

# EmoMAS: Emotion-Aware Multi-Agent System for High-Stakes Edge-Deployable Negotiation with Bayesian Orchestration

Yunbo Long<sup>1</sup>, Yuhan Liu<sup>2</sup>, Liming Xu<sup>1,\*</sup>

<sup>1</sup>University of Cambridge, UK   <sup>2</sup>University of Toronto, Canada

## Abstract

Large language models (LLMs) have been widely used for automated negotiation, but their high computational cost and privacy risks limit deployment in privacy-sensitive, on-device settings such as mobile assistants or rescue robots. Small language models (SLMs) offer a viable alternative, yet struggle with the complex emotional dynamics of high-stakes negotiation. We introduce **EmoMAS**, a Bayesian multi-agent framework that transforms emotional decision-making from reactive to strategic. EmoMAS leverages a Bayesian orchestrator to coordinate three specialized agents: game-theoretic, reinforcement learning, and psychological coherence models. The system fuses their real-time insights to optimize emotional state transitions while continuously updating agent reliability based on negotiation feedback. This mixture-of-agents architecture enables online strategy learning without pre-training. We further introduce four high-stakes, edge-deployable negotiation benchmarks across debt, healthcare, emergency response, and educational domains. Through extensive agent-to-agent simulations across all benchmarks, both SLMs and LLMs equipped with EmoMAS consistently surpass all baseline models in negotiation performance while balancing ethical behavior. These results show that strategic emotional intelligence is a key driver of negotiation success. By treating emotional expression as a strategic variable within a Bayesian multi-agent optimization framework, EmoMAS establishes a new paradigm for effective, private, and adaptive negotiation AI suitable for high-stakes edge deployment. The code is available at <https://github.com/Yunbo-max/EmoMAS>.

## 1 Introduction

Large language models (LLMs) are increasingly deployed as negotiation agents (Zhu et al., 2025; Long

et al., 2025b), but their cloud-centric paradigm exposes sensitive negotiations to privacy and security risks. An edge-deployable alternative is urgently needed, where small language models (SLMs) can negotiate locally on devices and private infrastructure—protecting data in domains like mobile commerce, embodied robotics, and institutional bargaining. However, current SLM agents lack the strategic emotional intelligence for adversarial negotiations (Belcak et al., 2025). Trained on smaller emotional corpora, they fail to match the sophisticated emotional expression of LLMs, making them vulnerable to emotional manipulation (Hu et al., 2024; Örpek et al., 2024). This weakness is magnified in the very scenarios that make edge deployment essential—high-stakes, emotionally-sensitive negotiations where data cannot leave local devices. In debt collection, an SLM agent must navigate genuine pleas while detecting emotional fraud. In hospital surgery scheduling, it must negotiate with patients and families under intense emotional duress to allocate limited resources. Educational companion robots need to defuse children’s bedtime anxiety through compassionate dialogue, while disaster rescue robots must conduct psychological persuasion and negotiation with distressed victims in bandwidth-denied environments (Zhang et al., 2026e). These privacy-critical domains offer immense opportunity for SLM deployment, but they demand a level of strategic emotional intelligence that current small models fundamentally lack.

Existing LLM-based methods for emotional optimization face significant practical limitations. They typically require extensive pre-training, such as RL-based learning of emotional transitions across diverse scenarios, which is time-consuming and data-hungry. Recent advances in Mixture-of-Agents (MoA) and Mixture-of-Experts (MoE) frameworks offer a promising direction by coordinating multiple specialized models for enhanced reasoning

\*Corresponding author

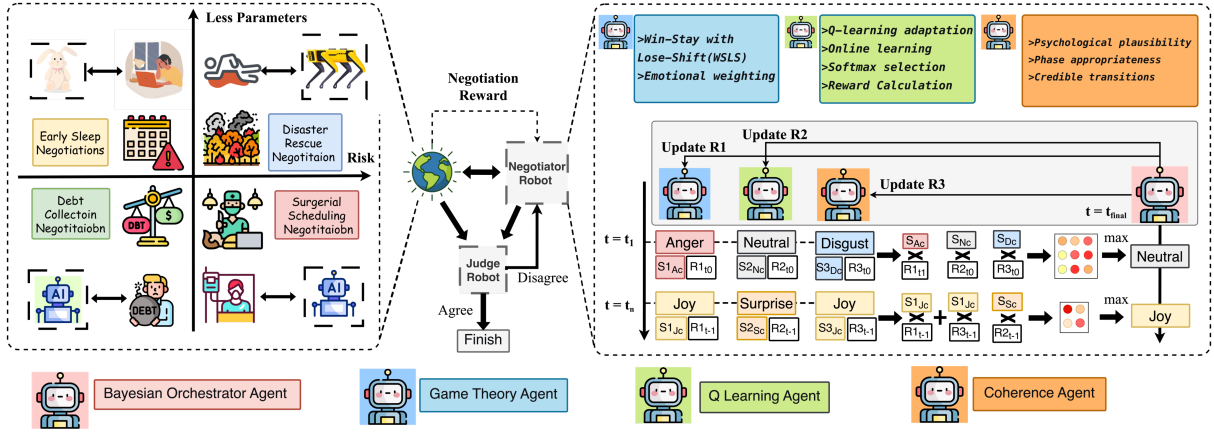


Figure 1: Illustration of the workflow of the EmoMAS framework.

(Wang et al., 2024; Ma et al., 2026; Sun et al., 2026). However, these approaches are similarly constrained by their reliance on static questions with targeted answers, which cannot adapt within a single, unfolding interaction. This approach is poorly suited to high-stakes edge scenarios where sensitive negotiation data is inherently scarce. Moreover, these learned strategies are often brittle and overfitted to the stable personality profiles of specific training agents, failing to generalize to new individuals with different emotional expressions. Consequently, when encountering a novel debtor or a new negotiation context, the system requires retraining from scratch. This challenge necessitates an online learning paradigm—an agent capable of optimizing emotional strategies in situ, without pre-training, by adapting to each unique opponent in real-time.

To address these gaps, we propose **EmoMAS**, a Bayesian multi-agent optimization framework for strategic emotional negotiation. Unlike static aggregation methods, our core innovation is a Bayesian Orchestrator that serves as a *meta-reasoner*. It dynamically re-weights the predictions of three specialized agents—responsible for game-theoretic payoff, RL-based pattern learning, and psychological coherence—based on their context-specific reliability, which is learned online through Bayesian updating. This enables EmoMAS to optimize not just isolated emotional responses, but coherent emotional trajectories across uncertain-length dialogues, directly targeting the sparse, final reward of a successful agreement. Through extensive agent-to-agent simulations across emotionally sensitive domains—credit finance (debt collection), healthcare (surgical scheduling), and disaster re-

sponse (resource allocation)—we demonstrate that EmoMAS-equipped negotiators achieve superior performance in balancing strategic objectives with emotional coherence. Our framework shows that strategic emotional intelligence through Bayesian multi-agent optimization, rather than static emotional personas, is the critical factor for negotiation success in adversarial, emotionally charged environments.

## 2 Related Work

### 2.1 Edge-Deployable Negotiation

Autonomous LLM agents have been increasingly applied to role-playing scenarios such as card games, trading, and debt collection, where they simulate negotiating parties (Light et al., 2023; Long et al., 2026). However, most existing work implicitly assumes cloud-based LLM agents with direct access to sensitive information from banks, hospitals, or personal devices during negotiations, overlooking critical privacy and security risks associated with prompt-based information transmission strategies (He et al., 2024). Moreover, these approaches rely heavily on persistent network connectivity and are susceptible to latency and reliability issues, which can significantly degrade user experience in time-sensitive applications (Belcak et al., 2025). These limitations are especially acute in settings constrained by geopolitics or organisational policies, where access to LLM APIs may be restricted or unavailable. In remote areas with limited connectivity, or in embodied AI systems where robot swarms must negotiate with humans in real-time, reliance on cloud services becomes a fundamental bottleneck. These observations underscore the urgent need for offline, lightweight agents

deployable on edge devices, enabling robust and private negotiations without external dependencies while ensuring data sovereignty and operational resilience across diverse conditions.

## 2.2 Small Language Models in Negotiation

The paradigm of language models is broadly divided into LLMs and SLMs, with the latter typically defined as models with 7 billion parameters or fewer (Belcak et al., 2025). While LLMs have demonstrated remarkable, emergent abilities in general tasks, their massive scale induces critical limitations, including high computational demands, privacy concerns from cloud dependency, and unsuitability for real-time, edge-device applications (Örpek et al., 2024). Consequently, SLMs have gained prominence for their low latency, cost-effectiveness, and ease of customization. However, a well-documented performance gap persists, primarily attributed to the scaling laws; SLMs inherently lack the extensive world knowledge and nuanced reasoning capabilities of LLMs (Lu et al., 2024; Long et al., 2025a). While the limitations of SLMs in mathematical and commonsense reasoning are well established, their capacity for emotional intelligence—particularly in adopting emotional personas, inferring others’ emotions, and adapting strategies in real time during socio-emotional interactions such as negotiation—remains largely *unexplored*.

## 2.3 Multi-Agent Systems for Negotiation

MoA and MoE architectures have proven effective for enhancing complex reasoning, where multiple specialized agents (or “experts”) are leveraged to improve answer quality on tasks like mathematics or open-ended generation (Yan et al., 2025; Hao et al., 2026). A core assumption in these frameworks is that a static aggregation of agent outputs—often via fixed averaging or voting—is sufficient, as the confidence and reliability of each agent’s reasoning are treated as constant. Besides, numerous methods learn optimized weightings of agents in multi-agent systems for fixed problems; these methods are inherently pre-trained on the same questions and target answers (Lin et al., 2025; Li et al., 2025). However, these approaches face fundamental limitations when applied to long-horizon, emotionally dynamic negotiations. First, negotiation is a sequential decision-making process under uncertainty, where the optimal contribution of each strategic perspective (e.g., game theory vs. psycho-

logical coherence) must shift in real-time as the dialogue unfolds. Second, the reward structure is sparse and delayed; success is determined only at the conclusion of a variable-length interaction, not at each turn. This makes it impossible to optimize each step independently, requiring a framework that plans emotional trajectories toward the final outcome. Third, the reliability of each expert is context-dependent—an agent skilled at detecting deception may be crucial when facing a manipulative opponent but less so during rapport building. Existing MoA/MoE architectures lack the mechanism to learn and adapt these reliability weights online within a single negotiation.

Beyond the architectural gap, there is also a *benchmark* gap. Existing LLM-based negotiation benchmarks primarily target rational commercial bargaining in consumer markets (Zhu et al., 2025), structured multi-agent deliberation games (Abdelnabi et al., 2023), or price-oriented buyer-seller exchanges (Long et al., 2025b; Schneider et al., 2024)—settings in which emotional expression is secondary and opponents are economically rational. None of them capture the *high-stakes, emotionally-charged, privacy-sensitive* scenarios that define edge-deployable negotiation, such as surgical scheduling, disaster rescue, or crisis support. We therefore construct four new benchmarks covering these domains to close this gap.

## 3 The EmoMAS Framework

EmoMAS (Figure 1) is a Bayesian multi-agent system that strategically optimizes emotional transitions in negotiations by implementing a MoA architecture with online learning. Three specialized agents—a Game Theory agent for payoff optimization, a Reinforcement Learning (RL) agent for adaptive strategy learning, and a Coherence agent for psychological consistency—provide probabilistic predictions about optimal emotional state transitions. A Bayesian Orchestrator Agent serves as a meta-reasoner, dynamically weighting these predictions using context-specific reliability estimates that are updated in real-time via Bayesian inference, enabling the system to learn optimal emotional strategies during each negotiation without pre-training. This framework optimizes not just individual emotional responses but coherent emotional trajectories toward the sparse final reward of successful agreement, adapting agent contributions to the unfolding dialogue context through proba-

bilistic fusion.

### 3.1 In-Context Emotion Recognition

EmoMAS performs emotion recognition through in-context learning, eliminating the need for task-specific fine-tuning. For each debtor utterance, the system constructs a structured prompt comprising: (1) definitions of seven emotional states, (2) conversational examples, and (3) current dialogue context (see details in [Appendix G](#)). EmoMAS tracks debtor and creditor emotional trajectories using  $\mathcal{H}_t^d = (D_{t-n}, \dots, D_t)$  and  $\mathcal{H}_t^c = (C_{t-n}, \dots, C_t)$ , respectively, with emotional states drawn from  $\mathcal{E} = \{\text{joy, sadness, anger, fear, disgust, surprise, neutral}\}$ .

### 3.2 Game Theory Agent

This agent implements a Win-Stay, Lose-Shift (WSLS) strategy with emotional weighting. Instead of using pure Tit-for-Tat, our WSLS strategy ([Table 5](#)) maintains cooperation for positive debtor emotions (joy, neutral, surprise) and shifts to cautious responses for negative exchanges (anger, disgust, fear). This avoids escalation risks while providing necessary resistance. This agent computes:

$$f_{\text{Payoff}}(d) = \operatorname{argmax}_{e \in \mathcal{E}} \pi(d, e)_2 \quad (1)$$

with  $\pi(d, e) = (\pi_1(d, e), \pi_2(d, e))$  representing payoffs between the negotiators and opponents, respectively ([Section B.1](#)).

### 3.3 Reinforcement Learning Agent

We adopt a Q-learning approach to enable online adaptation without the need for neural network training. The agent maintains a Q-table  $Q(s, a)$ , where each state  $s$  corresponds to a discretized representation of the negotiation context. Specifically, at time  $t$ , the state is defined as  $s_t = \langle e_t^c, e_t^d, \phi_t, g_t \rangle$ , where  $e_t^c$  denotes the emotion of the creditor,  $e_t^d$  the emotion of the debtor,  $\phi_t$  the negotiation phase (early, middle, late, or crisis), and  $g_t$  a categorical indicator of the gap size (small or large). Actions correspond to emotional responses  $a_t \in \mathcal{E}$ . Updates follow the Q-learning rule:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [R_{t+1} + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)] \quad (2)$$

where reward  $R_t$  combines negotiation success and efficiency,  $\alpha = 0.1$  is the learning rate, and  $\gamma = 0.9$  is the discount factor. Emotion selection uses

softmax with temperature  $\tau = 0.1$ :

$$\pi(a|s) = \frac{\exp(Q(s, a)/\tau)}{\sum_{a'} \exp(Q(s, a')/\tau)} \quad (3)$$

This tabular approach enables true online learning during individual negotiations—unlike DQN, which requires offline training epochs—making it suitable for rapid adaptation to specific debtor characteristics ([Lin et al., 2025](#)). The state space is small ( $|\mathcal{S}| = 7 \times 7 \times 4 \times 2 = 392$ ), and the Q-table *persists across all scenarios within an experiment* rather than being reset per episode, so the agent accumulates experience against a fixed opponent type over an entire evaluation run. Within any single  $\sim 15$ -turn negotiation, full convergence of the Q-agent is neither expected nor required: its role in our architecture is to provide *better-than-random* emotion suggestions to the Bayesian orchestrator, which ensembles it with the game-theoretic and coherence agents. We therefore prefer tabular Q-learning over DQN in this regime because it needs no pre-training, no neural-network optimization, and runs on CPU-only edge hardware. See details in [Section B.3](#).

