Towards AI-Powered Real-Time Negotiation Agent

Martin Žust[1,*], Marko Grobelnik[1], Adrian M. Grobelnik[1] and Abdul Sittar[1]

[1]Jožef Stefan Institute, Ljubljana, Slovenia

Abstract

Negotiation is a fundamental human skill critical to personal, business, and societal interactions, yet many lack the expertise to negotiate effectively, leading to significant economic losses. Research shows unskilled negotiators in industries lose large amounts of value due to inadequate real-time guidance, while existing AI tools focus on automation rather than dynamic support. This introductory study aims to help unskilled negotiators with accessible, real-time assistance to improve outcomes across diverse contexts. Here we introduce a web-based negotiation agent that combines AI's analytical power with human intuition, providing live recommendations during simulated conversations. Unlike prior tools limited to post-event analysis or structured tasks, our agent transcribes dialogue, builds a dynamic world model, and offers tailored advice beyond what standalone AI or humans typically achieve. This hybrid approach reveals that integrating AI into live negotiations can bridge the efficacy gap, making proficiency more attainable than previously thought with static or isolated solutions. Broadly, it suggests a shift toward collaborative AI-human systems to address possibly complex interpersonal challenges.

1. Introduction

Negotiation is fundamental to human interaction, shaping outcomes in personal, business, and societal contexts, but many lack the skills to negotiate effectively. This gap leads to significant economic losses: US SaaS companies lose $2.08 billion annually, FinTech firms $1.36 billion, and the fitness industry $1.373 billion due to negotiation failures [1, 2, 3]. These stakes underscore the need for accessible tools to enhance negotiation proficiency, especially in real-time settings requiring quick decisions. While AI excels in structured tasks, its limitations in soft skills like empathy and adaptability [4, 5] suggest a hybrid approach combining AI's analytical precision with human intuition could optimize outcomes. Existing tools—e.g., Statworx, SirionLabs, and Lindy for contracts [6, 7, 8], or Gong.io and Chorus.ai for post-event analysis [9, 10]—focus on automation or retrospective insights, not real-time negotiation support. Research shows human-AI collaboration outperforms standalone efforts in sales and decision-making [11], inspiring our proposed negotiation agent. Integrated into platforms like Zoom and Google Meet, it transcribes dialogue, models the counterpart's perspective, and offers live, tailored advice based on user goals. Figure 1 illustrates this workflow. By merging AI and human strengths, this accessible solution aims to democratize negotiation expertise and improve efficacy across diverse contexts.

2. Problem Statement

Negotiation is a critical skill where an individual, say A, aims to reach an agreement o with a counterpart B that satisfies both parties' goals, denoted G_A and G_B (e.g., price, terms). Despite clear objectives, unskilled negotiators often struggle to achieve favorable outcomes, resulting in agreements that fall short of their minimum expectations, θ_A.

The core problem is that unskilled negotiators receive no support to interpret B's responses or adjust their approach mid-dialogue, represented as $D(t)$, the conversation state at time t. Without assistance, A struggles to build an accurate understanding of B's perspective, M_B, or to propose offers that align with

SemGenAge: 1st Workshop on Semantic Generative Agents on the Web at ESWC 2025; Workshops and Tutorials Joint Proceedings

*Corresponding author.
marti.zust@gmail.com (M. Žust)
https://www.twon-project.eu/semgenage25 (M. Žust)

