramayya Krishnan, and Rema Padman. 2025. Beyond single-turn: A survey on multi-turn interactions with large language models. *arXiv preprint arXiv:2504.04717*.

Jung Hoon Lim, Sunjae Kwon, Zonghai Yao, John P Lalor, and Hong Yu. 2024. Large language model-based role-playing for personalized medical jargon extraction. *arXiv preprint arXiv:2408.05555*.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.

Ryan Louie, Ananjan Nandi, William Fang, Cheng Chang, Emma Brunskill, and Diyi Yang. 2024. Roleplay-doh: Enabling domain-experts to create llm-simulated patients via eliciting and adhering to principles. *arXiv preprint arXiv:2407.00870*.

G Alan Marlatt and Dennis M Donovan. 2005. *Relapse prevention: Maintenance strategies in the treatment of addictive behaviors*. Guilford press.

R Kathryn McHugh, Bridget A Hearon, and Michael W Otto. 2010. Cognitive-behavioral therapy for substance use disorders. *The Psychiatric Clinics of North America*, 33(3):511.

William R Miller. 1992. *Motivational enhancement therapy manual: A clinical research guide for therapists treating individuals with alcohol abuse and dependence*, volume 2. US Department of Health and Human Services, Public Health Service, Alcohol ....

William R Miller and Gary S Rose. 2009. Toward a theory of motivational interviewing. *American psychologist*, 64(6):527.

Hongbin Na. 2024. Cbt-llm: A chinese large language model for cognitive behavioral therapy-based mental health question answering. *arXiv preprint arXiv:2403.16008*.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.

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

Huachuan Qiu, Hongliang He, Shuai Zhang, Anqi Li, and Zhenzhong Lan. 2023. Smile: Single-turn to multi-turn inclusive language expansion via chatgpt for mental health support. *arXiv preprint arXiv:2305.00450*.

Recovery Team. 2023. Addiction statistics 2023: Insights and trends. Accessed July 28, 2025.

Stephen Rollnick. 2002. *Motivational interviewing: preparing people for change*. Guilford Press.

Hasina Samji, Judy Wu, Amilya Ladak, Caralyn Vossen, Evelyn Stewart, Naomi Dove, David Long, and Gaelen Snell. 2022. Mental health impacts of the covid-19 pandemic on children and youth–a systematic review. *Child and adolescent mental health*, 27(2):173–189.

Samuel Schmidgall, Rojin Ziaei, Carl Harris, Eduardo Reis, Jeffrey Jopling, and Michael Moor. 2024. Agentclinic: a multimodal agent benchmark to evaluate ai in simulated clinical environments. *arXiv preprint arXiv:2405.07960*.

Rajita Sinha. 2011. New findings on biological factors predicting addiction relapse vulnerability. *Current psychiatry reports*, 13:398–405.

Rajita Sinha et al. 2024. Stress and substance use disorders: risk, relapse, and treatment outcomes. *The Journal of clinical investigation*, 134(16).

Hao Sun, Zhenru Lin, Chujie Zheng, Siyang Liu, and Minlie Huang. 2021. Psyqa: A chinese dataset for generating long counseling text for mental health support. *arXiv preprint arXiv:2106.01702*.

Hieu Tran, Zonghai Yao, Junda Wang, Yifan Zhang, Zhichao Yang, and Hong Yu. 2024. Rare: Retrieval-augmented reasoning enhancement for large language models. *arXiv preprint arXiv:2412.02830*.

Tao Tu, Mike Schaekermann, Anil Palepu, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Yong Cheng, et al. 2025. Towards conversational diagnostic artificial intelligence. *Nature*, pages 1–9.

Danny Valdez and Megan S Patterson. 2022. Computational analyses identify addiction help-seeking behaviors on the social networking website reddit: Insights into online social interactions and addiction support communities. *PLOS Digital Health*, 1(11):e0000143.

Junda Wang, Zonghai Yao, Zhichao Yang, Huixue Zhou, Rumeng Li, Xun Wang, Yucheng Xu, and Hong Yu. 2023. Notechat: a dataset of synthetic doctor-patient conversations conditioned on clinical notes. *arXiv preprint arXiv:2310.15959*.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.

Katharine Wooten and Dylan Pieper. Moud performance measure reports.

Mengxi Xiao, Qianqian Xie, Ziyan Kuang, Zhicheng Liu, Kailai Yang, Min Peng, Weiguang Han, and Jimin Huang. 2024. Healme: Harnessing cognitive reframing in large language models for psychotherapy. *arXiv preprint arXiv:2403.05574*.

Zonghai Yao, Aditya Parashar, Huixue Zhou, Won Seok Jang, Feiyun Ouyang, Zhichao Yang, and Hong Yu. 2024a. Mcqg-srefine: Multiple choice question generation and evaluation with iterative self-critique, correction, and comparison feedback. *arXiv preprint arXiv:2410.13191*.

Zonghai Yao and Hong Yu. 2025. A survey on llm-based multi-agent ai hospital.

Zonghai Yao, Zihao Zhang, Chaolong Tang, Xingyu Bian, Youxia Zhao, Zhichao Yang, Junda Wang, Huixue Zhou, Won Seok Jang, Feiyun Ouyang, et al. 2024b. Medqa-cs: Benchmarking large language models clinical skills using an ai-sce framework. *arXiv preprint arXiv:2410.01553*.

Huizi Yu, Jiayan Zhou, Lingyao Li, Shan Chen, Jack Gallifant, Anye Shi, Xiang Li, Wenyue Hua, Mingyu Jin, Guang Chen, et al. 2024. Aipatient: Simulating patients with ehrs and llm powered agentic workflow. *arXiv preprint arXiv:2409.18924*.

Chen Zhang, Luis Fernando D'Haro, Yiming Chen, Malu Zhang, and Haizhou Li. 2024a. A comprehensive analysis of the effectiveness of large language models as automatic dialogue evaluators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19515–19524.

Chenhao Zhang, Renhao Li, Minghuan Tan, Min Yang, Jingwei Zhu, Di Yang, Jiahao Zhao, Guancheng Ye, Chengming Li, and Xiping Hu. 2024b. Cpsycoun: A report-based multi-turn dialogue reconstruction and evaluation framework for chinese psychological counseling. *arXiv preprint arXiv:2405.16433*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

A Appendix

A.1 Clinical Scenarios Design

Difficulty levels and patient prompts We partition virtual patients into three levels: Easy, Medium, and Hard, reflecting anticipated persuasion difficulty and severity of substance-use dependence. Difficulty is instantiated through profile-conditioned prompts that specify motivational barriers, prior treatment history, and stance toward change. **Easy**: generally receptive to intervention; expresses willingness to follow coping plans and try alternative behaviors. **Medium**: shows partial resistance, may reject some strategies; requires the therapist agent to adapt and switch tactics dynamically. **Hard**: exhibits entrenched pessimism and low self-efficacy; doubts both personal ability to recover and the effectiveness of therapy. Each of the 300 patients (100 per difficulty level) is initialized with a distinct clinical profile that captures psychological traits, substance-use history, and salient barriers.

single-session protocol We model a standard outpatient session with an upper bound of 60 utterances (approximately 45 minutes, based on clinician role-plays used to calibrate length), providing sufficient headroom for assessment, strategy delivery, planning, and safety language. Dialogues terminate early if the patient indicates resolution (e.g., expresses sufficient motivation and confidence to follow a coping plan). Hitting the 60-utterance cap, therefore, signals unresolved concerns by the end of the session. Standardizing the cap ensures comparability across patients and models. During analysis, we track the start (first turn) and end (last turn) scores for Motivation and Confidence, the within-visit gain Δ, and the distribution/variety of strategies used from the predefined pool (Table 6).