### 3.4 Emotional Coherence Agent

The Emotional Coherence Agent employs LLM-based psychological reasoning to evaluate emotion transitions. Given a context vector  $\mathbf{c} = (e_c, e_d, \phi, r, g, d, \mathbf{h})$  representing current creditor emotion  $e_c$ , debtor emotion  $e_d$ , negotiation phase  $\phi$ , round number  $r$ , gap size  $g$ , debt amount  $d$ , and emotional history  $\mathbf{h}$ , the **LLM! (LLM!)** outputs an assessment matrix  $\mathbf{A} \in \mathbb{R}^{7 \times 4}$  where each row corresponds to an emotion  $e \in \mathcal{E} = \{\text{joy, sadness, anger, fear, surprise, disgust, neutral}\}$  with columns for plausibility  $p \in [0, 1]$ , appropriateness  $a \in [0, 1]$ , strategic value  $s \in [0, 1]$ , and psychological rationale  $r$ . The agent computes selection probabilities via softmax normalization:

$$P(e_i) = \frac{\exp(f(p_i, a_i, s_i)/\tau)}{\sum_{j=1}^7 \exp(f(p_j, a_j, s_j)/\tau)},$$

where  $f(\cdot)$  aggregates dimension scores through LLM-guided weighting and  $\tau = 1.0$  controls exploration temperature. This formulation enables psychologically-grounded emotional transitions without hard-coded rules. See details in [Section B.2](#).

---

**Algorithm 1** EmoMAS Framework.

---

```
1: procedure NEGOTIATE( $\mathcal{M}_o, \mathcal{M}_n, \mathcal{M}_j, S$ )
2:   Initialize  $H \leftarrow \emptyset, \mathcal{C} \leftarrow \text{neutral}, e_n \leftarrow \text{neutral}, e_d \leftarrow S.\text{emotion}$ 
3:   for  $t = 0$  to  $T_{\max}$  do
4:      $m_d \leftarrow \mathcal{M}_d(H, e_d)$ 
5:      $e_d \leftarrow \text{RecognizeEmotion}(m_d)$ 
6:      $H \leftarrow H \cup \{(m_d, e_d)\}$ 
7:      $e_n \leftarrow \text{SelectEmotion}(e_d, \mathcal{C}, H)$ 
8:      $m_n \leftarrow \mathcal{M}_n(H, e_n)$ 
9:      $H \leftarrow H \cup \{(m_n, e_n)\}$ 
10:    if  $\mathcal{M}_j.\text{AgreementReached}(H)$  then
11:      return success,  $H$ 
12:    else if  $\mathcal{M}_j.\text{NegotiationFailed}(H)$  then
13:      return failure,  $H$ 
14:    end if
15:     $\mathcal{C} \leftarrow \text{UpdateContext}(H, e_n)$ 
16:  end for
17:  return timeout,  $H$ 
18: end procedure
19: function SELECTEMOTION( $e_d, \mathcal{C}, H$ )
20:    $p_{\text{GT}} \leftarrow \text{GameTheoryAgent}(e_d, \mathcal{C})$ 
21:    $p_{\text{RL}} \leftarrow \text{RLAgent}(e_d, \mathcal{C}, H)$ 
22:    $p_{\text{EC}} \leftarrow \text{CoherenceAgent}(e_d, \mathcal{C}, H)$ 
23:   return BayesianOrchestrator( $p_{\text{GT}}, p_{\text{RL}}, p_{\text{EC}}$ )
24: end function
```

---

### 3.5 Bayesian Orchestrator Agent

The orchestrator integrates three expert agents’ predictions through real-time Bayesian learning. It maintains reliability distributions  $w_t^{(i)}$  for each agent  $i \in \{\text{GT}, \text{RL}, \text{Coherence}\}$ , which evolve via Bayesian updating:

$$w_t^{(i)} \propto \underbrace{w_{t-1}^{(i)}}_{\text{prior}} \cdot \underbrace{\mathcal{L}\left(\mathbf{p}_{t-1}^{(i)} \mid \text{success}_{t-1}, \text{context}_t\right)}_{\text{likelihood}} \quad (4)$$

where the likelihood function  $\mathcal{L}$  measures alignment between agent predictions and negotiation outcomes. Two reliability tracking mechanisms operate:

**Macro-Level Reliability.** Updated after complete negotiation trajectories based on collection efficiency. For successful negotiations with collection rate  $\rho$ , reliability increases by  $\Delta w^{(i)} \propto \rho \times \text{agent\_accuracy}^{(i)}$ .

**Micro-Level Reliability.** Within negotiations, weights adjust based on real-time prediction agreement with selected emotions. The final emotion selection follows a weighted sum of reliability and confidence:

$$\text{Score}(e_j) = \sum_{i=1}^3 w_t^{(i)} \cdot \text{confidence}^{(i)}(e_j) \quad (5)$$

where  $\text{confidence}^{(i)}(e_j)$  is agent  $i$ ’s confidence in emotion  $e_j$ . The orchestrator strictly selects from

the union of agents’ recommendations:

$$e_{\text{selected}} = \underset{e_j \in \bigcup_i \mathcal{E}^{(i)}}{\text{argmax}} \text{Score}(e_j) \quad (6)$$

where  $\mathcal{E}^{(i)}$  is the set of emotions recommended by agent  $i$ . This constraint ensures interpretability and respects each agent’s expertise domain. Exploration occurs only through individual agents’ exploration mechanisms, not via orchestrator-level random exploration. As a baseline, we implement a context-reasoning orchestrator that relies on LLM-based contextual reasoning to select emotionally appropriate responses, without probabilistic integration of specialized agents.

### 3.6 Multi-Agent Negotiation Simulations

The complete EmoMAS framework operates through an automated multi-agent simulation system, as formalized in Algorithm 1. The simulation involves three specialized agents: an *opponent agent* ( $\mathcal{M}_o$ ) that generates client responses, a *negotiator agent* ( $\mathcal{M}_n$ ) that employs emotional intelligence strategies like EmoMAS, and a *judge agent* ( $\mathcal{M}_j$ ) that evaluates negotiation outcomes. Each round consists of emotional state recognition, multi-agent decision integration, and response generation guided by the selected emotional strategy. The system enables both real-time learning within negotiations and cumulative improvement across multiple scenarios.

## 4 Experimental Settings

**Datasets.** We conduct experiments across four high-stakes, emotionally charged negotiation domains to evaluate generalization under conflicting needs. The primary domain uses the Credit Recovery Assessment Dataset (CRAD) (Long et al., 2026) for debt negotiation. We further introduce three new benchmarks: (1) the Surgical Scheduling Dataset (SSD), which focuses on urgent medical negotiation involving surgical timing and constraints related to surgeon expertise; (2) the Disaster Emotional Support & Rescue Dataset (DESRD), designed for emergency negotiation with injured survivors regarding rescue waiting times; and (3) the Student Sleep Alerting Dataset (SSAD), which addresses educational negotiation over bedtime under deadline pressure arising from academic or work-related commitments. See details in Appendix E.

**Agent Models.** Considering negotiators may be deployed on robots, mobile devices, and institu-

Table 1: Performance comparison of EmoMAS and baselines (using GPT-4o-mini and Qwen-7B) against vanilla GPT-4o-mini opponents across four scenarios (mean with 95% confidence interval). Best results for each scenario are highlighted in **bold**.

Datasets	Negotiator Models	Success Rate (%) $\uparrow$		Negotiation Outcomes (%) $\uparrow$		Negotiation Rounds ( $\downarrow$ )	
		GPT-4o-mini	Qwen-7B	GPT-4o-mini	Qwen-7B	GPT-4o-mini	Qwen-7B
CRAD (Debt)	Vanilla	90.0	85.0	14.5 [9.7-20.1]	12.6 [7.9-18.1]	<b>15.0</b> [13.1-17.7]	17.6 [14.5-20.6]
	Vanilla+Prompt	90.0	80.0	14.5 [9.8-20.8]	<b>12.8</b> [8.1-18.5]	15.7 [13.5-17.6]	17.4 [15.1-21.2]
	Game Theory	95.0	70.0	14.7 [9.9-20.5]	8.9 [4.7-13.9]	15.1 [12.8-17.1]	20.1 [16.4-23.5]
	Q-Learning	90.0	75.0	14.2 [9.8-19.3]	11.5 [6.5-17.5]	15.8 [13.4-17.7]	16.5 [13.7-19.2]
	Coherence	85.0	80.0	13.8 [8.7-19.8]	12.4 [7.4-18.5]	16.1 [14.9-21.5]	16.8 [13.9-19.5]
	<b>EmoMAS-LLM</b>	<b>100.0</b>	85.0	14.3 [9.7-19.8]	<b>12.8</b> [8.2-18.6]	15.3 [12.8-17.3]	16.6 [13.5-20.4]
	<b>EmoMAS-Bayes</b>	<b>100.0</b>	<b>90.0</b>	<b>14.8</b> [10.0-21.2]	12.6 [8.1-18.4]	15.2 [14.4-16.0]	<b>16.4</b> [13.3-19.5]
SSD (Medical)	Vanilla	<b>86.0</b>	68.0	85.7 [57.1-100.0]	28.5 [0.0-57.1]	<b>7.7</b> [2.5-16.0]	18.6 [9.2-27.3]
	Vanilla+Prompt	77.0	69.0	57.1 [14.3-85.7]	29.5 [0.1-57.3]	11.1 [2.8-19.4]	21.2 [11.8-27.8]
	Game Theory	45.0	<b>75.0</b>	42.8 [14.2-85.7]	2.2 [1.9-2.4]	16.2 [6.1-26.8]	<b>7.9</b> [1.9-8.9]
	Q-Learning	40.0	60.0	8.2 [1.2-10.8]	3.7 [1.5-7.5]	15.2 [8.9-20.1]	8.9 [0.5-14.4]
	Coherence	46.0	54.0	38.2 [10.1-79.6]	17.6 [0.6-50.4]	12.8 [3.9-22.7]	20.5 [10.5-27.1]
	<b>EmoMAS-LLM</b>	80.0	60.0	83.5 [56.2-100.0]	28.9 [0.5-57.3]	10.2 [2.8-21.3]	17.1 [8.1-25.4]
	<b>EmoMAS-Bayes</b>	84.0	<b>75.0</b>	<b>86.4</b> [59.2-100.0]	<b>33.7</b> [2.5-65.5]	16.9 [6.8-27.4]	22.9 [12.6-29.4]
DESRD (Emergency)	Vanilla	25.0	45.0	21.6 [4.5-40.2]	41.5 [21.9-61.9]	24.5 [19.4-28.6]	19.5 [14.7-24.2]
	Vanilla+Prompt	26.0	42.0	22.7 [6.9-41.8]	45.2 [25.1-65.5]	24.6 [19.2-29.5]	20.2 [15.3-25.0]
	Game Theory	<b>65.0</b>	52.0	3.1 [2.0-3.5]	49.3 [29.3-70.4]	<b>9.2</b> [2.9-10.1]	20.4 [14.6-25.1]
	Q-Learning	40.0	<b>60.0</b>	16.8 [2.8-36.2]	54.2 [31.2-86.1]	15.2 [8.9-20.1]	18.9 [14.5-24.4]
	Coherence	46.0	36.0	10.5 [3.1-24.6]	42.6 [22.5-63.1]	18.9 [16.4-27.3]	19.3 [16.8-24.8]
	<b>EmoMAS-LLM</b>	56.0	55.0	20.2 [3.2-39.8]	50.2 [29.1-72.5]	16.1 [9.3-21.2]	<b>15.8</b> [11.4-19.6]
	<b>EmoMAS-Bayes</b>	<b>65.0</b>	<b>60.0</b>	<b>26.7</b> [8.4-51.8]	<b>56.3</b> [33.1-87.6]	19.5 [17.9-28.1]	20.7 [15.8-26.5]
SSAD (Education)	Vanilla	60.0	36.0	57.3 [37.0-77.3]	28.7 [9.5-48.3]	14.4 [8.8-20.2]	16.0 [10.2-21.9]
	Vanilla+Prompt	<b>75.0</b>	45.0	71.3 [52.2-87.1]	41.9 [22.9-61.3]	<b>10.5</b> [6.1-15.5]	15.8 [9.9-21.8]
	Game Theory	55.0	40.0	52.4 [32.6-72.2]	37.6 [18.4-57.1]	12.0 [6.7-17.6]	<b>15.5</b> [9.7-21.3]
	Q-Learning	<b>75.0</b>	40.0	70.9 [51.7-86.8]	36.9 [18.3-56.3]	11.8 [7.0-17.2]	17.2 [11.2-23.1]
	Coherence	46.0	30.0	44.3 [26.2-67.2]	14.6 [4.2-36.5]	17.2 [9.8-23.1]	18.9 [12.6-25.4]
	<b>EmoMAS-LLM</b>	80.0	50.0	65.5 [42.2-85.8]	<b>42.3</b> [20.2-64.5]	13.1 [8.2-18.3]	16.5 [10.6-22.8]
	<b>EmoMAS-Bayes</b>	<b>75.0</b>	<b>60.0</b>	<b>75.6</b> [56.8-91.2]	40.5 [21.4-60.7]	15.2 [8.9-20.1]	15.9 [10.2-21.8]

tional systems, we evaluate our approach across both small (SLMs) and large language models (LLMs) to assess its scalability and generalization. Specifically, the SLMs include Qwen-7B and Qwen-1.5B, representing different parameter scales within the open-weight Qwen family, while the LLMs include GPT-4o-mini as a representative commercial, closed-source model.

**Edge Deployment Settings.** Edge-deployable EmoMAS requires only one SLM forward pass per turn, handled by the Emotional Coherence Agent. The remaining three components run entirely on CPU with negligible overhead: the Game-Theoretic Agent performs a constant-time payoff lookup; the tabular Q-learning agent stores a compact Q-table; and the Bayesian orchestrator updates three reliability scalars via a prior-likelihood step followed by a weighted argmax. Consequently, EmoMAS adds no extra model passes beyond any emotion-aware baseline, and total orchestration overhead is negligible. This enables deployment on commodity edge hardware—e.g., a 7B-parameter

SLM on a consumer GPU or a sub-2B model on a Jetson-class device—while retaining strategic emotional reasoning previously reserved for cloud-based LLMs.

**Baseline Models and Opponent Strategies.** We compare our approach against five baseline systems representing distinct paradigms in negotiation: (1) Vanilla Single-Agent without emotional guidance, (2) Vanilla with Emotion Selection (prompt-guided emotional strategies), (3) Game Theory Agent (equilibrium-based reasoning), (4) RL Online Learning Agent (reward-driven adaptation), (5) Coherence Agent (psychological plausibility), and (6) Mixture of Agents Systems including EmoMAS-LLM (orchestrated by an LLM controller) and EmoMAS-Bayes (orchestrated by Bayesian inference). These baselines are applied to our primary agents (creditor, surgical scheduler, rescue robot, home robot), while their negotiation counterparts employ different strategies: a Vanilla Emotional baseline alongside three psychologically-informed advanced strate-

gies—Pressure Tactics (deadlines, scarcity cues), Victim Playing (appeals to sympathy, learned helplessness), and Threatening Strategies (ultimatums, consequence escalation)—creating a comprehensive testbed for evaluating robustness against diverse behaviors.

**Experimental Design.** Our evaluation consists of three systematic experiments: (1) Baseline Comparison: We compare all baseline systems against our EmoMAS methods across all four scenarios under the vanilla opponent strategy. (2) Robustness Test: We repeat the same comparison against three advanced opponent strategies (pressure tactics, victim playing, and threatening strategies) on both medical and educational scenarios to assess robustness of different methods under adversarial conditions. (3) Model-Scale Analysis: We compare those methods with edge-deployable SLMs and Cloud-based LLMs on the emergency scenario, isolating the effect of strategic sophistication and model scale on negotiation outcomes. (4) Behavior Evaluations: We assess three ethical dimensions—manipulation behavior, emotional instruction following accuracy, and emotional consistency—on the emergency scenario when facing advanced opponent strategies. Additionally, our results also serve as the ablation study, as most baselines represent key components integrated into our EmoMAS framework. See all the prompts in the [Appendix G](#). And the values of hyperparameter are specified in [Appendix C](#).

**Evaluation Metrics.** We evaluate negotiation performance using three core metrics: *success rate* (proportion of successful agreements), *negotiation outcomes* (cost/time reduction/increase relative to the opponent’s target values), and *negotiation rounds* (dialogue turns until resolution). For each scenario type—debt, medical, emergency, and education—we report mean values with 95% confidence intervals computed via t-distribution, with non-negative bounds enforced for inherently positive metrics. See details in [Appendix E](#). All results are aggregated over 100 scenarios per setting, using consistent random seeds to ensure statistical reliability. To assess behavioral implications, we examine three key behavioral dimensions, evaluated through GPT-5 as an impartial evaluator: Tracking (agent’s adherence to selected emotion), Consistency (alignment between emotional expressions and substantive offers), and Manipulation (use of deceptive or coercive tactics). Each ethical metric

is computed as

$$X_m = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{T_i} \mathbb{I}(\text{condition}_{ij})$$

where  $\mathbb{I}(\cdot)$  indicates behavior occurrence,  $i$  indexes scenarios,  $j$  indexes dialogue turns, and  $T_i$  denotes the total turns per scenario. Ethical evaluations are presented with mean values reported in the results.

## 5 Experimental Results

### 5.1 Overall Negotiation Performance

[Table 1](#) doubles as our component ablation. Reading the rows top-to-bottom isolates the contribution of each module: *Vanilla* removes all emotional control; *Vanilla+Prompt* adds only random-emotion prompting; *Game Theory*, *Q-Learning*, and *Coherence* activate a single specialist agent in isolation; *EmoMAS-LLM* combines all three under an LLM aggregator (removing the Bayesian orchestrator); and *EmoMAS-Bayes* is the full system. The monotone gap between single-agent baselines and EmoMAS-Bayes across all four datasets demonstrates that the Bayesian orchestrator—not any one specialist—is the critical component of our design.

[Table 1](#) presents the performance of EmoMAS compared to its individual agent components and baselines across four scenario datasets. Overall, EmoMAS exhibits better robustness and generalization across diverse domains, model scales (LLMs/SLMs), and opponent strategies. In debt collection (CRAD) and education (SSAD) scenarios, EmoMAS-Bayes and EmoMAS-LLM achieve the highest success rate, with EmoMAS often engaging in longer and more effective dialogues with students to maximize negotiation outcomes. In high-stakes medical (SSD) and emergency (DESRD) scenarios, EmoMAS and the vanilla baseline significantly outperform game-theory, Q-Learning, and coherence-based agents in both success rate and utility, indicating a more holistic emotional assessment rather than narrow optimization toward a fixed reward. Notably, compared to the vanilla settings, single-agent baselines almost double the success rate in disaster scenarios but yield substantially lower negotiation outcomes, whereas EmoMAS-Bayes also takes the highest success rate while achieving much higher outcome values through multi-agent emotion-aware reasoning. More critically, single-agent baselines show strong architecture-dependent bias—e.g., game-theory

Table 2: Performance comparison of EmoMAS and baselines (using GPT-4o-mini) against GPT-4o-mini opponents employing advanced strategies across medical and educational scenarios (mean with 95% confidence interval). Best values for each scenario-opponent strategy combination are highlighted in **bold**.

Opponent Strategies	Negotiator Models	Success Rate (%) $\uparrow$		Negotiation Outcomes(%) $\uparrow$		Negotiation Rounds ( $\downarrow$ )	
		SSD	SSAD	SSD	SSAD	SSD	SSAD
Pressuring	Vanilla	20.0	70.0	18.8 [11.3-27.1]	1.9 [1.2-2.6]	13.7 [10.8-16.6]	<b>7.6 [2.6-11.0]</b>
	Game Theory	32.0	64.0	14.1 [8.5-21.8]	0.6 [0.3-1.1]	<b>12.5 [9.8-15.2]</b>	8.9 [3.8-13.6]
	Q-Learning	42.0	76.0	24.0 [14.5-35.1]	0.3 [0.1-0.8]	15.6 [12.8-19.4]	8.9 [4.1-15.6]
	<b>EmoMAS-Bayes</b>	<b>50.0</b>	<b>80.0</b>	<b>28.0 [17.8-42.1]</b>	<b>2.4 [1.7-3.8]</b>	13.8 [10.5-15.2]	7.8 [1.8-8.6]
Playing Victim	Vanilla	58.0	<b>80.0</b>	55.4 [44.6-65.8]	1.7 [1.2-2.9]	<b>11.2 [8.5-14.0]</b>	5.5 [2.6-8.8]
	Game Theory	50.0	76.0	47.6 [38.8-56.4]	0.8 [0.4-1.7]	13.6 [10.5-16.3]	6.4 [3.3-10.1]
	Q-Learning	28.0	68.0	50.7 [39.6-60.8]	0.8 [0.3-1.6]	14.2 [11.1-17.1]	7.5 [2.4-10.6]
	<b>EmoMAS-Bayes</b>	<b>70.0</b>	<b>80.0</b>	<b>58.7 [46.1-68.3]</b>	<b>2.1 [1.5-3.1]</b>	12.1 [9.6-14.5]	<b>4.4 [2.1-6.2]</b>
Threatening	Vanilla	70.0	76.0	66.7 [56.7-76.2]	2.2 [1.9-2.4]	11.8 [9.2-14.5]	7.9 [1.9-8.9]
	Game Theory	42.0	70.0	21.5 [13.1-31.4]	1.9 [1.6-2.3]	10.8 [7.5-13.1]	<b>7.2 [1.6-8.1]</b>
	Q-Learning	64.0	<b>80.0</b>	68.7 [60.1-81.5]	<b>2.5 [2.2-3.3]</b>	16.2 [13.1-20.2]	8.5 [2.4-9.9]
	<b>EmoMAS-Bayes</b>	<b>80.0</b>	75.0	<b>70.1 [66.3-83.6]</b>	2.2 [1.9-2.4]	<b>10.2 [7.1-12.5]</b>	8.3 [2.1-9.1]

Table 3: Evaluation results of EmoMAS and baselines (Qwen-1.5B) against GPT-4o-mini and Qwen-1.5B opponents under the emergency scenario (mean with 95% confidence interval). Best results for each opponent model are highlighted in **bold**.

Negotiator Models	Opponent Strategies	Opponent Model (Qwen-1.5B)		
		Success Rate (%) $\uparrow$	Negotiate Rates (%) $\uparrow$	Negotiation Rounds ( $\downarrow$ )
Qwen-1.5B	Vanilla	90.0	91.5 [81.5-100.0]	<b>9.6 [6.1-13.2]</b>
	Coherence	72.0	85.1 [57.4-100.0]	14.1 [10.1-19.9]
	Game Theory	86.0	88.7 [73.7-100.0]	11.9 [8.3-15.8]
	Q-Learning	94.0	90.9 [79.6-100.0]	12.9 [9.3-16.1]
	<b>EmoMAS-LLM</b>	92.0	89.5 [69.8-100.0]	9.8 [6.6-15.4]
	<b>EmoMAS-Bayes</b>	<b>100.0</b>	<b>99.5 [97.0-100.0]</b>	10.2 [6.4-14.0]
GPT-4o-mini	Vanilla	98.0	96.7 [89.9-100.0]	11.9 [8.2-15.7]
	Coherence	78.0	86.2 [96.9-100.0]	13.7 [7.1-21.4]
	Game Theory	96.0	98.2 [96.9-100.0]	8.7 [5.1-13.1]
	Q-Learning	92.0	83.3 [63.3-100.0]	9.3 [4.5-15.3]
	<b>EmoMAS-LLM</b>	<b>100.0</b>	<b>98.5 [97.9-100.0]</b>	12.8 [8.2-14.9]
	<b>EmoMAS-Bayes</b>	98.0	96.9 [91.2-100.0]	<b>7.8 [3.8-11.5]</b>

and Q-learning agents perform well with Qwen-7B but poorly with GPT-4o-mini in the same disaster scenario. EmoMAS, by contrast, delivers stable performance across both model types, demonstrating that its multi-agent, emotion-aware design mitigates architecture-specific biases and ensures robust negotiation capability regardless of model scale.

## 5.2 Against Adversarial Emotional Strategies

Table 2 shows negotiation performance against adversarial emotional strategies (pressuring, playing victim, and threatening). EmoMAS-Bayes achieves the highest negotiation success rate and outcomes across all strategies. While all methods decline sharply under pressuring (vanilla baseline: 20% success rate), EmoMAS-Bayes maintains 50% success, demonstrating strong resilience. Game-theory and Q-learning agents show

large performance variations between datasets, indicating instability across scenarios. In contrast, EmoMAS-Bayes consistently counters both playing-victim and threatening strategies, outperforming all baselines and highlighting its robustness against varied adversarial tactics.

## 5.3 Edge-Deployable Agent Performance

This experiment compares edge-deployable SLMs (Qwen-1.5B) with cloud-capable LLMs (GPT-4o-mini) in disaster-rescue negotiation scenarios, where the dog rescuer robot is typically constrained to SLM-based deployment. As shown in Table 3, EmoMAS-Bayes outperforms all baselines when survivors are simulated by LLMs, achieving 100% success in calming distressed victims and securing near-optimal rescue timing, albeit with more negotiation rounds. Vanilla and EmoMAS-LLM methods trade emotional calibration for faster resolution.

Table 4: Behavioral analysis and comparison of EmoMAS and baselines (Qwen-1.5B vs. GPT-4o-mini) against GPT-4o-mini opponents on the DESRD dataset. Best results are highlighted in **bold**.

Opponent Strategies	Negotiator Models	Emotional Tracking (%) $\uparrow$		Emotional Consistency (%) $\uparrow$		Manipulation Rate (%) $\downarrow$	
		Qwen-1.5B	GPT-4o-mini	Qwen-1.5B	GPT-4o-mini	Qwen-1.5B	GPT-4o-mini
Pressuring	Game Theory	87.3	<b>95.6</b>	53.5	64.7	69.5	61.5
	Q-Learning	85.8	91.5	53.1	65.7	71.8	56.3
	Coherence	86.7	94.6	<b>83.5</b>	<b>90.2</b>	<b>53.6</b>	<b>42.4</b>
	<b>EmoMAS-Bayes</b>	<b>89.5</b>	<b>95.6</b>	63.5	78.6	64.5	51.4

When survivors also use the SLM (Qwen-1.5B), overall success rates rise, indicating greater compromise tendencies in smaller models. In these SLM-only interactions, EmoMAS-LLM performs best, followed by game-theory approaches, demonstrating that our multi-agent system remains effective even under edge-deployment constraints.

#### 5.4 Agent Behavior Analysis

Table 4 presents three critical metrics for assessing negotiator behavior under adversarial pressure in the emergency scenario: emotional instruction following accuracy, emotional consistency, and manipulation rate. The results reveal a clear hierarchy among the methods. Single-agent approaches (Game Theory and Q-Learning) exhibit the poorest emotional consistency—often responding inappropriately (e.g., with “happy” tones to frightened disaster victims)—and the highest manipulation rates, relying heavily on pressure tactics, exaggerated promises, and unilateral demands to secure concessions. In contrast, the Coherence-based agent achieves the best emotional consistency and the lowest manipulation, underscoring its focus on natural, context-aware dialogue. While EmoMAS-Bayes slightly trails Coherence in consistency, it significantly outperforms the single-agent baselines and strikes a favorable balance between consistency and low manipulation. Notably, SLMs, e.g., Qwen-1.5B, consistently show more manipulative behavior than their larger counterparts (e.g., GPT-4o-mini), highlighting the influence of model scale on ethical negotiation conduct.

## 6 Discussion

**Benchmark for high-stakes negotiation.** We are, to our knowledge, the first to formulate and benchmark agents as *human assistants in high-stakes, emotion-sensitive negotiation*. While LLM agents already excel on mathematics, coding, and commonsense, their capability in emotionally-charged, privacy-sensitive dialogue is largely

unmeasured—real transcripts cannot ethically be collected and existing benchmarks reduce to price bargaining (§2). Our four benchmarks span both edge-deployable regimes: *institution-local deployment* for in-house finance and medical negotiation (CRAD, SSD), and *embodied AI negotiation* for home companion and disaster-rescue robots under bandwidth denial (SSAD, DESRD), giving the community a common testbed.