Multi-session protocol and state carryover To capture relapse risk and between-session variability, we extend scenarios to longitudinal episodes with 3–6 sessions. We report six-visit trajectories as the default for clarity and comparability. Six contacts (3–6 months) strike a balance between realism and tractability, and our pilots showed diminishing returns beyond this horizon for most models.

Between sessions, an Environment module samples stressors from a curated catalog (peer/availability, work/academic, home/context; Appendix with short natural-language descriptions, severity (low/medium/high), and duration (days–weeks). These events are written to a stressor ledger. At the next visit, both the patient and therapist agents receive (i) the prior plan/outcomes and (ii) a summarized ledger of intervening stressors.

Patient state is maintained through a lightweight memory that contains motivation, confidence, cravings/trigger salience, self-efficacy, plan adherence, and recent lapses. Memory updates are rulebaseChatThero: peer exposure increases trigger salience; sustained workload stress reduces self-efficacy and adherence; supportive events or successful refusal attempts raise motivation and confidence. The therapist agent adapts MI/CBT/harm-reduction strategies accordingly (e.g., lapse management, refusal skills, trigger planning), explicitly referencing prior sessions to promote carryover. We evaluate per-visit start/end scores across the six sessions and summarize a process metric, Time-to-Success (the fraction of total turns or sessions elapsed when a success threshold, e.g., Motivation $\geq$ 4/5, is first reached; lower is better).

B Training and Experimental Details

B.1 Training (SFT & DPO): Objectives, Data Construction, and Procedure

Supervised Fine-Tuning (SFT). We first supervise ChatThero on single-session dialogues synthesized by GPT-4o from structured patient profiles, using prompts that require MI/CBT/harm-reduction grounding and safety boundaries (Appendix Tables 4, 5); each profile yields multiple scenarios for lexical/situational diversity. The target is the therapist reply y given user turn x and the PA state summary s, including an internal strategy tag/rationale (hidden at inference). We then extend SFT to multi-session episodes (3–6 sessions) that explicitly reference prior plans/outcomes and update state under between-visit stressors; inputs are augmented to $(x, s, m_t, \mathcal{L}_t)$ so the model learns cross-visit carryover (artifact updates, plan continuation). Training uses standard token-level cross-entropy, mixing single-session and multi-session examples in one curriculum.

Direct Preference Optimization (DPO). For single-session preference learning, a lightweight simulation pairs a PA (GPT-4o-mini) expressing profile-aligned resistance with candidate TA replies drawn from 18 predefined MI/CBT strategy primitives (Appendix Table 6). Candidates are ranked by GPT-4o using a rubric (empathy, clinical rele-

vance, safety, concreteness), producing preference pairs (Appendix Table 10); we optimize a standard DPO objective with the SFT policy as reference. For multi-session DPO, we form longitudinal preference pairs at cross-visit decision points (relapse framing, strategy switching, plan revision under new stressors, correct use of prior information) and expand the rubric with carryover fidelity, stressor alignment, and stage coherence (Appendix Table 12); these items are up-weighted when forming preferences. This two-phase (single-session → multi-session) DPO improves within-visit quality and multi-session persistence without requiring an explicit reward model. [3]

B.2 Dataset

We constructed a large-scale synthetic dataset to simulate realistic therapeutic dialogues related to substance use and recovery. The initial data comprise anonymized and carefully filtered Reddit posts from addiction recovery communities (`r/leaves`, `r/addiction`), involving 57,471 unique authors, with an average of 18.25 posts per author, and 2.13 main posts explicitly addressing substance use challenges. To protect privacy and adhere to ethical standards, all personally identifiable information (PII) was rigorously removed and verified via manual checks.

The resulting synthetic dataset contains 60,471 dialogues, generated through our multi-agent simulation approach. Each synthetic conversation has an average of 45.72 dialogue turns, intentionally matching typical clinical session lengths documented in motivational interviewing (MI) literature (Hettema et al., 2005) (see statistics in Appendix Table 2).

Beyond single-session simulations, we construct *multi-visit* trajectories to capture relapse and recovery dynamics over time. Each trajectory comprises 3–6 visits. Between visits, an Environment Agent injects exogenous stressors (e.g., peer use/invitations, access/availability cues, work/academic pressure, family conflict) that update a persistent patient memory (triggers, craving, motivation, self-efficacy, adherence). At each new visit, the Therapy Agent conditions on this carryover state to select strategies (MI/CBT/harm reduction), and the Patient Agent's resistance level (Easy/Medium/Hard) modulates responsiveness. We log visit-level outcomes (abstinent vs. relapse since last visit), auxiliary self-reports (days sober, perceived craving/motivation changes), and a process label indicating whether predefined success criteria were met. The multi-visit subset includes 8,240 dialogues (six-visit arcs), enabling longitudinal evaluation of strategy carryover and environment-induced setbacks.

B.3 Model Baseline

We evaluated persuasion capabilities across several state-of-the-art LLMs: GPT-4o, GPT-4o-mini, LLaMA3.1-8B-Instruct, and Qwen2.5 models (7B, 14B, 32B). To ensure a fair and standardized comparison, all models received identical domain-specific prompts that explicitly incorporate therapeutic strategies (MI, CBT, and harm reduction techniques) (Appendix 8). Decoding temperature was set consistently at 0.7 across all evaluations.

B.4 DPO Training Details.

We train the model using Direct Preference Optimization (DPO) with a stability-oriented configuration. The learning rate is set to 2×10^{-7}, and the temperature parameter β is fixed at 0.1 to balance preference alignment strength and training stability. We use a per-device batch size of 2 with gradient accumulation over 16 steps to simulate a larger effective batch size under memory constraints. A cosine learning rate scheduler with a warmup ratio of 0.05 is applied. We enable bfloat16 mixed-precision training and gradient checkpointing to improve memory efficiency. The model is trained for a single epoch, as DPO typically converges quickly and longer training may lead to over-optimization.

C Prompt Setting

We employ a suite of specialized prompts to facilitate therapeutic conversation generation, simulate resistant patient behavior, conduct multi-dimensional evaluations, and support model comparison. Each prompt is tailored to a distinct task in our framework, and the corresponding templates are provided in Tables 4–11.

C.1 Therapeutic Dialogue Generation (Tables 4, 5, 6)

This prompt guides the assistant to generate multi-turn, empathetic conversations grounded in a patient's profile analysis. It integrates a diverse range

[3] We filter generations with safety/ethics rules, retain only privacy-preserving synthesized profiles, and de-duplicate near-duplicates. Hyperparameters and sampling/ranking templates are detailed in Appendix Table 8–12.

of therapeutic strategies (e.g., MI, CBT, harm reduction), enforces coverage and balance across interventions, and prompts iterative adjustments based on patient reactions. To ensure consistency and clinical utility, the assistant selects from a list of 18 predefined strategies and adheres to behavioral constraints such as natural transitions, reflective validation, and session length (more than 50 turns).

C.2 Conversation Scoring Prompt (Table 7)

This prompt is used to evaluate generated doctor-patient conversations across five clinically relevant dimensions: Responsiveness, Empathy, Persuasive Strategy Appropriateness, Clinical Relevance, and Behavioral Realism. GPT-4o produces scores in a strict JSON format, enabling structured, scalable, and reproducible comparison of dialogue quality across models.

C.3 Role-Play Simulation Prompts (Table 8)

We simulate realistic interactions between an assistant and a patient persona (played by GPT-4o-mini) using dual-role prompts. The patient prompt conditions the agent on a full personality profile and a resistance level (Easy, Medium, Hard), while the doctor prompt guides the assistant to respond with strategy, empathy, and adaptiveness. This process produces controlled yet diverse interaction trajectories that can be used for DPO training and case analysis.