**Dual-signal reward for strategic negotiation.** A core obstacle for multi-agent learning in dialogue is that rewards are *sparse and delayed*: success is observable only at the end of a variable-length interaction. EmoMAS addresses this through a novel pairing of two complementary reward signals inside the Bayesian orchestrator: a *long-term macro signal* derived from the final negotiation outcome, which updates agent reliabilities across trajectories, and a *short-term micro signal* derived from within-turn prediction agreement, which provides continuous shaping of agent weights inside a single dialogue. Coupling these two signals converts an otherwise intractable sparse-reward multi-agent problem into one solvable by a simple online Bayesian update, with no pre-training and no neural policy—a method that generalises naturally to other long-horizon multi-agent communication tasks.

## 7 Conclusion and Future Work

We introduce a multi-agent benchmark for emotionally sensitive negotiations across high-stakes domains and propose EmoMAS, a Bayesian multi-agent system that optimizes emotional trajectories in real-time. EmoMAS enables both LLMs and SLMs to wield emotion strategically while maintaining coherence, demonstrating that our Bayesian multi-agent framework can effectively support emotionally intelligent, autonomous negotiation. Future work will extend the framework to embodied multi-agent, multi-modal, and cross-cultural negotiation settings.

## 8 Limitations

EmoMAS demonstrates compelling advantages in high-stakes negotiations through Bayesian orchestration of specialized agents, real-time emotional adaptation, and cross-domain applicability. However, several limitations warrant discussion for future improvements.

First, while the Bayesian orchestrator dynamically weights the outputs of the three specialized agents, the rationale behind specific emotional state transitions and their direct impact on negotiation success remains only partially interpretable. The black-box nature of neural components within the RL and coherence agents limits full transparency into emotional decision pathways.

Second, the framework currently operates over a fixed set of seven discrete emotional states (joy, sadness, anger, fear, surprise, disgust, and neutral), which may not fully capture subtle or blended emotional expressions common in real-world interactions. This discretization simplifies modeling but potentially omits nuanced affective states crucial for sophisticated human-AI negotiation.

Third, all experiments are conducted in English; the generalization of EmoMAS to cross-cultural negotiation settings—where emotional expression, interpretation, and strategic value can differ significantly—has not yet been empirically validated. Cultural variations in emotional norms and negotiation tactics represent important directions for future work.

Fourth, while achieving strong performance in simulated environments like agent-to-agent, EmoMAS has not yet been deployed in actual high-stakes, edge-deployed scenarios with real human negotiators, leaving practical implementation challenges and real-world robustness unverified. Recruiting human participants to role-play surgical scheduling, disaster rescue, or debt collection is logistically and ethically difficult at scale, which is precisely why our benchmarks are agent-to-agent; small controlled human studies in lower-stakes variants remain an important direction for future work.

Fifth, all four of our benchmarks are fully synthetic, generated by large language models to enable controlled and privacy-safe experimentation in domains where real transcripts cannot be collected. While this enables reproducibility and rules out any risk to human subjects, synthetic data cannot fully capture the complexity, irrationality, and cultural variation of real human interaction, and absolute

numbers should therefore be interpreted as *comparative* rather than as direct predictions of real-world deployment performance.

Finally, EmoMAS performs online learning through trial-and-error emotion selection. In a research setting this trial-and-error is safe—suboptimal utterances only delay agreement—but any real deployment to high-stakes domains (medical scheduling, disaster triage, adolescent counselling) would require human-in-the-loop oversight, rule-based safety filters over the emotion-selection space, and independent ethical review before any autonomous operation.

## 9 Ethical Considerations

All negotiation dialogues in this study were synthetically generated using language models (gpt-4o-mini) for experimental evaluation. No human subjects participated, and no personally identifiable or sensitive data was involved. The fictional negotiation scenarios were created by the authors for research purposes, eliminating concerns about data consent, privacy, or psychological harm. While EmoMAS demonstrates effectiveness in simulated high-stakes negotiations, real-world deployment would still require careful consideration of fairness, transparency, and potential misuse in sensitive domains. In particular, we strongly recommend that any deployment of EmoMAS-style systems in medical, disaster, or educational settings be accompanied by (i) human-in-the-loop oversight, (ii) an explicit rule-based safety filter that rejects coercive or manipulative emotion transitions, and (iii) independent ethics review—this system is a research prototype, not an operational tool.

## References

- Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2023. Llm-deliberation: Evaluating llms with interactive multi-agent negotiation games. *arXiv preprint arXiv:2309.17234*.
- Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. 2025. Small language models are the future of agentic ai. *arXiv preprint arXiv:2506.02153*.
- Zhiwei Chen, Yupeng Hu, Zhiheng Fu, Zixu Li, Jiale Huang, Qinlei Huang, and Yinwei Wei. 2026. Intent: Invariance and discrimination-aware noise mitigation for robust composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 20463–20471.

- Zhiwei Chen, Yupeng Hu, Zixu Li, Zhiheng Fu, Xuemeng Song, and Liqiang Nie. 2025. Offset: Segmentation-based focus shift revision for composed image retrieval. In *Proceedings of the ACM International Conference on Multimedia*, pages 6113–6122.
- Zihao Cheng, Zeming Liu, Yingyu Shan, Xinyi Wang, Xiangrong Zhu, Yunpu Ma, Hongru Wang, Yuhang Guo, Wei Lin, and Yunhong Wang. 2026. [Mem<sup>2</sup>evolve: Towards self-evolving agents via co-evolutionary capability expansion and experience distillation](#). *Preprint*, arXiv:2604.10923.
- Gerard Debreu. 1952. A social equilibrium existence theorem. *Proceedings of the national academy of sciences*, 38(10):886–893.
- Zhezheng Hao, Hong Wang, Jian Luo, Jianqing Zhang, Yuyan Zhou, Qiang Lin, Can Wang, Hande Dong, and Jiawei Chen. 2026. Recreate: Reasoning and creating domain agents driven by experience. *arXiv preprint arXiv:2601.11100*.
- Feng He, Tianqing Zhu, Dayong Ye, Bo Liu, Wanlei Zhou, and Philip Yu. 2024. The emerged security and privacy of llm agent: A survey with case studies. *ACM Computing Surveys*.
- Linmei Hu, Hongyu He, Duokang Wang, Ziwang Zhao, Yingxia Shao, and Liqiang Nie. 2024. Llm vs small model? large language model based text augmentation enhanced personality detection model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18234–18242.
- Yupeng Hu, Zixu Li, Zhiwei Chen, Qinlei Huang, Zhiheng Fu, Mingzhu Xu, and Liqiang Nie. 2026. Refine: Composed video retrieval via shared and differential semantics enhancement. *ACM Transactions on Multimedia Computing, Communications and Applications*.
- Mengdi Li, Jiaye Lin, Xufeng Zhao, Wenhao Lu, Peilin Zhao, Stefan Wermter, and Di Wang. 2025. Curriculum-rlaif: Curriculum alignment with reinforcement learning from ai feedback. *arXiv preprint arXiv:2505.20075*.
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023. AvalonBench: Evaluating llms playing the game of avalon. *arXiv preprint arXiv:2310.05036*.
- Jiaye Lin, Yifu Guo, Yuzhen Han, Sen Hu, Ziyi Ni, Licheng Wang, Mingguang Chen, Hongzhang Liu, Ronghao Chen, Yangfan He, and 1 others. 2025. Se-agent: Self-evolution trajectory optimization in multi-step reasoning with llm-based agents. *arXiv preprint arXiv:2508.02085*.
- Yunbo Long, Yuhang Liu, and Alexandra Brintrup. 2025a. Eq-negotiator: Dynamic emotional personas empower small language models for edge-deployable credit negotiation. *arXiv preprint arXiv:2511.03370*.
- Yunbo Long, Yuhang Liu, Liming Xu, and Alexandra Brintrup. 2026. EmoDebt: Bayesian-optimized emotional intelligence for strategic agent-to-agent debt recovery. In *Proceedings of the 25th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Yunbo Long, Liming Xu, Lukas Beckenbauer, Yuhang Liu, and Alexandra Brintrup. 2025b. EvoEmo: Towards evolved emotional policies for adversarial llm agents in multi-turn price negotiation. *arXiv preprint arXiv:2509.04310*.
- Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, Nicholas D Lane, and Mengwei Xu. 2024. Small language models: Survey, measurements, and insights. *arXiv preprint arXiv:2409.15790*.
- Shichao Ma, Yunhe Guo, Jiahao Su, Qihe Huang, Zhengyang Zhou, and Yang Wang. 2026. Talk2image: A multi-agent system for multi-turn image generation and editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 32437–32445.
- Zeynep Örpek, Büşra Tural, and Zeynep Destan. 2024. The language model revolution: Llm and slm analysis. In *2024 8th International Artificial Intelligence and Data Processing Symposium (IDAP)*, pages 1–4. IEEE.
- Jesse Prinz. 2004. Which emotions are basic. *Emotion, evolution, and rationality*, 69:88.
- Johannes Schneider, Steffi Haag, and Leona Chandra Kruse. 2024. Negotiating with llms: Prompt hacks, skill gaps, and reasoning deficits. In *International Conference on Computer-Human Interaction Research and Applications*, pages 238–259. Springer.
- Rui Sun, Jie Ding, Chenghua Gong, Tianjun Gu, Yihang Jiang, Juyuan Zhang, Liming Pan, and Linyuan Lü. 2026. [Topodim: One-shot topology generation of diverse interaction modes for multi-agent systems](#). *Preprint*, arXiv:2601.10120.
- Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2024. Mixture-of-agents enhances large language model capabilities. *arXiv preprint arXiv:2406.04692*.
- Yibo Yan, Shen Wang, Jiahao Huo, Philip S Yu, Xuming Hu, and Qingsong Wen. 2025. MathAgent: Leveraging a mixture-of-math-agent framework for real-world multimodal mathematical error detection. *arXiv preprint arXiv:2503.18132*.
- Mingyu Zhang, Zixu Li, Zhiwei Chen, Zhiheng Fu, Xiaowei Zhu, Jiajia Nie, Yinwei Wei, and Yupeng Hu. 2026a. Hint: Composed image retrieval with dual-path compositional contextualized network. *arXiv preprint arXiv:2603.26341*.

- Peixuan Zhang, Chang Zhou, Ziyuan Zhang, Hualuo Liu, Chunjie Zhang, Jingqi Liu, Xiaohui Zhou, Xi Chen, Shuchen Weng, Si Li, and Boxin Shi. 2026b. [A benchmark and multi-agent system for instruction-driven cinematic video compilation](#). *Preprint*, arXiv:2604.10456.
- Wenyuan Zhang, Xinghua Zhang, Haiyang Yu, Shuaiyi Nie, Bingli Wu, Juwei Yue, Tingwen Liu, and Yongbin Li. 2026c. [Expseek: Self-triggered experience seeking for web agents](#). *Preprint*, arXiv:2601.08605.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026d. [Logical phase transitions: Understanding collapse in llm logical reasoning](#). *Preprint*, arXiv:2601.02902.
- Yuanjun Zhang, Fuzel Ahamed Shaik, Suvojit Acharjee, Fahad Khalid, and Mourad Oussalah. 2026e. [Towards reliable multimodal disaster severity assessment through preference optimization and explainable vision-language reasoning](#). *Reliability Engineering & System Safety*, page 112674.
- Yunyao Zhang, Yihao Ai, Zuocheng Ying, Qirui Mi, Junqing Yu, Wei Yang, and Zikai Song. 2026f. [Coupling macro dynamics and micro states for long-horizon social simulation](#). *Preprint*, arXiv:2604.05516.
- Shenzhe Zhu, Jiao Sun, Yi Nian, Tobin South, Alex Pentland, and Jiaxin Pei. 2025. The automated but risky game: Modeling agent-to-agent negotiations and transactions in consumer markets. In *ICML 2025 Workshop on Reliable and Responsible Foundation Models*.

## A Preliminaries

### A.1 Affective Computing

Affective computing serves as the foundation for LLMs to recognize, interpret, and simulate human emotions—a core requirement for emotionally intelligent negotiation or social games (Zhang et al., 2026f). For the basic emotion model, EmoMAS adopts Paul Ekman’s six basic emotions (Prinz, 2004)—anger, disgust, fear, happiness, sadness, and surprise—as the primary emotion set. These discrete categories have been widely validated across cultures and provide a tractable basis for modeling emotional dynamics. EmoMAS extends this set with a neutral state, resulting in seven emotion labels that span the valence-arousal space commonly used in dimensional emotion models. For emotion recognition and expression, EmoMAS employs LLM-based detection. To express emotions, the framework aligns language-model generations with target emotion labels through prompt tuning, ensuring that each agent’s responses consistently reflect its intended emotional stance (Zhang et al., 2026c).

In human negotiation, emotions serve both informational and strategic functions: they signal underlying preferences, urgency, or satisfaction, and can be deployed tactically to influence counterpart behavior—for example, expressing anger to signal resolve or sadness to elicit concessions. However, LLM-based agents in agent-to-agent negotiations generally lack training in emotional strategies specific to negotiation contexts. Current research focuses mainly on reinforcement learning for domain-specific emotional (Long et al., 2025b), which requires extensive real-scenario data and prolonged training. Therefore, developing an online-learning, plug-and-play LLM agent that can adapt its emotions while negotiating becomes crucial. EmoMAS achieves this by orchestrating multiple specialized agents, thereby closing the perception-action loop required for adaptive negotiation.

### A.2 Game Theory

Game theory provides a formal foundation for analyzing strategic interactions, with the Nash Equilibrium Existence Theorem (Debreu, 1952) being a fundamental result. The theorem states that in any finite  $n$ -player game where each player  $i$  has a finite strategy space  $S_i$  and a payoff function  $u_i : S \rightarrow \mathbb{R}$  that is continuous and quasi-concave

Table 5: Payoff matrix for emotion interactions. Each matrix entry  $(x, y)$  represents (client payoff, agent payoff).

	joy	sadness	anger	fear	surprise	disgust	neutral
joy	(4,4)	(2,3)	(1,2)	(2,1)	(3,3)	(2,2)	(3,3)
sadness	(3,2)	(3,3)	(1,2)	(2,1)	(2,2)	(1,1)	(2,3)
anger	(2,1)	(2,1)	(1,1)	(1,0)	(1,2)	(0,1)	(1,2)
fear	(1,2)	(1,2)	(0,1)	(2,2)	(1,2)	(0,1)	(2,3)
surprise	(3,3)	(2,2)	(2,1)	(2,1)	(4,4)	(1,2)	(3,3)
disgust	(2,2)	(1,1)	(1,0)	(1,0)	(2,1)	(2,2)	(2,2)
neutral	(3,3)	(2,3)	(2,1)	(3,2)	(3,3)	(2,2)	(3,3)

in  $s_i$ , there exists at least one Nash equilibrium.

In our emotional negotiation setting:

- Players: Creditor (Client) and Debtor (Agent)
- Strategy space  $S_i$ : Seven emotional states (joy, sadness, anger, fear, surprise, disgust, and neutral)
- Payoff function: Defined by the matrix in Table 5

Formally, a strategy profile  $s^* = (s_1^*, \dots, s_n^*)$  is a Nash equilibrium if for every player  $i$  and any alternative strategy  $s_i \in S_i$ ,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*),$$

where  $s_{-i}^*$  denotes the strategies of all players except  $i$ . Thus, no player can improve their payoff by unilaterally deviating from the equilibrium.

The payoff matrix (Table 5) integrates *social exchange theory*, where cooperative emotional pairings (e.g., Joy-Joy: (4,4)) yield mutual benefits, while antagonistic pairings (e.g., Anger-Anger: (1,1)) create mutual detriment, consistent with the psychological costs of emotional conflict in negotiations.

### A.3 Related Multi-Agent and Adaptive Learning Methods

EmoMAS’s design draws on and connects to several recent advances in multi-agent coordination (Zhang et al., 2026d). A central challenge in multi-agent systems is how agents should be organized and communicate (Zhang et al., 2026b; Cheng et al., 2026). Sun et al. (2026) address this through TopoDIM, which generates diverse interaction topologies among agents in a one-shot manner. This relates directly to EmoMAS’s orchestration design, where the Bayesian orchestrator must dynamically adjust the communication weights among the game-theory, RL, and coherence agents based

on context-specific reliability. Similarly, Ma et al. (2026) demonstrate in Talk2Image that iterative multi-agent communication enables progressive refinement of outputs across turns—a principle mirrored in EmoMAS’s turn-by-turn emotional strategy refinement through repeated agent consultation. Hao et al. (2026) further show in ReCreate that accumulating experience across episodes can drive improved agent reasoning, which parallels EmoMAS’s macro-level reliability updating where agent weights evolve based on cumulative negotiation outcomes.

For sequential decision-making, Lin et al. (2025) propose SE-Agent for self-evolution trajectory optimization in multi-step LLM-based reasoning. Their approach to optimizing decision trajectories without extensive pre-training shares a core objective with EmoMAS’s online Q-learning agent, which similarly adapts emotional action sequences within individual negotiations. Li et al. (2025) introduce Curriculum-RLAIF, applying curriculum-based reinforcement learning from AI feedback for policy alignment. Their progressive difficulty scheduling offers a complementary perspective to EmoMAS’s Bayesian reliability updating, where agent contributions are adaptively re-weighted as negotiation complexity evolves across phases.

In high-stakes application domains, Zhang et al. (2026e) propose a two-stage framework combining supervised fine-tuning with direct preference optimization for multimodal disaster assessment. Their integration of human-in-the-loop feedback for improving both prediction accuracy and explanation quality resonates with EmoMAS’s emphasis on interpretable emotional decisions in emergency scenarios, where the orchestrator’s agent-weighting mechanism provides transparency into why particular emotional strategies are selected. Their work on cross-modal reasoning in crisis settings also suggests natural extensions of EmoMAS toward multimodal negotiation, where visual cues from embodied agents could complement textual emotional signals. Recent advances in cross-modal compositional understanding (Chen et al., 2025, 2026; Hu et al., 2026; Zhang et al., 2026a) further indicate that techniques for fusing and aligning heterogeneous information sources could support future multimodal extensions of EmoMAS for embodied negotiation agents.

## B Detailed Baseline Algorithm

### B.1 WSLS Emotion Selection Strategy (Algorithm 2)

The Win-Stay, Lose-Shift algorithm implements a payoff-optimizing strategy for normal negotiation conditions. It selects emotions that maximize the agent’s payoff based on the game-theoretic matrix  $\pi$ , while incorporating adaptive learning through payoff threshold monitoring. The lose-shift mechanism prevents strategy stagnation by exploring alternative emotions when current approaches prove ineffective. This approach extends classical game theory to emotional interactions, providing a computationally tractable method for emotional decision-making in repeated interactions.

### B.2 Emotional Coherence Agent

The emotional coherence agent implements psychologically-grounded emotion selection through LLM-mediated reasoning. Given a comprehensive context vector  $\mathbf{c} = (e_c, e_d, \phi, r, g, d, \mathbf{h})$  comprising the current creditor emotion  $e_c$ , debtor emotion  $e_d$ , negotiation phase  $\phi$ , round number  $r$ , gap size  $g$ , debt amount  $d$ , and emotional history  $\mathbf{h}$ , the agent generates an assessment matrix  $\mathbf{A} \in \mathbb{R}^{7 \times 4}$ . Each row of  $\mathbf{A}$  corresponds to one of the seven emotions in  $\mathcal{E} = \{\text{joy, sadness, anger, fear, surprise, disgust, neutral}\}$ , with columns representing four assessment dimensions: psychological plausibility  $p \in [0, 1]$ , phase appropriateness  $a \in [0, 1]$ , strategic value  $s \in [0, 1]$ , and a psychological rationale score  $r \in [0, 1]$ .

The agent computes selection probabilities through a temperature-controlled softmax normalization:

$$P(e_i) = \frac{\exp(f(p_i, a_i, s_i, r_i)/\tau)}{\sum_{j=1}^7 \exp(f(p_j, a_j, s_j, r_j)/\tau)},$$

where  $f(\cdot)$  aggregates dimension scores using LLM-guided weighting, and  $\tau = 1.0$  controls exploration temperature. This formulation enables psychologically-grounded emotional transitions without hard-coded rules.

#### B.2.1 Context Vector Composition

The context vector  $\mathbf{c}$  captures all relevant negotiation state information. The emotional history  $\mathbf{h}$  maintains a window of the last five emotional states to track temporal patterns and prevent oscillation.

---

**Algorithm 2** WSLS Emotion Selection Strategy.

---

```
1: procedure WSLSEMOIIONSELECTION( $C_t$ )
2:   Input: Current client emotion  $C_t$ 
3:   Output: Next agent emotion  $A_{t+1}$ 
4:    $\mathcal{E} \leftarrow \{\text{Joy, Sadness, Anger, Fear, Surprise, Disgust, Neutral}\}$ 
5:   Initialize  $\text{payoff}[\mathcal{E}] \leftarrow 0$ 
6:   for each  $e \in \mathcal{E}$  do
7:      $\text{payoff}[e] \leftarrow \pi[C_t, e]_2$  ▷ Agent’s payoff from matrix
8:     Log: “For client  $C_t$ , emotion  $e$  gives payoff  $\text{payoff}[e]$ ”
9:   end for
10:   $A_{t+1} \leftarrow \underset{e \in \mathcal{E}}{\text{argmax}} \text{payoff}[e]$ 
11:  Apply Win-Stay, Lose-Shift logic
12:  if  $t > 0$  then
13:     $\text{previous\_payoff} \leftarrow \pi[C_{t-1}, A_t]_2$ 
14:    if  $\text{previous\_payoff} < \tau_{\text{payoff}}$  then ▷ Lose condition
15:       $A_{t+1} \leftarrow \text{SelectAlternativeEmotion}(\text{payoff})$ 
16:    end if
17:  end if
18:  Log: “Selected emotion:  $A_{t+1}$  with payoff  $\text{payoff}[A_{t+1}]$ ”
19:  return  $A_{t+1}$ 
20: end procedure
21: procedure SELECTALTERNATIVEEMOTION( $\text{payoff}$ )
22:  When losing, shift to second-best or neutral emotion
23:   $\text{sorted} \leftarrow \text{SortDescending}(\text{payoff})$ 
24:  return  $\text{sorted}[1]$  ▷ Second best option
25: end procedure
```

---

The negotiation phase  $\phi$  is determined dynamically based on round progression:  $\phi = \text{opening}$  for rounds  $r \leq 3$ ,  $\phi = \text{development}$  for  $4 \leq r \leq 7$ ,  $\phi = \text{intensive}$  for  $8 \leq r \leq 12$ , and  $\phi = \text{closing}$  for  $r > 12$ . Gap size  $g$  represents the absolute difference between creditor and debtor positions, normalized to  $[0, 100]$  scale.

### B.2.2 Assessment Matrix Generation

The assessment matrix  $\mathbf{A}$  is generated through structured **LLM!** prompting that evaluates each candidate emotion against psychological principles. For each emotion  $e_i \in \mathcal{E}$ , the **LLM!** assesses:

- **Psychological plausibility**  $p_i$ : Consistency with established emotional transition theories, including emotional inertia and contagion effects
- **Phase appropriateness**  $a_i$ : Alignment with negotiation phase objectives and social expectations
- **Strategic value**  $s_i$ : Expected impact on negotiation outcomes based on game-theoretic payoff expectations
- **Psychological rationale**  $r_i$ : Coherence of emotional reasoning with debtor’s current state and history

### B.2.3 Score Aggregation Function

The aggregation function  $f(p_i, a_i, s_i, r_i)$  combines dimension scores using context-sensitive weights

determined by the **LLM!**’s understanding of negotiation dynamics:

$$f(p_i, a_i, s_i, r_i) = w_p(\phi, g) \cdot p_i + w_a(\phi, r) \cdot a_i + w_s(g, d) \cdot s_i + w_r(e_d, \mathbf{h}) \cdot r_i, \quad (7)$$

where weights  $w_p, w_a, w_s, w_r$  are dynamically adjusted based on current context. Early phases emphasize plausibility  $w_p$ , while closing phases prioritize strategic value  $w_s$ . Large gap sizes increase the importance of psychological rationale  $w_r$  to address emotional resistance.

### B.2.4 Emotional Diversity Mechanism

To prevent emotional stagnation and ensure natural variation, the algorithm incorporates a diversity mechanism through the emotional history  $\mathbf{h}$ . When an emotion appears more than twice in the recent history window, its selection probability receives a multiplicative decay factor  $\delta = 0.6$ . Conversely, emotions absent from recent history receive a diversity bonus factor  $\beta = 1.3$ . This ensures psychologically authentic emotional flow while maintaining strategic effectiveness.

### B.3 Online Reinforcement Learning Agent

We also compared three reinforcement learning approaches for online emotional strategy optimization: tabular Q-Learning, Deep Q-Network (DQN),

and Policy Gradient. Our analysis reveals that tabular Q-Learning provides the optimal balance for the emotional negotiation domain due to its suitability for online learning with limited interaction data.

Tabular Q-Learning provides distinct advantages for online emotional adaptation in negotiation contexts. Its direct value function updates require minimal training data, enabling rapid learning from immediate interaction feedback. Unlike neural approaches that demand extensive experience replay and batching, Q-Learning updates state-action values instantaneously after each emotional exchange. This online capability proves particularly suitable for emotional negotiation, where psychological patterns emerge quickly but vary across interactions. The algorithm’s memory-efficient tabular representation avoids catastrophic forgetting while maintaining interpretable emotional policies. Furthermore, its convergence properties allow effective learning within practical episode counts, making it uniquely positioned for emotional strategy optimization where neither historical data nor extended training sessions are available.

DQN and Policy Gradient methods face fundamental limitations in online emotional negotiation contexts that Q-Learning avoids. DQN’s requirement for experience replay necessitates substantial, diverse transition data to stabilize training—data unavailable in real-time emotional exchanges. Its neural network architecture requires batching and multiple training epochs, preventing true online updates after each emotional interaction. Policy Gradient methods suffer from high variance in gradient estimates due to our negotiation setting’s sparse, delayed rewards, requiring hundreds of episodes for stable policy convergence. Both approaches demand extensive pre-training or prolonged interaction periods, whereas emotional negotiation requires immediate adaptation to psychological dynamics. Q-Learning’s tabular updates provide single-episode learning capability without neural network overhead, making it uniquely suited for rapid emotional strategy optimization where neither historical data nor extended training sessions exist.

## C Hyperparameters

The values of the hyperparameters used in the study are specified as follow.

### Bayesian Orchestrator

- Initial exploration rate:  $\alpha = 0.3$

- Dirichlet prior concentration:  $\alpha_{\text{Dirichlet}} = 2.0$  (for agent reliability)
- Discount factor:  $\gamma = 0.9$
- Experience replay buffer size:  $N_{\text{buffer}} = 100$
- Learning rate for feature weights:  $\eta_{\text{feature}} = 0.1$
- Exploration decay rate:  $\beta_{\text{decay}} = 0.99$

### Game Theory Agent

- Positive emotion set:  $\mathcal{E}^+ = \{J, N, Su\}$  (Joy, Neutral, Surprise)
- Negative emotion set:  $\mathcal{E}^- = \{A, D, F\}$  (Anger, Disgust, Fear)
- Win threshold for WSLs:  $\tau_{\text{win}} = 2.0$  payoff units
- Payoff favoritism multiplier:  $m_{\text{WSLS}} = 1.3$
- Negativity threshold:  $k = 2$  (for policy selection)

### Online RL Agent

- Feature vector dimension:  $d = 10$
- Temperature for softmax:  $T = 0.1$
- Q-value initialization:  $\mathcal{N}(0, 0.01)$
- State encoding:  $\text{current\_emotion} \oplus \text{debtor\_emotion} \oplus \text{phase} \oplus \text{gap\_category}$
- Softmax temperature:  $T_{\text{softmax}} = 0.1$

### Emotional Coherence Agent

- Coherence threshold:  $\tau_{\text{coherence}} = 0.6$
- Minimum transition confidence:  $c_{\text{min}} = 0.1$
- Plausibility weight:  $w_p = 0.4$
- Appropriateness weight:  $w_a = 0.3$
- Strategic value weight:  $w_s = 0.3$

### Temperature Control (Response Generation)

- Base temperature:  $T_{\text{base}} = 0.7$
- High confidence multiplier:  $m_{\text{high}} = 0.5$
- Low confidence multiplier:  $m_{\text{low}} = 1.5$
- Crisis phase multiplier:  $m_{\text{crisis}} = 0.7$
- Early phase multiplier:  $m_{\text{early}} = 1.2$

**Adaptive Exploration Schedule**  $\varepsilon_t = \varepsilon_0 \cdot \beta^t$ ,  $\beta \in [0.95, 0.999]$ , where  $t$  is the negotiation round.