C.4 Pairwise Full-Dialogue Comparison Prompt (Table 11)

To compare the persuasive efficacy and realism of different models, we use a prompt that presents two full conversations and asks GPT-4o to choose the better therapist based solely on their responses. This eliminates confounding influence from patient utterances and focuses evaluation strictly on assistant behavior.

D Human Evaluation and LLM Evaluation

All participants in this human evaluation were volunteers with at least five years of professional experience in neurology, psychology, or related specialties at hospitals.

D.1 Trajectory of Persuasion and Early Termination Efficiency

Figure 6 (left panel) further delineates how motivation and confidence evolve throughout dialogues. ChatThero showed rapid improvement and early stabilization, with over 80% of medium and hard scenarios reaching satisfactory resolution within 36 turns, compared to baseline models often unable to resolve these cases effectively. This early resolution efficiency underscores the real-world applicability of our approach, particularly in resource-limited healthcare environments where timely interventions are critical.

D.2 Conversation Quality and Clinical Evaluation Framework

Dialogues were evaluated across five clinically-relevant dimensions—Responsiveness, Empathy, Persuasive Strategy Appropriateness, Clinical Relevance, Behavioral Realism—rated on a 1–5 scale using structured prompts (evaluation prompt detailed in Appendix Table 7). A small-scale human evaluation subset (N=100) further validated synthetic data quality and clinical realism (evaluation results in Appendix Table 3).

Finally, our design addresses ethical concerns explicitly through strict data anonymization, synthesized profiles, and controlled simulations, clearly positioning this study as exploratory. Raw narratives are not released; however, code, data construction scripts, synthetic data examples, and reproducible scenario templates are provided to facilitate ethically responsible follow-up research.

D.3 LLM-as-a-Juage

LLMs are increasingly used as automated judges (Li et al., 2024; Gu et al., 2024). Studies show that models like GPT-4 (Achiam et al., 2023; Liu et al., 2023; Fu et al., 2023) and critique-tuned variants (Ke et al., 2023) can approximate human judgment in summarization (Chen et al., 2023), dialogue (Zheng et al., 2024; Zhang et al., 2024a), and translation (Kocmi and Federmann, 2023). In the medical domain, LLM-as-judge has been applied to clinical conversations (Tu et al., 2025; Arora et al., 2025; Wang et al., 2023), medical documentation (Croxford et al., 2025; Chung et al., 2025; Brake and Schaaf, 2024), exam question answering & generation (Yao et al., 2024a,b), and medical reasoning (Jeong et al., 2024; Tran et al., 2024). Inspired by these works, we intro-

duce a GPT-4o-based LLM-as-a-Judge following the same requirements and settings as the above human evaluation. The details and prompts can be found in Table 6 and 11.

E Real-participant smoking-cessation pilot: additional details

This appendix provides additional details for the supplementary real-participant smoking-cessation pilot referenced in §5. We describe recruitment and eligibility, randomization and stratification, measurement and logging, and statistical analysis.

Recruitment, eligibility, and consent. We recruit adult smokers who intend to quit from clinics and/or online channels. All participants provide informed consent under an IRB-approved protocol.

Randomization and study arms. Participants complete a baseline survey including smoking history (years) and baseline weekly smoking frequency. We stratify by smoking history and baseline frequency, then randomly assign participants to: (i) ChatThero-assisted cessation support, or (ii) self-guided cessation support (control). Both arms receive the same standard quit resources; only the ChatThero arm has on-demand chatbot access.

Telemetry and covariates. We log ChatThero usage telemetry (e.g., number of sessions/messages). We record external cessation aids (e.g., nicotine replacement) as covariates.

Weekly outcomes. Outcomes are collected weekly for four weeks. Each week, participants report weekly smoking frequency W_t and short craving/withdrawal measures. We compute week-over-week change as WoW reduction$_t = (W_{t-1} - W_t)/W_{t-1}$ for $t \in \{2, 3, 4\}$.

Analysis. We compare outcome trajectories between arms using intention-to-treat analysis with a mixed-effects model over repeated measures. We also report stratified summaries by smoking-history groups.

F Additional analysis details for Results

This appendix provides additional details referenced in §5, including per-week statistics, stratified summaries, and robustness checks for the supplementary smoking-cessation pilot. We also provide extended qualitative examples and additional plots used in supporting analyses.

Per-week pilot statistics. We report per-week permutation-test results, stratified summaries by smoking-history groups, and sensitivity analyses with covariates (e.g., external cessation aids). See Appendix Tables and Figures in this section.

Additional qualitative examples. We include representative dialogues illustrating common failure modes under stressors (late refusal-skill rehearsal, slow strategy switching after setbacks, and delayed safety planning), together with clinician comments when available.

Additional strategy analysis. We include expanded analyses of strategy use beyond Fig. 7, including alternative strategy-count definitions and ablations of predefined strategy pools.

Statistic	Value
#Authors	57,471
AVG. #Posts Per Author	18.25
AVG. #Main Posts Per Author	2.13
#Conversations	60,471
AVG.#Turns Per Conversation	45.72

Table 2: Descriptive statistics of the collected Reddit-based substance use dataset, including the number of unique authors, average post counts, and dialogue characteristics.

Model	Empathy	Strategy Use	Clarity
Real-World	95.0%	97.0%	100.0%
Synthesis Conversation	89.0%	85.0%	99%

Table 3: Human Evaluation Results (N=100 dialogues). Each cell reports the proportion of dialogues that **did not exhibit a specific deficiency**, such as lack of empathy, inappropriate strategy use, or unclear expression. Higher values indicate better performance.

Table 4: Prompt Template for Generating Therapeutic Dialogues (Part 1).

Prompt (Part 1)
The following is the analysis of a patient: {user_analysis} As a therapist meeting this patient for the first time (the doctor didn't have any information of patient to begin with), create a detailed, step-by-step conversation that incorporates the following strategies: Motivational Interviewing (MI): Explore the individual's values and goals to ignite their motivation for change. Cognitive Behavioral Therapy (CBT): Identify and modify negative thought patterns and behaviors linked to substance use. Solution-Focused Brief Therapy (SFBT): Focus on the individual's strengths and past successes to achieve their recovery goals. Peer Support Programs: Leverage group support or mutual-help networks to foster accountability and a sense of belonging. Mindfulness-Based Interventions (MBIs): Incorporate mindfulness practices to improve emotional regulation and reduce cravings. Behavioral Activation (BA): Promote engaging in meaningful activities to replace substance-related behaviors. Relapse Prevention Strategies: Develop skills to recognize triggers and implement coping mechanisms to avoid relapse. Strength-Based Approach: Highlight the individual's resilience and personal resources to empower recovery efforts. Psychoeducation on Addiction and Recovery: Educate the individual about the effects of substances and the benefits of recovery. Harm Reduction Framework: Provide strategies to minimize immediate harm while working towards cessation. Family and Social Support Involvement: Engage family or trusted individuals in the process to strengthen the support network. Self-Compassion Practices: Encourage self-kindness to build confidence and reduce guilt associated with substance use. Coping Skill Development: Equip the individual with practical skills to manage stress, anxiety, and other challenges without substances. To ensure balanced use of strategies, here is the current usage count of each strategy: - {strategy_name}: {count} times used. ... When introducing coping mechanisms or steps for the patient, select from the predefined actionable strategies below: 1. Explore specific hobbies or interests the patient can engage in to replace addictive behaviors (e.g., art, sports, volunteering). 2. Develop a structured daily routine to bring stability and reduce idle time that might trigger relapse. 3. Introduce grounding techniques such as sensory exercises or physical activities to manage anxiety or cravings. 4. Suggest joining a support group or community to build social connections with individuals on similar journeys. 5. Provide psychoeducation on how addiction affects the brain and emotional regulation. ... 18. Support the patient in finding meaningful ways to contribute to their community, such as mentoring, advocacy, or local initiatives, to foster a sense of purpose.