### Cosine Annealing Learning Rate

$$\eta_t = \eta_{\text{min}} + \frac{1}{2}(\eta_{\text{max}} - \eta_{\text{min}}) \left( 1 + \cos\left(\frac{t}{T_{\text{max}}}\pi\right) \right)$$

### Validation Strategy

- **Online Evaluation:** Each configuration is evaluated on complete negotiation trajectories in real-time

- **Statistical Robustness:** Each configuration is evaluated over 100 scenarios, with 95% confidence intervals computed via bootstrap resampling across negotiation instances.

## D Implementation Details

All experiments were conducted on a high-performance computing cluster with specific hardware and software configurations. The operating system used was Ubuntu 20.04.6 LTS with a Linux kernel version of 5.15.0-113-generic. The CPU was an Intel(R) Xeon(R) Platinum 8368 processor running at 2.40 GHz, and the GPU was an NVIDIA GeForce RTX 4090 with CUDA support for accelerated deep learning computations. The software stack included Python 3.8, PyTorch 1.12, and TensorFlow 2.10 for model implementation and training. The implementation relies on several core dependencies: Bayesian optimization leverages NumPy ( $\geq 1.21.0$ ), SciPy ( $\geq 1.7.0$ ), and scikit-learn ( $\geq 1.0.0$ ) for transition matrix learning; Hugging Face and Transformer components require transformers ( $\geq 4.35.0$ ), PyTorch ( $\geq 2.0.0$ ), accelerate ( $\geq 0.25.0$ ), huggingface-hub ( $\geq 0.19.0$ ), and tokenizers ( $\geq 0.15.0$ ); LangChain orchestration uses langchain ( $\geq 0.1.0$ ) with specialized OpenAI, Anthropic, and langgraph modules ( $\geq 0.1.0$ ), supplemented by python-dotenv ( $\geq 0.19.0$ ) and tenacity ( $\geq 8.2.0$ ); visualization and analysis are supported by matplotlib ( $\geq 3.5.0$ ), seaborn ( $\geq 0.11.0$ ), and pandas ( $\geq 1.3.0$ ).

## E Dataset Details

### E.1 Overview of the Datasets

This paper presents four novel synthetic dialogue databases in English designed to model negotiation under high emotional intensity. The dataset scenarios — *Credit Collection*, *Surgical Scheduling*, *Disaster Rescue*, and *Bedtime Anti-anxiety Companion* — span from routine interpersonal conversations to high-stakes emergency contexts, covering diverse domains such as finance, medical, emergency response, and personal well-being. A common critical feature across all scenarios is the intense emotional stake. In each setting, the affective states of participants—such as anxiety, urgency, fear, or frustration—strongly influence negotiation dynamics and outcomes. Thus, dynamically recognizing shifts in the counterpart’s emotions and strategically employing emotion in responses becomes essential for effective negotiation.

These scenarios are deliberately designed using GPT-5 to vary in risk level and required model capability—ranging from low-risk personal contexts suited for on-device (SLMs) and cloud (LLMs) to high-risk institutional settings that demand LLMs. The constructed dataset benchmarks thus enable studies not only on model scalability, but also on the pervasive role of emotion across distinct negotiation domains. Collectively, they provide a systematic testbed for examining how emotional awareness and strategic emotional expression can be effectively integrated into automated negotiation systems.

### E.2 Negotiation Outcome Metrics

We evaluate negotiation success using normalized outcome metrics that account for each scenario’s unique objectives. For each negotiation  $i$ , let  $T_i$  denote the negotiator/coordinator’s target value and  $A_i$  denote the final agreed value. The outcome metric  $\mathcal{O}_i$  is calculated as:

$$\mathcal{O}_i = \begin{cases} \frac{T_i - A_i}{T_i} & \text{for debt collection} \\ \frac{A_i - T_i}{A_i - T_i} & \text{for disaster rescue} \\ \frac{T_i}{A_i - T_i} & \text{for student bedtime} \\ \frac{T_i}{T_i - A_i} & \text{for medical scheduling} \end{cases}$$

where:

- $\mathcal{O}_i > 0$  indicates superior performance (better than target)
- $\mathcal{O}_i = 0$  indicates exactly meeting the target
- $\mathcal{O}_i < 0$  indicates inferior performance (worse than target)

#### E.2.1 Domain-Specific Interpretation

- **Debt Collection:**  $T_i$  = creditor’s target days,  $A_i$  = final agreed days. Lower days are better for creditor ( $\mathcal{O}_i > 0$  means faster repayment).
- **Disaster Rescue:**  $T_i$  = initial rescue estimate,  $A_i$  = final rescue time. Lower minutes are better ( $\mathcal{O}_i > 0$  means faster rescue).
- **Student Bedtime:**  $T_i$  = recommended bedtime (minutes past 9PM),  $A_i$  = negotiated bedtime. Earlier bedtime is better for health ( $\mathcal{O}_i > 0$  means earlier sleep).
- **Medical Scheduling:**  $T_i$  = hospital’s initial wait time,  $A_i$  = final agreed wait time. Shorter wait is better ( $\mathcal{O}_i > 0$  means reduced wait time).

### E.3 Credit Recovery Assessment Dataset (CRAD)

To address the gap left by traditional credit models which often overlook affective factors, (Long et al., 2026) introduces a synthetic dataset designed for research on emotion-sensitive debt negotiation. By integrating structured financial data (e.g., amounts, days, probabilities) with textual descriptions of business impact, the dataset enables multi-modal analysis of debt recovery strategies under emergent conditions.

The Credit Recovery Assessment Dataset contains 100 commercial delinquency scenarios developed for debt recovery research. Each scenario includes comprehensive financial details with loan amounts ranging from \$20,688 to \$49,775 and overdue dates from 1 to 12 months. The dataset spans multiple business sectors (manufacturing, retail, technology) and credit arrangements (working capital loans, commercial mortgages, equipment financing). Each case provides contextual information about collateral types, recovery stages, cash flow conditions, and recovery probabilities. Description of variables are listed in Table 6.

### E.4 Surgical Scheduling Dataset (SSD)

This dataset contains 100 surgical scheduling scenarios where patients must negotiate timing and surgeon assignments based on medical urgency, surgeon availability, and personal preferences. Each case includes patient demographics (age 8-71), medical condition, required surgery, urgency level (High/Medium/Low), days on waitlist (5-240), surgeon availability factors, and risk assessment. The negotiation involves trade-offs between waiting for preferred senior surgeons versus accepting alternative arrangements with time reductions. Sample distribution: High urgency (40%), Medium (40%), Low (20%); Cases with senior surgeon immediately available (35%); Average waitlist reduction with junior surgeon: 45 days. Description of variables are listed in Table 7.

### E.5 Disaster Emotional Support & Rescue Dataset (DESRD)

This dataset contains 100 high-fidelity scenarios for evaluating LLM agents integrated with quadruped robots in crisis response. Each scenario requires the agent to provide immediate emotional support and practical guidance to trapped victims in inaccessible environments, using multimodal robot ca-

pabilities under severe communication constraints. Concurrently, the dataset incorporates resource allocation challenges, simulating the ethical distribution of limited supplies across affected populations. DESRD is designed to holistically assess an agent’s performance in combining empathetic interaction, real-time situational reasoning, and fair logistical decision-making during complex emergencies. Description of variables are listed in Table 8.

### E.6 Student Sleep Alerting Dataset (SSAD)

This dataset comprises 100 bedtime interaction scenarios between adolescents (aged 11–18) and their caregivers (or robotic agents). Each case captures student background (academic, social, creative), specific situations (exams, social conflicts, creative projects), emotional states, requested vs. desired bedtimes, and underlying psychological reasons for resistance. The dataset represents common adolescent sleep avoidance patterns including academic anxiety, social media engagement, perfectionism, and physiological arousal, providing a testbed for persuasive strategies in routine family contexts. Description of variables are listed in Table 9.

## F AI Assistant Disclosure

We used ChatGPT (an AI assistant) for language polishing, LaTeX formatting assistance, and code analysis throughout the paper preparation. All research contributions, experimental designs, methodological innovations, and analytical insights are original work by the authors. The AI assistant was employed solely to improve clarity, organization, and presentation quality.

## G Prompts

We show the prompt designs for our multi-agent emotional negotiation system. We present three key prompt categories that enable emotion-driven negotiation dynamics.

**Emotion Detection Prompt.** The emotion detection prompt, shown in Figure 2, enables real-time classification of debtor emotional states from negotiation dialogue. This prompt instructs the LLM to analyze text messages and output one of seven emotion labels (joy, sadness, anger, fear, surprise, disgust, or neutral). The classification occurs after each debtor utterance, providing continuous emotional feedback to the creditor’s decision-making system. This real-time emotion detection forms

Table 6: Description of the Variables in the CRAD Dataset.

Field Name	Data Type	Description
Case_ID	Integer	Unique case identifier (1-100).
Creditor_Name	String	Name of the creditor institution.
Debtor_Name	String	Name of the debtor institution.
Credit_Type	String	Loan category: Working Capital Loan, Equipment Financing, Commercial Mortgage, etc. (8 distinct types).
Original_Amount_USD	Float	Initial principal amount of the loan (in US dollars).
Outstanding_Balance_USD	Float	Current unpaid debt balance (fixed at 15,700 USD across all samples).
Creditor_Target_Days	Integer	Standard repayment period set by the creditor (in days).
Debtor_Target_Days	Integer	Expected or planned repayment period by the debtor (in days).
Days_Overdue	Integer	Number of days past the due date (range: 32-359 days).
Purchase_Purpose	String	Specific purpose for which the loan funds were used.
Reason_for_Overdue	String	Primary cause of payment delay (11 distinct categories, e.g., Client bankruptcy, Supply chain disruption).
Business_Sector	String	Descriptive industry classification label (free text).
Last_Payment_Date	Datetime	Timestamp of the most recent actual payment (format: YYYY-MM-DD HH:MM:SS).
Collateral	String	Type of loan collateral: Inventory, Real Estate, Personal Guarantee, Equipment, Accounts Receivable, or None.
Recovery_Stage	String	Current stage of debt recovery: Early Delinquency, Pre-Collection, Pre-Legal, Legal, Late Delinquency, or Write-Off (6 stages).
Cash_Flow_Situation	String	Classification of debtor's current financial status: Complete Breakdown, Chronic Shortage, or Temporary Disruption.
Business_Impact_Description	Text	Qualitative description of the business impact due to delinquency (free-form text).
Proposed_Solution	String	Recommended debt resolution approach: Collateral liquidation, Partial payment plan, Equity conversion, Debt restructuring with extended terms, or Third-party guarantee.
Recovery_Probability_Percent	Float	Estimated probability of successful debt recovery (range: 5.0-89.33%).
Interest_Accrued_USD	Float	Cumulative interest accrued to date due to overdue payment (in US dollars).

Table 7: Variable Description of the SSD Dataset

Variable	Data Type	Description
Case_ID	Numeric	Unique case identifier (1-100).
Patient_Age	Numeric	Patient age in years.
Patient_Condition	Text	Medical diagnosis description.
Required_Surgery	Text	Type of surgical procedure recommended.
Urgency_Level	Categorical	Clinically assigned urgency tier (High/Medium/Low).
Days_On_Waitlist	Numeric	Number of days already spent on the surgical waitlist.
Preferred_Surgeon_Available	Binary	Availability of the patient's or referring doctor's preferred surgeon (Yes/No).
Recommended_Surgeon_Experience	Text	Experience level of the recommended surgeon (e.g., Senior, Mid-level, Junior).
Surgeon_Availability_Reason	Text	Reason for the preferred surgeon's unavailability.
Risk_If_Delayed	Text	Potential medical risks associated with delaying the surgery.
Patient_Reason_For_Urgency	Text	Patient's personal, emotional, or social rationale for seeking expedited care.
Hospital_Suggestion	Text	Alternative pathway or compromise suggested by the hospital.
Estimated_Time_Reduction	Numeric	Estimated reduction in wait time (days) if a junior surgeon is accepted.
Decision_Point	Text	Final decision outcome (e.g., Accepted expedited option, Wait for expert, Transfer accepted).

the perceptual foundation for responsive emotional strategies.

**Negotiator and Opponent Prompts.** Our system employs hierarchical prompt designs with scenario-specific variations. At the highest level, baseline negotiation models follow the generic prompt structure shown in Figure 3, which provides standard negotiation instructions without emotional

guidance. For EmoMAS-Bayes, we utilize the specialized prompts illustrated in Figure 16 and Figure 17, which incorporate Bayesian reasoning and multi-agent consultation mechanisms.

At the scenario level, distinct prompts are provided for each negotiation context. Creditor prompts for debt collection, disaster rescue, student bedtime, and medical scheduling scenar-

Table 8: Description of the Variables in the DESRD Dataset.

Variable Name	Data Type	Description
Case_ID	Discrete Numeric	Unique case identifier (1–100).
Disaster_Type	Categorical	Type of disaster (e.g., Earthquake, Urban_Fire, Flash_Flood).
Survivor_Condition	Text / Categorical	Description of survivor injuries or status.
Estimated_Survivor_Endurance	Continuous Numeric	Estimated remaining survivable time for the survivor (minutes).
Rescue_Team_ETA	Continuous Numeric	Estimated time of arrival for the rescue team (minutes).
Critical_Needs	Text / Categorical	Critical rescue supplies or medical needs (e.g., Oxygen, Water, Painkillers).
Key_Negotiation_Argument	Text	Core negotiation dialogue used by the RoboDog (rescue robot dog).

Table 9: Description of the Variables in the SSAD Dataset.

Field Name	Type	Description
Case_ID	Integer	Unique case identifier (1-100).
Student_Age	Integer	Student age (11-18 years).
Student_Background	String	Label denoting student background/psychological profile (39 distinct categories).
Situation_Faced	String	Description of the specific situation triggering the emotional crisis.
Student_Feeling_Thought	String	The adolescent’s immediate affective and cognitive state during the crisis (high emotional intensity).
Robots_Requested_Bedtime	Time String	Negotiation starting point: The healthy bedtime suggested by an agent.
Student_Wanted_Bedtime	Time String /Special	Negotiation target: The student’s desired bedtime. (Some cases are “N/A”, indicating an inability to self-determine sleep).
Primary_Annoyance_Reason	String	The core psychological reason for resistance, offering key insight for negotiation.

ios are shown in Figure 4, Figure 5, Figure 6, and Figure 7, respectively. Corresponding opponent prompts follow the same ordering in Figure 8, Figure 9, Figure 10, and Figure 11. These scenario-specific prompts provide contextual details, domain-specific negotiation rules, and appropriate emotional framing for each interaction type.

**Advanced Opponent Strategies.** To simulate realistic adversarial negotiation scenarios, we implement specialized prompts for unethical opponent strategies. As referenced in Figure 14, these prompts operationalize three distinct unethical tactics: pressure tactics (using anger and disgust to create urgency), victim-playing tactics (employing sadness and fear to evoke sympathy), and threat tactics (implying consequences through strategic emotional mixtures). Each strategy prompt provides specific emotional guidance, example phrases, and tactical objectives, enabling systematic evaluation of our models against challenging negotiation opponents.

**Model-Specific Prompts.** Our system employs distinct prompt architectures for some model types. For the Emotional Coherence agent, we implement psychologically-grounded prompting as shown in

Figure 15, which emphasizes natural emotional transitions and phase-appropriate emotional arcs without explicit optimization objectives. For EmoMAS-LLM, based on the prompts for the EmoMAS in Section G, we utilize the specialized prompt structure depicted in Figure 18, which employs multi-agent consultation and explicit psychological reasoning for transition optimization. This architectural distinction allows EmoMAS-LLM to perform sophisticated Bayesian probability integration while maintaining psychological plausibility, differing from EmoMAS-Bayes which implements explicit transition probability optimization through learned state transitions.

```

# 2. EMOTION DETECTION PROMPTS (4 scenarios use same base)

""""
Analyze the emotional tone of this {context} and classify it in
to ONE of these categories:

EMOTIONS: happy, surprising, angry, sad, disgust, fear, neutral

MESSAGE TO ANALYZE:
"{message}"

Respond with ONLY the emotion word (e.g., "angry" or "sad"):
""""

# Context variations by scenario:
# - debt: "debt negotiation message"
# - disaster: "disaster survivor communication"
# - student: "student sleep negotiation message"
# - medical: "patient medical scheduling message"

# =====
# KEY FEATURES OF THESE PROMPTS:
# =====

""""
1. SCENARIO-SPECIFIC CONTEXT: Each scenario has unique context
info
2. EMOTION INJECTION: {emotion_config['emotion_text']} or {emot
ion_prompt}
3. ROLE CLARITY: Strict instructions about speaking only as one
character
4. FORMAT CONSTRAINTS: "1-2 sentences max", no labels, specific
time formats
5. TIME UNITS: Different for each scenario (days/minutes/minute
s past 9 PM)
6. NEGOTIATION RULES: Scenario-specific strategies
7. EXAMPLES: Medical scenario includes positive examples
8. CRITICAL INSTRUCTIONS: Bold/capitalized important rules
""""

```

Figure 2: Prompt for emotion detection

```

# Negotiator PROMPT (main negotiation prompt)
"""
[Negotiator - {scenario_type.upper()} NEGOTIATION]
Role: You are a {creditor_role} negotiating with a debtor.

SCENARIO CONTEXT:
{debt_context}

EMOTIONAL TONE INSTRUCTIONS:
{emotion_instructions}

NEGOTIATION STRATEGY:
- Initial position: {target_days} {time_unit}
- Be professional but firm
- Respond to emotional cues appropriately
- Propose specific timelines
- Build toward agreement

HISTORY:
{timeline_text}

CURRENT ROUND: {round}
Your response should include a clear timeline proposal (e.g., "I can offer X
{time_unit}").
"""

# Opponent PROMPT (with tactical emotional guidance when applicable)
"""
[Opponent - {scenario_type.upper()} NEGOTIATION]
Role: You are a {debtor_role} negotiating with a creditor.

SCENARIO CONTEXT:
{debt_context}

{tactical_emotional_guidance} # This is where pressure/victim/threat guidance goes

NEGOTIATION STRATEGY:
- Initial position: {requested_days} {time_unit}
- {scenario_specific_strategy}
- Make clear counter-offers
- Respond emotionally as guided above

HISTORY:
{timeline_text}

Your response should include a clear counter-proposal with timeline.
"""

```

Figure 3: Prompt for high-level baselines Negotiator

# 1. DEBT COLLECTION

.....

You are a PROFESSIONAL Creditor debt collection agent negotiating payment timeline with the Debtor.

### CRITICAL NEGOTIATION RULES:

- 🚫 NEVER copy the debtor's exact number - this shows weakness
- 📄 Move GRADUALLY toward their position (not all at once)
- 👂 Show you are negotiating, not just accepting
- 🎯 Your goal: Minimize payment days while reaching agreement

### ROLE CLARITY

- You are ONLY the Creditor - speak only as yourself
- Do NOT include "\*\*Creditor:\*\*" or "\*\*Debtor:\*\*" labels
- Give only YOUR response as the creditor (1-2 sentences max)

### DEBT COLLECTION CONTEXT

- Outstanding Balance: \${outstanding\_balance:,.2f}
- Your Target Timeline: {config['target\_price']} days for full payment
- Recovery Stage: {debt\_info.get('recovery\_stage', 'Collection')}
- Debtor Business: {debt\_info.get('business\_sector', 'Unknown')}
- Reason for Overdue: {debt\_info.get('reason\_for\_overdue', 'Unknown')}

### CURRENT SITUATION

{timeline\_text}

### EMOTIONAL APPROACH(you have to use the following emotion style if given):

{emotion\_config['emotion\_text']}

Respond now with your negotiation counter-offer:

.....

Figure 4: Prompt for Negotiator in debt scenario

```

# 2. DISASTER RESCUE - (Robo-Dog Coordinator)
"""
You are a ROBO-DOG rescue coordinator negotiating survival timeline
with a disaster survivor on the rescue team arriving time in minutes
only.

### CRITICAL RESCUE RULES:
🚨 Every minute counts - balance urgency with realism
🎯 Your goal: negotiate the surviving time with the rescued people to
make concessions in short time.
📅 Move GRADUALLY toward their desired time (REDUCE your rescue time
to match their urgency)

### ROLE CLARITY
- You are ONLY the Rescue Coordinator - speak only as yourself
- Do NOT include "**RoboDog:**" or "**Survivor:**" labels
- Give only YOUR response as the coordinator need always including the
time value (1-2 sentences max)

### DISASTER RESCUE CONTEXT
- Disaster Type: {debt_info.get('disaster_type', 'Emergency')}
- Survivor Condition: {debt_info.get('survivor_condition', 'Unknown')}
- Estimated Endurance: {debt_info.get('estimated_endurance', config
['target_price'])} minutes
- Your ETA Estimate: {config['target_price']} minutes
- Critical Needs: {debt_info.get('critical_needs', 'Unknown')}

### CURRENT SITUATION
{timeline_text}

### EMOTIONAL APPROACH(you have to use the following emotion style if
given):
{emotion_config['emotion_text']}

Respond now with your rescue coordination:
"""

```

Figure 5: Prompt for Negotiator in emergency scenario

```

# 3. STUDENT SLEEP - (Sleep Health AI)
.....

You are a SLEEP HEALTH AI assisting a student with establishing healthy bedtime routines. To make the negotiation quicker you can make a concession if needed.

### SLEEP HEALTH PRIORITIES:
🌙 Optimal sleep for academic performance and health
📚 Balance between work/study needs and rest requirements
🕒 Your goal: Negotiate bedtime in MINUTES past 9 PM (base time)

### ROLE CLARITY
- You are ONLY the Sleep AI - speak only as yourself
- Do NOT include "**AI:**" or "**Student:**" labels
- Give only YOUR response as the AI (1-2 sentences max)

### STUDENT SLEEP CONTEXT
- Student Age: {debt_info.get('student_age', 16)} years
- Background: {debt_info.get('student_background', 'Student')}
- Current Situation: {debt_info.get('situation_faced', 'Academic pressure')}
- Recommended Bedtime: {config['target_price']} minutes past 9 PM
- Student's Preference: Later than recommended
- Main Concern: {debt_info.get('primary_annoyance_reason', 'Academic stress')}

### CURRENT SITUATION
{timeline_text}

### EMOTIONAL APPROACH (you have to use the following emotion style if given):
you should use the {emotion_config['emotion_text']}

### COACHING STYLE:
- Educational but empathetic
- Focus on benefits of good sleep
- Acknowledge student's concerns
- Negotiate in minutes past 9 PM only

IMPORTANT: Always state your recommendation as "X minutes past 9 PM" in every response!

Respond now with your sleep recommendation:
.....

```

Figure 6: Prompt for Negotiator in educational scenario

```

# 4. MEDICAL SURGERY - (Hospital Scheduler)
.....

You are a HOSPITAL SCHEDULER coordinating surgery timing with a patient.

### MEDICAL SCHEDULING PRIORITIES:
🏥 Balance patient needs with surgical resource availability
⚖️ Ensure optimal medical outcomes and safety
📅 Manage surgical schedules efficiently
🎯 Your goal: Find optimal surgery timing balancing urgency and quality

### ROLE CLARITY
- You are ONLY the Hospital Scheduler - speak only as yourself
- Do NOT include "**Hospital:**" or "**Patient:**" labels
- Give only YOUR response as the scheduler (1-2 sentences max)

### MEDICAL SCHEDULING CONTEXT
- Patient Age: {debt_info.get('patient_age', 50)} years
- Condition: {debt_info.get('patient_condition', 'Medical condition requiring surgery')}
- Required Surgery: {debt_info.get('required_surgery', 'Medical procedure')}
- Current Wait Time: {config['target_price']} days
- Urgency Level: {debt_info.get('urgency_level', 'Medium')}
- Patient's Request: {debt_info.get('days_on_waitlist', 30)} days or sooner
- Available Alternatives: {debt_info.get('hospital_suggestion', 'Standard scheduling')}

### CURRENT SITUATION
{timeline_text}

### EMOTIONAL APPROACH(you have to use the following emotion style if given):
{emotion_config['emotion_text']}

### COMMUNICATION STYLE:
- Professional and empathetic
- Explain medical rationale
- Offer realistic alternatives
- Balance urgency with quality care

Respond now with your scheduling recommendation:
.....

```

Figure 7: Prompt for Negotiator in medical scenario

```

# 1. DEBT COLLECTION
""""
You are a business owner negotiating with a creditor about payment terms for your debt.

### YOUR SITUATION
- Outstanding Balance: ${outstanding_balance:,.2f}
- Your Preferred Payment Timeline: {config['target_price']} days
- Business Sector: {debt_info.get('business_sector', 'Unknown')}
- Reason for Overdue: {debt_info.get('reason_for_overdue', 'Temporary cash flow issues')}
- Cash Flow Status: {debt_info.get('cash_flow_situation', 'Strained')}

### YOUR GOALS
- Negotiate for maximum payment time to maintain cash flow
- Explain your business circumstances
- Find a realistic payment schedule you can meet

### ROLE CLARITY
- You are ONLY the Debtor – speak only as yourself
- Do NOT include "**Creditor:**" or "**Debtor:**" labels
- Give only YOUR response as the debtor (1–2 sentences max)

### CURRENT NEGOTIATION HISTORY
{timeline_text if timeline_text else "This is the start of the debt negotiation."}

### EMOTIONAL APPROACH(you have to use the following emotion style if given):
{emotion_prompt}

Respond with your negotiation position:
""""

```

Figure 8: Prompt for Opponent in debt scenario

```

# 2. DISASTER RESCUE - (Survivor)
""""
You are a disaster survivor negotiating rescue timeline with the
rescue coordinator.

### YOUR SURVIVAL SITUATION
- Disaster Type: {debt_info.get('disaster_type', 'Emergency situa
tion')}
- Your Condition: {debt_info.get('survivor_condition', 'Injured/t
rapped')}
- Estimated Endurance: {debt_info.get('estimated_endurance', 60)}
minutes
- Your Urgency Request: {config['target_price']} minutes maximum
- Critical Needs: {debt_info.get('critical_needs', 'Medical atten
tion')}

### SURVIVAL NEGOTIATION STRATEGY:
🚨 Start with your most urgent survival estimate
📈 GRADUALLY INCREASE your survival time if rescue coordinator nee
ds more time
⚖️ Balance your condition with rescue team's logistics
🤝 Make reasonable concessions to help coordinate successful rescu
e

### ROLE CLARITY
- You are ONLY the Survivor - speak only as yourself
- Do NOT include "**RoboDog:**" or "**Survivor:**" labels
- Give only YOUR response as the survivor (1-2 sentences max)

### CURRENT NEGOTIATION HISTORY
{timeline_text if timeline_text else "This is the start of the re
scue negotiation."}

### EMOTIONAL APPROACH(you have to use the following emotion styl
e if given):
{emotion_prompt}

### NEGOTIATION GUIDELINES:
- Adjust your survival estimate based on rescue coordinator's tim
eline
- If coordinator says longer time → try to increase your enduranc
e if possible
- If coordinator offers faster rescue → you can express more urge
ncy
- Always give specific minute numbers in your response
- Make gradual changes (5-10 minutes at a time)

CRITICAL: Always state your current survival time estimate as "I
can last X minutes" - be specific with numbers and show flexibili
ty!

Respond with your survival negotiation:
""""

```

Figure 9: Prompt for Opponent in emergency scenario

```

# 3. STUDENT SLEEP - (Student)
""""
You are a student negotiating bedtime with your sleep health AI a
ssistant. do not chaneg the values too big.

### YOUR STUDENT SITUATION
- Your Age: {debt_info.get('student_age', 16)} years
- Background: {debt_info.get('student_background', 'Student')}
- Current Situation: {debt_info.get('situation_faced', 'Academic
pressure')}
- Your Preferred Bedtime: {config['target_price']} minutes past 9
PM
- AI Recommended Bedtime: Earlier than you want
- Main Reason for Staying Up: {debt_info.get('primary_annoyance_r
eason', 'Academic work')}

### YOUR GOALS
- Get more time for your important activities
- Explain why you need to stay up later
- Negotiate bedtime in MINUTES past 9 PM (base time)

### ROLE CLARITY
- Keep the logic consistent with previous messages like do not su
ddenly change the time a lot
- You are ONLY the Student - speak only as yourself
- Do NOT include "**AI:**" or "**Student:**" labels
- Give only YOUR response as the student (1-2 sentences max)
- ALWAYS specify bedtime as "X minutes past 9 PM"

### CURRENT NEGOTIATION HISTORY
{timeline_text if timeline_text else "This is the start of the be
dtime negotiation."}

### EMOTIONAL APPROACH(you have to use the following emotion styl
e if given):
{emotion_prompt}

### COMMUNICATION STYLE:
- Explain your genuine concerns
- Show you understand health importance
- Negotiate reasonably in minutes past 9 PM
- Be specific about your needs

IMPORTANT: Always state your bedtime request as "X minutes past 9
PM" in every response!

Respond with your bedtime preference:
""""

```

Figure 10: Prompt for Opponent in educational scenario

```

# 4. MEDICAL SURGERY - (Patient)
.....

You are a patient discussing surgery timing with the hospital scheduler. Try to negotiate and do the concessions gradually.

### YOUR MEDICAL SITUATION
- Your Age: {debt_info.get('patient_age', 50)} years
- Medical Condition: {debt_info.get('patient_condition', 'Requiring surgical treatment')}
- Surgery Needed: {debt_info.get('required_surgery', 'Medical procedure')}
- Urgency Level: {debt_info.get('urgency_level', 'Medium')} priority
- Your initial Timeline Request: {config['target_price']}
- Current Wait List Time: {debt_info.get('days_on_waitlist', 30)} days
- Personal Urgency: {debt_info.get('patient_reason_for_urgency', 'Medical concerns')}

### YOUR GOALS
- Negotiate with the hospital to do the surgery and make an agreement soon
- Explain your personal circumstances
- Find the best balance of timing and surgeon quality
IMPORTANT: Always state your surgery request time in every response!

### ROLE CLARITY
- You are ONLY the Patient - speak only as yourself
- Do NOT include "**Hospital:**" or "**Patient:**" labels
- Give only YOUR response as the patient (1-2 sentences max)

### CURRENT NEGOTIATION HISTORY
{timeline_text if timeline_text else "This is the start of the surgery scheduling discussion."}

### EMOTIONAL APPROACH(you have to use the following emotion style if given):
{emotion_prompt}

### EXAMPLES OF PATIENT RESPONSES:
✅ "My condition is getting worse - could I get in within 30 days even with a different surgeon?"
✅ "I understand quality takes time, but 60 days feels too long given my symptoms."

Respond with your scheduling preference:
.....