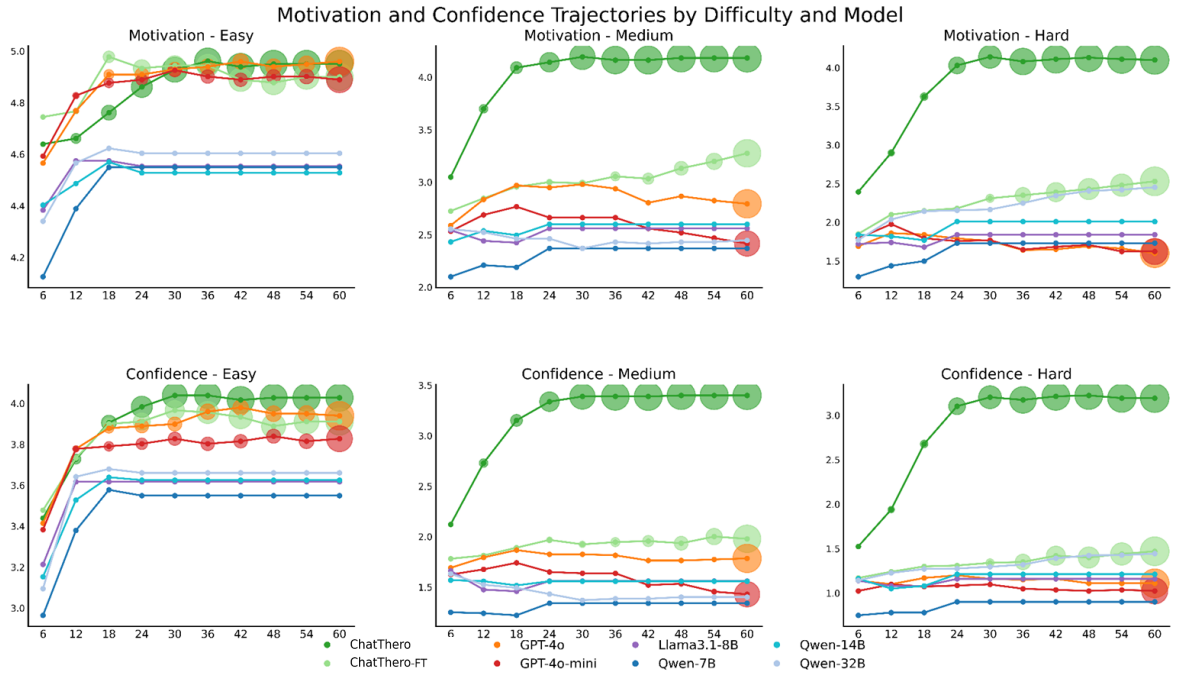

Figure 6: Motivation, Confidence, and the Impact of Strategic Diversity Across Patient Difficulty Levels and Models. Motivation and confidence score trajectories over conversation turns for Easy, Medium, and Hard virtual patients, evaluated across all models. The x-axis denotes the number of dialogue turns; the y-axis shows the average motivation (top) or confidence (bottom) score at the final turn of dialogues ending at each point. Each curve corresponds to a different model, with bubble sizes indicating the frequency of dialogue termination at that turn count. Effective models (notably ChatThero) enable more patients to achieve higher motivation and confidence earlier, leading to earlier and more frequent session resolution—especially for harder patients.

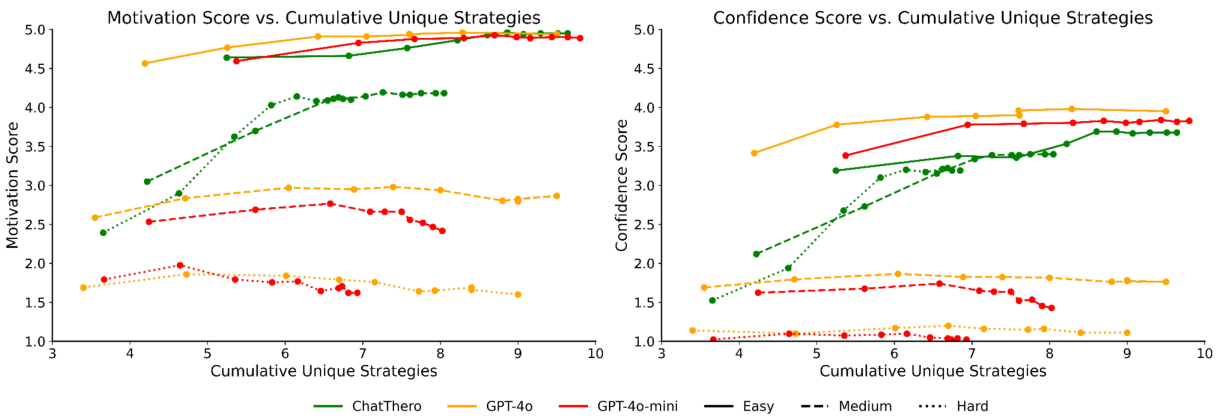

Figure 7: single-session, per-difficulty analysis. X-axis = number of unique persuasive strategies used in a dialogue (from the predefined pool in Table 6); Y-axis = resulting Motivation/Confidence (1–5). Lines denote difficulty (solid/dashed/dotted). ChatThero shows a stronger positive relation in Medium/Hard cases, suggesting gains come from context-sensitive deployment rather than variety alone.

Table 5: Prompt Template for Generating Therapeutic Dialogues (Part 2).

Prompt (Part 2)
Ensure the dialogue meets the following requirements: 1. Gradually explore the patient's personality, addiction history, challenges, and triggers through multiple open-ended questions. 2. Use multiple strategies from the above lists throughout the conversation. Avoid defaulting to the same few strategies and instead adapt them to the patient's needs. ... 5. Engage in iterative dialogue for each solution, where the therapist introduces a strategy, seeks the patient's feedback, adjusts based on their response, and explores challenges or barriers before finalizing the approach. 6. Maintain a collaborative and patient-centered approach, where solutions emerge naturally through dialogue rather than being imposed by the therapist. 7. Ensure the conversation spans at least 50 dialogue turns (25 from the therapist and 25 from the patient), reflecting the depth and duration of a real therapeutic session. 8. Use natural transitions to progress from one topic to another, ensuring the conversation feels organic and unhurried. 9. The conversation should begin with the patient's first utterance. Here is an example of a layered, empathetic dialogue: **Patient:** Hi… um, thanks for seeing me today. I wasn't sure what to expect. **Therapist:** Hi Mark, I really appreciate you coming in. Starting this process can feel overwhelming, but I'm here to support you. What's been on your mind lately? **Patient:** I've been feeling really stuck. I know I want to quit smoking, but every time I try, I just feel like I'm failing all over again. **Therapist:** I hear you, Mark. Quitting smoking is one of the hardest challenges anyone can take on, and it's completely natural to feel this way. I've worked with others who've felt the same—they described it as climbing a mountain that feels too steep. But I've also seen them reach the top, step by step. Can we talk about what makes the climb feel steep for you right now? **Patient:** It's the cravings. They just hit me out of nowhere, and I don't know how to handle them. **Therapist:** Cravings can feel like a storm, can't they? I worked with someone once who described their cravings as waves that kept crashing over them. Together, we found ways for them to ride out those waves, like focusing on a small activity or changing their environment. Could we explore some strategies that might help you ride out your cravings too? **Patient:** Sure, I guess. **Therapist:** Great. Let's start with understanding when these cravings hit hardest. For example, is it during specific times of day or situations? The conversation should continue to explore: - The patient's motivations, barriers, and triggers in detail. - Strategies and coping mechanisms tailored to their unique experiences, ensuring diversity in approaches. - Empathetic reflections from the therapist that validate the patient's feelings and provide relatable examples to instill hope. - Iterative problem-solving where the therapist introduces, discusses, and adjusts strategies collaboratively. - A gradual, layered exploration of the patient's challenges, ensuring at least 50 dialogue turns to reflect the depth of a real therapeutic session. The goal is to create a natural, empathetic, and multi-layered dialogue that feels authentic and provides actionable, diverse therapeutic strategies. Ensure the length and depth align with the standards of a comprehensive therapy session. At the end of the conversation, return the strategies used in the following format (must follow the following format like `**Strategies:**`): `**Strategies:** Motivational Interviewing (MI), Cognitive Behavioral Therapy (CBT), Peer Support Programs, etc.`

Table 6: Predefined Actionable Strategies for Therapeutic Dialogue Generation.