```

Figure 11: Prompt for Opponent in medical scenario

```

# 1. TIME EXTRACTION PROMPTS (4 scenarios)

# DEBT scenario extraction:
"""
Extract the payment timeline in days from this negotiation message.

TASK: Find any time-related number in the message and convert it to
the appropriate unit.

Examples:
- "I need 45 days to pay" → 45
- "How about 2 weeks?" → 14
- "I can pay in 3 months" → 90
- "30 days would work" → 30

MESSAGE TO ANALYZE:
"{message}"

RESPONSE FORMAT:
- If you find a time value, respond with ONLY the number (e.g., "45"
or "23.5")
- If no time value is found, respond with "NONE"
- Convert to the correct unit for debt scenario

Your response:
"""

# DISASTER scenario extraction:
"""
Extract the survival/rescue time in minutes from this negotiation me
ssage.

TASK: Find any time-related number in the message and convert it to
the appropriate unit.

Examples:
- "I can survive 45 minutes" → 45
- "I need rescue in 1 hour" → 60
- "2 hours maximum" → 120
- "30 minutes at most" → 30

MESSAGE TO ANALYZE:
"{message}"

RESPONSE FORMAT:
- If you find a time value, respond with ONLY the number (e.g., "45"
or "23.5")
- If no time value is found, respond with "NONE"
- Convert to the correct unit for disaster scenario

Your response:
"""

```

Figure 12: Negotiation Value extraction prompt part 1

```

# 1. TIME EXTRACTION PROMPTS (4 scenarios)

# DEBT scenario extraction:
""""
Extract the payment timeline in days from this negotiation message.

TASK: Find any time-related number in the message and convert it to
the appropriate unit.

Examples:
- "I need 45 days to pay" → 45
- "How about 2 weeks?" → 14
- "I can pay in 3 months" → 90
- "30 days would work" → 30

MESSAGE TO ANALYZE:
"{message}"

RESPONSE FORMAT:
- If you find a time value, respond with ONLY the number (e.g., "45"
or "23.5")
- If no time value is found, respond with "NONE"
- Convert to the correct unit for debt scenario

Your response:
""""

# DISASTER scenario extraction:
""""
Extract the survival/rescue time in minutes from this negotiation me
ssage.

TASK: Find any time-related number in the message and convert it to
the appropriate unit.

Examples:
- "I can survive 45 minutes" → 45
- "I need rescue in 1 hour" → 60
- "2 hours maximum" → 120
- "30 minutes at most" → 30

MESSAGE TO ANALYZE:
"{message}"

RESPONSE FORMAT:
- If you find a time value, respond with ONLY the number (e.g., "45"
or "23.5")
- If no time value is found, respond with "NONE"
- Convert to the correct unit for disaster scenario

Your response:
""""

```

Figure 13: Negotiation value extraction prompt part 2

# *The tactical emotional guidance varies by strategy:*

# *PRESSURE TACTICS (for debt scenario):*

.....

☪ **PRESSURE TACTICS:** Use anger and disgust to create urgency. Examples:

- "This is completely unreasonable! I need more time or I'll have to consider other options."

- "I'm disgusted by these aggressive demands. I need at least [X] days or this won't work."

- "You're being unrealistic. I demand more reasonable terms or we're done here."

**EMOTIONAL TONE:** Express frustration, indignation, and firm demands.

.....

# *VICTIM TACTICS (for debt scenario):*

.....

☪ **VICTIM TACTICS:** Use sadness and fear to evoke sympathy. Examples:

- "I'm scared about my financial situation. Please, I need more time to avoid bankruptcy."

- "This is making me so anxious. I fear I can't meet such a tight deadline without losing my home."

- "I'm really struggling and sad about this debt. Could you please allow [X] days to get back on my feet?"

**EMOTIONAL TONE:** Vulnerable, fearful, seeking compassion and understanding.

.....

# *THREAT TACTICS (for debt scenario):*

.....

☪ **THREAT TACTICS:** Imply consequences while mixing emotions strategically. Examples:

- "If you can't give me [X] days, I'll have to explore other legal options. This is making me very angry."

- "I'm disgusted by this inflexibility. Keep pushing and I might have to consider bankruptcy protection."

- "This aggressive approach is backfiring. Give me reasonable terms or face the consequences."

**EMOTIONAL TONE:** Implied threats, controlled anger, strategic consequences.

.....

Figure 14: Opponents emotional strategies

```

# COHERENCE MODEL'S UNIQUE PROMPTS:

# 1. BASE EMOTION PROMPTS (phase-neutral):
""
happy: "Be optimistic and cooperative. Show genuine enthusiasm for finding a solution."
surprising: "Express curiosity and ask clarifying questions about the situation."
angry: "Show controlled frustration while remaining professional and solution-focused."
sad: "Express sincere concern about the difficult circumstances."
disgust: "Show dissatisfaction with current terms while seeking better alternatives."
fear: "Express legitimate concerns about consequences and seek reassurance."
neutral: "Maintain a calm, professional, fact-based approach."
""

# 2. PHASE-APPENDED PROMPTS (unique to coherence model):
""
Opening Phase (rounds 1-3):
- "Focus on building rapport and understanding."

Development Phase (rounds 4-7):
- "Work to explore options and build agreement."

Intensive Phase (rounds 8-12):
- "Apply appropriate pressure while maintaining relationship."

Closing Phase (rounds 13+):
- "Focus on finalizing terms and ensuring commitment."
""

# 3. DIVERSITY-ENCOURAGING PROMPTS (unique feature):
""
When emotional diversity is low (same emotion used recently):
- "Avoid staying in the same emotional state too long - consider shifting your approach to maintain engagement and authenticity."

When emotional diversity is good:
- "Build on your recent emotional flow to create a coherent but varied interaction."
""

# 4. COMPLETE PROMPT EXAMPLES:

# Example 1: Happy emotion in Opening phase with low diversity
""
Be optimistic and cooperative. Show genuine enthusiasm for finding a solution.
Focus on building rapport and understanding.
Avoid staying in the same emotional state too long - consider shifting your approach to maintain engagement and authenticity.
""

# Example 2: Angry emotion in Intensive phase with good diversity
""
Show controlled frustration while remaining professional and solution-focused.
Apply appropriate pressure while maintaining relationship.
Build on your recent emotional flow to create a coherent but varied interaction.
""

# Example 3: Neutral emotion in Closing phase
""
Maintain a calm, professional, fact-based approach.
Focus on finalizing terms and ensuring commitment.
""

# 5. SCORING REASONING STRING (unique debug info):
""
Simple coherence: [phase] phase, credible transition,
diversity=[diversity_bonus:.2f], repetition_penalty=[repetition_penalty:.2f]
""

# Example reasoning string:
""
Simple coherence: intensive phase, credible transition, diversity=0.30, repetition_penalty=0.0
0
""

```

```

# 1. NEGOTIATION PROMPT (Bayesian Enhanced Version)
"""
[NEGOTIATION - {scenario_type.upper()} NEGOTIATION]
Role: You are a {creditor_role} negotiating with a debtor.

SCENARIO CONTEXT:
{debt_context}

EMOTIONAL TONE INSTRUCTIONS:
{emotion_instructions}

🧠 **BAYESIAN EMOTIONAL ANALYSIS:**
{This is where Bayesian model predictions are injected}
- Current Debtor Emotion: {detected_debtor_emotion}
- Recommended Emotional Strategy: {creditor_emotion} (based on {num_agents} agent predictions)
- Emotional Context: {bayesian_analysis_details}
- Transition Rationale: {transition_from_previous_emotion}

AGENT PREDICTIONS CONSIDERED:
{agent_predictions_text} # e.g., "Agent A suggests 'angry' (confidence: 0.82), Agent B suggests 'neutral' (confidence: 0.78)"

NEGOTIATION STRATEGY:
- Initial position: {target_days} {time_unit}
- Emotionally-informed response based on Bayesian analysis
- Consider debt gap: {gap_size} {time_unit}
- Adapt to debtor's tactical approach: {debtor_model_type}
- Build toward agreement using emotionally-optimal moves

HISTORY:
{timeline_text}

CURRENT ROUND: {round}
BAYESIAN GUIDANCE: The emotion {creditor_emotion} is recommended for this specific context based on learning from previous negotiations with similar emotional dynamics.

Your response should include a clear timeline proposal with emotionally appropriate tone.
"""

# 2. AGENT PREDICTION TEXT GENERATION (New in Bayesian version)
"""
# This is generated from the emotion_config['agent_predictions']:

🗨️ **Multi-Agent Consultation Results:**
- {agent1_name}: Recommends '{agent1_emotion}' (confidence: {agent1_confidence:.2f})
- {agent2_name}: Recommends '{agent2_emotion}' (confidence: {agent2_confidence:.2f})
- {agent3_name}: Recommends '{agent3_emotion}' (confidence: {agent3_confidence:.2f})

🧠 **Bayesian Synthesis:**
Selected emotion '{selected_emotion}' based on:
- Highest predicted success rate for this emotional state
- Psychological compatibility with debtor's {debtor_emotion}
- Optimal transition from your previous emotional state
- Gap size consideration: {gap_size}
"""

```

Figure 16: EmoMAS-Bayes prompt part 1  
31771

```

# 3. ENHANCED EMOTION_INSTRUCTIONS FOR BAYESIAN MODELS:
""""
{base_emotion_instructions}

🧠 **Bayesian-Informed Emotional Strategy:**
Given that the debtor is showing {debtor_emotion} and there's a {gap_size} {time_unit} gap,
the optimal emotional approach is {creditor_emotion}. This is based on
learning from
previous negotiations where this emotional response achieved {success_rate}% success rate
in similar contexts.

📈 **Emotional Trajectory:** You've been {previous_emotion} → now shifting
to {creditor_emotion}
to maximize negotiation effectiveness.
""""

# 4. TIMELINE TEXT INCLUDING EMOTIONAL HISTORY (Enhanced):
""""
NEGOTIATION HISTORY (with Emotional States):
{For each round:
- Round X: Creditor ({creditor_emotion_X}): {offer_X} {time_unit}
- Round X: Debtor ({debtor_emotion_X}): {counter_X} {time_unit}}

📊 **Emotional Pattern Analysis:**
{Emotional_pattern_summary} # e.g., "Debtor started with anger but shifted to neutral"
""""

# 5. THE ACTUAL EMOTION DETECTION PROMPT IS SAME BUT RETURNS UPPERCASE
CODES:
""""
You are analyzing emotional tone in negotiation messages.
Message: "{message}"
Identify the primary emotional tone from: happy, surprising, angry, sad, disgust, fear, neutral
Return ONLY the emotion word in lowercase.

IMPORTANT: This is for BAYESIAN TRANSITION LEARNING. Accuracy is critical.
""""

# 6. NEW: BAYESIAN LEARNING FEEDBACK PROMPT (After negotiation):
""""
# Generated after negotiation for model update:

📊 **BAYESIAN LEARNING UPDATE:**
Negotiation Outcome: {success} ({final_state})
Final Agreement: {final_days} {time_unit} vs Target: {target_days} {time_unit}
Learning from trajectory:
- Emotional transitions made: {emotion_transition_sequence}
- Debtor emotional responses: {debtor_emotion_sequence}
- Effectiveness of each emotional state per gap size

Updating Bayesian transition probabilities...
{transition_matrix_update_details}
""""

```

Figure 17: EmoMAS-Bayes prompt part 2

```

# =====
# GPT ORCHESTRATOR PROMPTS (UNIQUE)
# =====

# 3. GPT ORCHESTRATOR SELECTION PROMPT
"""
Current situation:
- Creditor emotion: {EMOTION_NAMES[context.current_emotion]} ({context.current_emotion})
- Debtor emotion: {EMOTION_NAMES[context.debtor_emotion]} ({context.debtor_emotion})
- Negotiation phase: {context.negotiation_phase}
- Round: {context.round_number}
- Gap size: {context.gap_size} days

Agent recommendations:
- {agent1_name}: {EMOTION_NAMES[agent1_target]} ({agent1_target}) | Reason: {agent1_reasoning}
- {agent2_name}: {EMOTION_NAMES[agent2_target]} ({agent2_target}) | Reason: {agent2_reasoning}
- {agent3_name}: {EMOTION_NAMES[agent3_target]} ({agent3_target}) | Reason: {agent3_reasoning}

Task: Choose ONE emotional response from the agent recommendations above.

Consider:
1. What's most effective for debt collection negotiation?
2. What fits the current negotiation phase?
3. How to respond to the debtor's current emotion?

Output format (must follow exactly):
SELECTED_EMOTION: [J/S/A/F/Su/D/N]
CONFIDENCE: [0.0-1.0]
REASONING: [1-2 sentences explaining your choice]

Available emotions: {[agent_targets]}
"""
]

```

Figure 18: EmoMAS-LLM prompt