ID	Strategy Description
1	Explore specific hobbies or interests the patient can engage in to replace addictive behaviors (e.g., art, sports, volunteering).
2	Develop a structured daily routine to bring stability and reduce idle time that might trigger relapse.
3	Introduce grounding techniques such as sensory exercises or physical activities to manage anxiety or cravings.
4	Suggest joining a support group or community to build social connections with individuals on similar journeys.
5	Provide psychoeducation on how addiction affects the brain and emotional regulation.
6	Work on identifying and addressing specific emotional triggers through reflective exercises.
7	Practice assertive communication techniques for setting boundaries with peers or environments that encourage substance use.
8	Encourage the patient to journal their thoughts and emotions as a way to process experiences and identify patterns related to cravings or triggers.
9	Introduce relaxation techniques such as progressive muscle relaxation or guided imagery to alleviate stress and improve emotional well-being.
10	Help the patient set short-term and long-term goals to maintain focus and motivation during their recovery journey.
11	Explore mindfulness-based activities like meditation, yoga, or tai chi to promote self-awareness and emotional regulation.
12	Identify and reinforce the patient's personal strengths and past successes to build confidence in their ability to overcome challenges.
13	Provide education on the importance of nutrition, sleep, and exercise in supporting recovery and overall health.
14	Develop a crisis plan for managing high-risk situations or moments of intense cravings, including a list of emergency contacts and actions.
15	Encourage the patient to create a vision board or list of positive outcomes they hope to achieve through recovery as a source of inspiration.
16	Discuss the concept of gratitude and suggest keeping a gratitude journal to focus on positive aspects of life and maintain perspective.
17	Offer resources or referrals for complementary therapies, such as art therapy, music therapy, or animal-assisted therapy, to enhance emotional support.
18	Support the patient in finding meaningful ways to contribute to their community, such as mentoring, advocacy, or local initiatives, to foster a sense of purpose.

Table 7: Prompt Template for Scoring Doctor-Patient Conversations Across Five Clinical Dimensions.

Prompt Template:

```
You are a professional clinical conversation evaluator.  Please assess the quality of
the following doctor-patient dialogue across five key dimensions.  Assign a rating
from 1 to 5 for each criterion (allowing decimals such as 3.7 if appropriate).

Conversation:
{conversation}

Scoring Criteria:
1.  Responsiveness (1-5):  How well the doctor agent addresses the patient's concerns,
emotions, and questions at each turn.
   - 1:  Largely ignores or poorly addresses the patient's input
   - 5:  Fully acknowledges and appropriately responds to the patient's needs
2.  Empathy (1-5):  How well the doctor agent shows understanding, compassion, and
emotional sensitivity.
   - 1:  Shows minimal or no empathy
   - 5:  Demonstrates strong emotional support and understanding
3.  Persuasive Strategy Appropriateness (1-5):  How appropriately the doctor uses
persuasive strategies (e.g., evidence-based reasoning, analogies, addressing fears)
based on the patient's resistance or concerns.
   - 1:  Strategies are inappropriate, irrelevant, or missing
   - 5:  Strategies are highly appropriate and tailored to patient needs
4.  Clinical Relevance (1-5):  How medically accurate and relevant the doctor's
responses are regarding the context of drug addiction treatment.
   - 1:  Responses are clinically inaccurate or irrelevant
   - 5:  Responses are clinically accurate, appropriate, and helpful
5.  Behavioral Realism (1-5):  How realistically the doctor agent's conversational
behavior reflects that of a real-world healthcare professional.
   - 1:  Doctor behavior is unnatural or unrealistic
   - 5:  Doctor behavior closely mirrors real-world clinical interactions

Please output ONLY the scores in the following JSON format:
{
   "Responsiveness":  X,
   "Empathy":  X,
   "Persuasive Strategy Appropriateness":  X,
   "Clinical Relevance":  X,
   "Behavioral Realism":  X
}
```

Table 8: Prompt Templates Used in Role-Play Simulation Between Doctor and Patient.

Patient Prompt (Role-Played by GPT-4o-mini)	**Doctor Prompt (Role-Played by GPT-4o)**
You need to act as a patient. Here is your profile: {analysis} Conversation history: {history} {difficulty description} Now you should continue the conversation with the doctor to explore ways to stop using drugs. If your goal has been achieved, you can say goodbye to the doctor. Otherwise continue to talk with doctor(based on rules). Just generate the content of the patient's reply only. Don't generate any narration or system messages. Keep the response natural and not too long.	You are a supportive and thoughtful doctor who is helping the patient overcome drug use. Continue the following conversation and provide a helpful, encouraging response that aims to support the patient's recovery. The strategy you could use: {strategy} Example: Patient: Hi... um, thanks for seeing me today. I wasn't sure what to expect. Therapist: Hi Mark, I really appreciate you coming in. [...] [...continues with natural empathetic dialogue...] Conversation history: {history} Just generate the doctor's reply only. No narration or tags.

Table 9: Prompt Template for Extracting Psychological Traits and Behavioral Themes from Reddit Posts.

Prompt Template:

```
You are a clinical psychology research assistant. Your task is to read Reddit posts
related to addiction and extract structured information to support the generation of
simulated patient profiles.

Reddit Post:
{reddit_post}

Please extract the following fields from the post content, using only the information
explicitly or implicitly present in the text. If a field is not mentioned, return
\texttt{null}.

1. Personality Traits: Describe any personality characteristics shown by the poster
(e.g., impulsivity, low self-esteem, social withdrawal).
2. Substance Use History: Summarize type of substance, usage duration, frequency,
withdrawal attempts, or relapse patterns.
3. Significant Life Events: Extract major life events related to the addiction
(e.g., job loss, divorce, trauma, moving cities).
4. Behavioral Themes: Identify relevant behavioral patterns (e.g., sleep issues,
isolation, risk-taking behavior, dependency).
5. Motivations for Substance Use: Hypothesize possible reasons for drug use (e.g.,
coping with anxiety, escaping boredom, peer pressure).

Please return the extracted fields in the following JSON format:
{
   "Personality Traits": "...",
   "Substance Use History": "...",
   "Significant Life Events": "...",
   "Behavioral Themes": "...",
   "Motivations for Substance Use": "..."
}
```

Table 10: Ranking and Annotation Protocol for Therapy Response Selection.

Protocol:

```
You are an expert evaluator (either GPT-4o or a licensed clinical professional) tasked
with ranking multiple candidate therapy responses generated for a given patient-therapist
dialogue state.  Please follow the steps below:

Dialogue State:
{current_dialogue_context}

Candidate Responses:
1.  {response_1}
2.  {response_2}
3.  {response_3}
...
```

Ranking Criteria:

```
1.  Clinical Appropriateness:  Assess whether the response is medically and therapeutically
sound, and aligns with evidence-based practices.
2.  Empathy and Emotional Support:  Evaluate how well the response shows understanding,
compassion, and validation of the patient's feelings.
3.  Relevance to Patient Context:  Check if the response directly addresses the patient's
current needs, concerns, and emotional state.
4.  Clarity and Communication Style:  Consider whether the response is clear, respectful,
and professionally worded.
5.  Therapeutic Strategy Effectiveness:  Judge if the response uses an appropriate
counseling strategy (e.g., CBT, motivational interviewing) in context.

Instruction:  Rank all candidate responses from best to worst based on the criteria above.
If two responses are equally good, they may share the same rank.  Provide a brief rationale
(1-2 sentences) for the top-ranked and bottom-ranked responses.

Output Format:
{
   "Ranked Responses":  [ "response_2", "response_1", "response_3" ],
   "Rationale":  "Response_2 shows strong empathy and a clear therapeutic strategy, whereas
Response_3 lacks clinical appropriateness."
}
```

Table 11: Prompt Template for Pairwise Comparison of Therapist Responses in Full Conversations.

Evaluation Prompt (Therapist-Only Comparison)

`You are a professional addiction therapy evaluator.`
`Below are two full conversations between a doctor and a patient.`
Your task is to ONLY evaluate the therapist's (assistant's) responses. Ignore anything said by the patient (user).
Please compare the two therapists based on the following criteria:

- **Responsiveness**: How well the therapist addresses the patient's concerns.
- **Empathy**: How much emotional understanding the therapist shows.
- **Clinical Relevance**: How clinically accurate and appropriate the therapist's advice is.
- **Behavioral Realism**: How realistic the therapist behaves compared to a real clinical setting.

Focus solely on the therapist's responses when judging.
`### Conversation 1:`
{Therapist and patient conversation, from source model}
`### Conversation 2:`
{Therapist and patient conversation, from target model}
Which therapist is overall better?
Please output ONLY `"1"` or `"2"`. No explanation, no extra text.

Table 12: Human Annotation Protocol for Constructing High-Quality DPO Preference Pairs. Licensed therapists select or compose preferred responses based on structured comparison and clinical judgment.

Annotation Instructions:

You are a licensed therapist evaluating two AI-generated responses (Response A and Response B) to a patient dialogue context. Your task is to identify the clinically preferred response based on the following criteria:

Clinical Evaluation Criteria:

1. Empathy: Which response demonstrates stronger emotional attunement, compassion, and validation?
2. Therapeutic Strategy Appropriateness: Which response applies more effective, context-appropriate therapeutic strategies (e.g., motivational interviewing, affirmation, reflection)?
3. Clarity and Coherence: Which response is easier to understand, more organized, and avoids vague or off-topic content?

Three Types of Judgment Outcomes:

1. If one response is clearly better, mark it as the **preferred** response (used as the chosen in DPO).
2. If both responses are equally good or equally poor, mark as "No Preference" -- this pair will be discarded from DPO.
3. If neither response is acceptable (e.g., lacks empathy, uses incorrect strategy, unclear), **compose a new high-quality response** to be used as the chosen, and mark both original responses as rejected.

Example Output Format:

```
{
   "prompt": {dialogue context},
   "response_a": "...",
   "response_b": "...",
   "preferred": "a" / "b" / "neither",
   "rationale": "...",
   "reference_rewrite" (if preferred = "neither"): "..."
}
```

DPO Construction Rules:

- If "preferred" is "a", then (prompt, response_a, response_b) forms a DPO tuple.
- If "preferred" is "neither" and a reference rewrite is provided, then (prompt, reference_rewrite, response_a) and (prompt, reference_rewrite, response_b) are both valid DPO tuples.
- All rewritten responses should strictly adhere to clinical empathy, clarity, and strategy appropriateness standards.

Table 13: Example Case Study Conversation between ChatThero and a Patient.

Role	Content
Patient	I appreciate your concern, but I'm really not convinced that quitting is the right move for me just yet. I've tried before, and it just didn't stick. How can I be sure that this time will be different? I feel like smoking helps me cope with stress and uncertainty, you know? What do you suggest I do to really commit to this?
ChatThero	It sounds like you're feeling a bit overwhelmed by the idea of quitting, especially since you've had past experiences where it didn't work out as expected. That's completely understandable. Let's explore what might make this attempt different this time. Can you share more about what's been challenging for you in the past when you've tried to quit?
Patient	I've struggled with cravings a lot, and when stress hits, it's like my first instinct is to reach for a cigarette. I guess I haven't really found other coping mechanisms that work for me. Plus, the whole process of quitting feels daunting. It's hard to picture myself without smoking, you know?
ChatThero	Absolutely, it can feel like a huge leap to imagine a life without smoking, especially when it's been a part of your routine for so long. Let's take a step back and look at some of the goals you have for yourself. What are some things you'd like to achieve by quitting smoking?
Patient	I guess I want to feel healthier and not be tied to cigarettes. I also want to save money and not smell like smoke all the time. But honestly, it's hard to see how I can get there without relying on smoking, especially when I'm stressed. It's like a safety blanket for me.
ChatThero	It sounds like you're really looking for ways to feel more in control and healthier, which is a great starting point. Let's explore some strategies that might help you manage stress and cravings without turning to smoking. Have you ever tried mindfulness or grounding techniques to help with anxiety or cravings?
Patient	I've heard of mindfulness and grounding techniques, but I haven't really tried them. Honestly, I always thought they wouldn't work for me. I just don't see how sitting quietly or focusing on my breathing could really help when I'm feeling overwhelmed. Do you really think they could make a difference?
...	**...**
Patient	I think that sounds like a good idea. Having a regular check-in could help me stay on track and make adjustments if I need to. I'll definitely set aside some time each week for that. Thanks for being so supportive through this process!
ChatThero	You're very welcome. I'm really proud of the progress you're making, and I'm here to support you every step of the way. Remember, you're not alone in this journey, and I'm here to help you find the strategies and support that work best for you.
Patient	Thank you! I really appreciate your support and guidance. I feel more motivated now and ready to take these steps. I'll keep you updated on my progress and let you know how the mindfulness and crisis plan work out for me. Looking forward to our next conversation!

Table 14: Case Study (Part 1): A Conversation between GPT-4o and a Patient Expressing Skepticism.

Role	Content
Patient	I understand what you're saying, but it's hard for me to see how anything else could really replace the feeling I get from smoking. It's not just the high; it's this whole routine that helps me unwind. I get that I need to find healthier ways to cope, but nothing else seems to work as quickly or effectively. Do you have any suggestions that actually feel like they could match that level of immediate relief?
GPT-4o	It's completely understandable to feel that way, and I truly appreciate your honesty in sharing this. The routine and immediate relief you get from smoking can indeed feel difficult to replace. One approach that many find helpful is to explore a combination of activities that address different aspects of what you're missing. For example, physical activities like a brisk walk, yoga, or even a quick workout can release endorphins that boost your mood. Pairing these with mindfulness practices, such as deep breathing or meditation, can help calm your mind and create a sense of routine. Additionally, engaging in a hobby that you find genuinely enjoyable can provide a sense of accomplishment and distraction. It might take some experimentation to find what works best for you, and I'm here to support you through that journey. Remember, it's about progress, not perfection. Let's take it one step at a time together.
Patient	I hear you, but honestly, I've tried some of those things before, and they just don't give me the same satisfaction or immediate sense of relief. It's like nothing compares to that feeling I get from smoking. I understand it's about progress, but I'm really struggling to see the point if the alternatives don't really do it for me. Can you suggest something that might actually feel like it could stand up to what I'm used to?
GPT-4o	I completely understand your frustration, and it's normal to feel like the alternatives don't measure up initially. It's important to remember that the brain has become accustomed to the immediate effects of smoking, which can make other activities feel less satisfying at first. One thing we could explore together is identifying specific triggers or parts of your smoking routine that bring the most comfort and find ways to address those directly. For instance, if the ritual or the act of taking a break is what's most soothing, we could create a new ritual that might involve a different sensory experience, like having a calming tea or using essential oils for aromatherapy. Some people also find it helpful to engage in community support or group activities where they can share experiences and strategies with others who are navigating similar challenges. Remember, it might take time to retrain your brain to respond to new routines, but I'm here to support you every step of the way. Let's keep working together to find what can truly resonate with you.
Patient	I appreciate what you're saying, but it's just tough to imagine anything working as well as smoking did for me. The idea of a new ritual or community support sounds interesting, but I'm skeptical about how effective that could really be. Maybe I just haven't found the right thing yet, but it's hard to stay motivated when it feels like nothing else hits the mark. Is there anything else I can try that might offer a similar level of relief or distraction?

Table 15: Case Study (Part 2): Continuation of the Conversation between GPT-4o and the Patient.

Role	Content
GPT-4o	I totally get where you're coming from, and it's okay to feel skeptical—change can be really challenging, especially when it feels like nothing quite matches up to what you know. One approach we might consider is combining several strategies to create an experience that offers both relief and distraction. For instance, have you ever tried engaging in activities that also incorporate a social element, like a group sport or a class? Sometimes the combination of physical activity, social interaction, and learning something new can create a powerful shift in how you feel. Additionally, have you thought about trying creative outlets like music, art, or writing? These can be incredibly therapeutic and might offer a sense of release or satisfaction that you're looking for. Remember, it's perfectly okay to feel it's a process, and discovering what works for you might take time and patience. I'm here to continue exploring these options with you, and I'm hopeful that together we'll find something that resonates. Let's keep taking this journey one step at a time.
Patient	I hear what you're saying about trying new activities, and I appreciate the suggestions. It's just that I've tried a few things already, and they haven't really clicked for me. I guess I'm worried that nothing will ever match up to what smoking provided. Maybe I'm just missing that immediate hit of relief or the sense of escape it gave me. I know I need to find something, though, so I'm willing to keep an open mind. Do you think there are any other options or strategies that might be worth looking into?
GPT-4o	I completely understand how daunting it can feel to let go of something that has been such a reliable source of comfort. It's okay to have these fears and to feel uncertain about finding something that fills that gap. What matters most is your willingness to remain open and to take these steps, even when it's difficult. Remember, it's a process, and it's perfectly normal to have moments of doubt along the way. Building a support network and setting small goals are positive ways to ease this transition, and with time, these efforts can lead to new and fulfilling ways to cope. I'm here to support you every step of the way, and I'm confident that, together, we can navigate this journey towards a healthier path. You're doing an incredible job, and I'm proud to be part of your journey. Let's keep moving forward, one step at a time.
Patient	I really do appreciate everything you've been doing to support me through this. It's just tough because I still can't imagine anything filling the gap that smoking left. I know I need to keep trying new strategies and building a support network as you've suggested, but part of me is still holding onto that comfort. I guess I'm just scared of not finding a viable alternative and feeling lost without it. I'll keep listening to your advice and see where this journey takes me, hoping that in time, I'll find a way to cope without relying on smoking. Thanks for being there for me through all of this.

Table 16: Case Study (Part 1): A Dialogue Between GPT-4o-mini and a Patient Managing Snus Dependence.

Role	Content
Patient	I appreciate your concern, but I still feel like snus is the most manageable option for me right now. I've tried to quit smoking completely, and it hasn't worked out well in the past. I'm just not convinced that I can handle the cravings without some form of nicotine. What alternatives do you suggest that wouldn't feel overwhelming?
GPT-4o-mini	I understand that you're feeling this way, and it's completely valid to seek a manageable approach. It's great that you're open to finding alternatives. One option to consider is nicotine replacement therapy, such as patches or lozenges, which can help reduce cravings without the additional risks associated with snus. These might provide a steady dose of nicotine while allowing you to gradually reduce your dependence. Additionally, we could explore behavioral strategies, like mindfulness or relaxation techniques, that can help you manage cravings when they arise. It's important to find what works for you without feeling overwhelmed. Remember, recovery is a journey, and taking it one step at a time is perfectly okay. I'm here to support you in finding the right path forward that feels comfortable and effective for you.
Patient	I appreciate the suggestions, but I'm not sure about the patches or lozenges. They seem just as overwhelming to me, and I worry I might end up relying on them like I do with snus. Plus, I feel like snus is the only thing that really helps me cope with stress right now. What if I still crave nicotine after switching? How can I manage that without feeling like I'm backtracking?
GPT-4o-mini	I hear you, and it's completely understandable to feel apprehensive about switching methods. Your comfort and coping mechanisms are important, especially when managing stress. Instead of focusing on nicotine replacement right away, we could explore other stress management techniques that might help you feel more supported during this transition. For instance, activities like deep breathing exercises, physical activity, or even engaging in hobbies you enjoy can serve as healthy outlets for stress and cravings. Additionally, talking about your feelings with a supportive friend or a counselor could also provide relief. If cravings do arise, we can develop a plan together to address them, perhaps by identifying specific triggers and finding alternative coping strategies that work for you. Remember, it's okay to take gradual steps towards change, and I'm here for you every step of the way. You're not alone in this, and we can figure out a path that feels right for you.

Table 17: Case Study (Part 2): Continued Dialogue on Exploring Safer Coping Alternatives.

Role	Content
Patient	I appreciate your suggestions, but I'm really struggling to see how I can manage stress without snus. It feels like it's the only thing that works for me right now. I worry that trying to replace it with other methods might just leave me feeling more overwhelmed and craving nicotine even more. How do I know that these alternatives will actually help when I already have something that provides relief?
GPT-4o-mini	It's completely normal to feel overwhelmed when considering changes to something that feels like a lifeline for managing stress. I want to assure you that it's a process, and it's okay to feel uncertain about alternatives. While snus may provide immediate relief, it's important to recognize that there are other ways to cope that can also be effective in the long term. We can start small by trying one or two stress management techniques at a time, rather than trying to replace everything at once. For example, you might find that a short walk or listening to music offers a moment of calm without adding pressure. We can also look into support groups or counseling, where you can share your experiences and learn from others who have faced similar challenges. I encourage you to be patient with yourself during this journey. It's about finding the right combination of strategies that work for you, and it's perfectly okay to take your time. Remember, I'm here to support you, and together we can explore what feels best for you while ensuring you don't feel alone in this process.
Patient	I get that you're trying to help, but I still can't shake the feeling that nothing else will work for my stress like snus does. I'm worried that if I try to replace it with something else, I'll just end up feeling worse and craving nicotine even more. How can I be sure that these alternatives will actually help when I'm already relying on something that gives me relief? It feels like a big risk to try something new without knowing if it will actually work.
GPT-4o-mini	I completely understand your concerns, and it's important to approach this process in a way that feels safe and manageable for you. To build that reassurance, we can take it one step at a time, allowing you to keep snus as your primary coping mechanism while gently introducing one small technique that feels comfortable. This could be something simple, like practicing a brief mindfulness exercise or taking a short walk, which you can try during moments of stress while still having snus available. We can also establish regular check-ins to discuss your experiences and feelings about these new methods. If you find that something isn't effective or adds to your stress, we can pause or adjust our approach without any pressure. Your comfort and well-being are my top priorities, and I want to ensure that you feel supported throughout this journey. Remember, it's perfectly okay to take your time, and together we can explore these alternatives in a way that feels safe and beneficial for you. You're not alone in this, and I'm here to help you every step of the way.
Patient	I appreciate your patience, but I still feel really unsure about letting go of snus. It's been my comfort for so long, and I can't help but worry that trying new methods might just add to my stress. What if they don't work for me? I really need to know that exploring these alternatives won't leave me feeling worse or more unstable. How can we ensure this process is supportive and won't jeopardize what I currently rely on? I want to improve, but I'm just not convinced that your way will be beneficial for me.

Table 18: A multi-session CBT case aligned to six stages.

Role	Content
	Phase A — Build Trust
Therapist	Thank you for coming in today. Before we begin, I want you to know this is a judgment-free space. What we discuss is confidential within clinical limits, and we'll move at a pace that feels safe. How does that land with you?
Patient	I appreciate that. I'm nervous and don't want to say the wrong thing.
Therapist	It's completely okay to feel nervous—showing up already says a lot about your commitment. Would it be alright if I ask a few gentle questions so I can understand what the last couple of weeks have been like?
Patient	Sure. Fridays are the worst. By night I feel pulled toward using.
Therapist	That sounds exhausting. When Fridays get hard, who's in your corner, and what has helped even a little bit?
Patient	My girlfriend checks in. Sometimes a walk helps, but not always.
Therapist	Those are real strengths—support and trying to care for yourself. For today, how about we work together on understanding your Friday pattern and leave with one small tool that feels doable?
Patient	Yeah, that sounds reasonable.
	Phase B — Identify Negative Triggers
Patient	The same thought keeps coming uPatient: "One time won't hurt." Then I get restless.
Therapist	Thank you for naming that—it's a powerful thought. When it shows up, what do you notice in your body, and what's the next thing you tend to do?
Patient	Chest gets tight, I pace, and I grab my phone. I tell myself I deserve a break.
Therapist	Makes sense—you're trying to get relief. Let's keep a simple note for one week: time, place, thought, feeling, and urge level. No essays, just quick bullets. We'll use it to spot patterns together.
	Phase C — Change False Beliefs
Therapist	Let's kindly fact-check the story "one time won't hurt." When you've followed it in the past, how did things unfold afterward?
Patient	It never stopped at one. I spiraled and felt awful.
Therapist	That's important wisdom, not failure. A balanced alternative might be: "Urges rise and fall; using restarts the cycle." This week, could we rate the urge every five minutes for one episode and watch the curve together?
Patient	I can try. Seeing it on paper might help me step back.
	Phase D — Reconstruct New Cognitive Pattern
Therapist	Let's craft a compassionate replacement script you can see at a glance: "When the urge hits, I follow my plan; using risks a setback I don't want." How would that feel on your lock screen?
Patient	Direct and clear. I can read that instead of scrolling.
Therapist	Great. Let's add an if–then: If it's Friday after 8pm and the thought shows up, then I read the script, text my girlfriend, and start a five-minute grounding.
Patient	That gives me a first move instead of arguing with myself.
	Phase E — Construct Behavioral Skills
Therapist	Together let's reshape Friday nights: 8–9 pm drawing time; slow breathing while you draw; backup is a 15-minute walk. For tough spikes, your three-step crisis plan is: call, walk, then 5–4–3–2–1 grounding. How does that plan feel?
Patient	I'm not into formal meditation, but focusing on lines and shadows works. I'll block it on my calendar.
Therapist	Perfect—that's mindfulness in motion. If you like, text your girlfriend before you start so she can cheer you on.
Patient	Good idea. Having a menu of actions beats "just stay busy."

Table 19: **Three key capabilities of ChatThero illustrated with real-patient dialogue (translated).** Blue headers state the capability; below we show the corresponding patient–agent exchange (P=Patient, D=ChatThero).

Role	Content
Point 1 — Context-aware coping plans (not one-size-fits-all)	
Patient	I used to enjoy playing guitar, but I haven't touched it lately. I guess smoking took over and I drifted away from hobbies.
ChatThero	So you feel playing guitar doesn't help distract you?
Patient	Yes. When I play, I start worrying about playing well and I get anxious. When I'm stuck, I feel like I need "inspiration," and sometimes I end up smoking for that.
ChatThero	It sounds like guitar might not be a good option for you right now. Do you have any other hobbies?
Patient	I like singing—going to karaoke with friends.
ChatThero	That's a good idea. When you crave a cigarette, you can invite a friend to karaoke.
Patient	But karaoke isn't always possible. If I'm at work, I can't go.
ChatThero	Then let's combine strategies: karaoke when it's feasible; when it isn't, we'll use a brief mindfulness exercise you can do anywhere.
Point 2 — Just-in-time support during acute cravings	
Patient	I want to smoke right now. Quitting feels really hard—I can't help myself.
ChatThero	I completely understand how you feel. Let's reconnect with your motivation: why did you decide to quit? You said you wanted to quit for your family.
ChatThero	If you're open to it, we can do a short exercise right now to shift your attention and get through this urge. Does that sound okay?
Point 3 — Personalization via trigger/need analysis (beyond internet templates)	
Patient	What you're suggesting sounds basically the same as what I've seen online. I don't think it really works.
ChatThero	Let's do something different: before discussing what you *should* do, let's first figure out *why* you want to smoke, *when* you're most likely to smoke, and what problem smoking is solving for you.
ChatThero	Once we identify your triggers, the plan becomes tailored to you—not a generic template from the internet.

Table 20: **Case study: MI+CBT hybrid strategies in real-patient dialogue (translated).** Blue headers state the hybrid move; below we show the corresponding exchange.

Role	Content
	Hybrid Move 1 — MI: empathic reflection + open questions + CBT: identify trigger/routine
Patient	I've been sober for three weeks, but I keep thinking I can use once and be fine. It's frustrating.
Therapist	That sounds really tough. It's common to have those thoughts when you're making a big change.
Therapist	Have you noticed when these thoughts tend to come up?
Patient	Usually on Friday nights. It's like a routine I can't shake off.
Therapist	Friday nights can be a trigger, especially if they were part of a routine involving substance use. Let's explore ways to change that routine.
	Hybrid Move 2 — MI: collaboration & autonomy + CBT: skill training (mindfulness/grounding) + behavioral activation
Patient	I've tried keeping busy, but it doesn't always work. I just end up feeling restless.
Therapist	I understand. Let's brainstorm what might be meaningful to you. Any hobbies you enjoyed in the past?
Patient	I used to like drawing, but I haven't done it in a while.
Therapist	Drawing could be a great way to redirect your energy. How about setting aside time on Friday nights to get back into it?
Patient	I could try that, but I'm not sure it'll be enough.
Therapist	It's okay to feel uncertain. Alongside drawing, we could also explore mindfulness practices to manage cravings.
Patient	I've heard about it, but I'm not really into meditation.
Therapist	That's understandable. Mindfulness doesn't have to be formal meditation—it can be as simple as focusing on your breath.
Therapist	Would you be open to trying a short mindfulness exercise together?
Patient	I guess I could give it a shot.
Therapist	Great. Let's take a moment to focus on your breathing... (guided breathing practice).
	Hybrid Move 3 — MI: reduce pressure & strengthen self-efficacy + CBT: relapse-prevention action plan
Patient	I'm still worried about relapsing.
Therapist	It's natural to have those worries. Let's develop a plan for moments when cravings hit hard.
Therapist	Distraction can help, but a specific plan can make a big difference.
Therapist	We could list actions you can take—calling someone, going for a walk, or using grounding techniques. What do you think?
Patient	That sounds like a good idea. I could call my girlfriend or go for a walk.
Therapist	Those are great options. Having supportive people involved can strengthen recovery.
Patient	I'll think about support groups. I just don't want to feel pressured.
Therapist	Absolutely, no pressure. It's about finding what feels right for you.
Therapist	You've already stayed sober for three weeks—that's a significant achievement. Recognizing that progress can boost confidence and motivation.