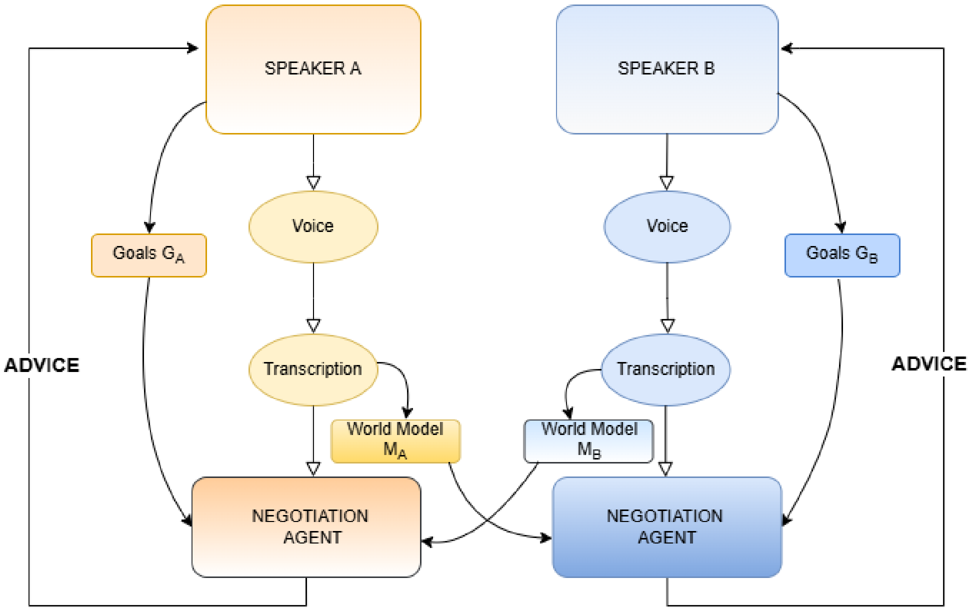

Figure 1: A diagram illustrating the negotiation agent's workflow.

G_A while meeting θ_B, B's acceptability threshold. Existing tools fall short here: they either automate structured tasks like contract drafting or analyze conversations after the fact, missing the opportunity to provide live, tailored recommendations. Moreover, AI alone cannot fully handle the soft skills—empathy, trust, adaptability—essential for reading B and responding effectively in real time. This leaves a critical need for a system that delivers strategic advice during $D(t)$, combining AI's analytical power with human judgment to improve outcomes and reduce losses.

3. Solution

In this work, we introduce a forthcoming web application[1] designed to showcase the capabilities of a novel negotiation agent. This application serves as a proof of concept, not yet as a production-ready tool, and focuses on simulating and analyzing a Zoom-like conversation between two individuals, A and B. By integrating our negotiation agent, it aims to demonstrate how AI can enhance negotiation outcomes, particularly in real-time settings, offering a practical testbed for achieving an agreement o that aligns with the goals G_A and G_B.

The application includes a test interface where users can upload a single video featuring a two-person conversation, similar to a YouTube podcast, with the active speaker displayed on screen. Users then use a graphical tool to segment this video into distinct sections, marking when A or B is speaking. Once segmented, the system presents two synchronized video streams side by side, replicating the layout of Zoom or Google Meet, providing a familiar visual context for the simulated negotiation dialogue $D(t)$.

Before the conversation can be played, users specify negotiation goals, G_A and G_B, for both participants. This feature allows the application to simulate the negotiation agent's functionality from either perspective—though in a production version, it would focus solely on A's goals, G_A. This setup enables us to explore how the agent interprets and responds to differing objectives, laying the groundwork for real-world applications where A seeks to meet θ_A while accommodating θ_B.

During playback, the speaking participant (e.g., A) is highlighted with a blue border to clearly indicate who is active, while the non-speaking participant's video (e.g., B) freezes at their last frame before pausing. After each monologue—or every 30 seconds for longer segments—the system generates

[1]Source code available at: Backend — https://github.com/MartinZust123/NegotiationAgentBackend, Frontend — https://github.com/MartinZust123/NegotiationAgentFrontend

a transcription of $D(t)$. Users can click a "Transcription" button beneath each speaker to view all monologues along with their start and end times, offering a detailed record of the dialogue at time t.

The application further leverages these transcriptions and specified goals, G_A and G_B, to build and update a world model, M_A or M_B, for each speaker after every transcribed segment. A "World Model" button displays this evolving model, reflecting the speaker's perspective and intent (e.g., M_B for B). Additionally, a "Give Advice" button activates the negotiation agent, which provides tailored recommendations based on G_A and the current state of $D(t)$. Together, these features illustrate the agent's potential to support A in dynamic, interpersonal scenarios, ensuring o meets or exceeds θ_A.

4. Experiments

4.1. Overview

This section outlines two experiments showcasing possible interactions conducted via web-based negotiation agent, designed to assist users like A in real-time conversations. Using a simulated car sale negotiation and simulated buying of an apartment, we demonstrate how the agent transcribes dialogue $D(t)$, builds a world model M_B, and provides strategic advice, testing its ability to enhance outcomes for unskilled negotiators aiming to achieve o above θ_A.

4.2. Experiment 1: Car Sale

4.2.1. Experimental Setup

The application's front-end is developed in React with Bootstrap as the CSS library to enhance user experience, while the back-end is built in Python using Flask for seamless front-end communication. Audio transcription of $D(t)$ is handled via Whisper. The negotiation agent leverages OpenAI's GPT-4o for constructing M_A or M_B and generating advice, selected for its robust natural language processing capabilities, with plans to explore optimal models in future iterations.

4.2.2. Experimental Design and Procedure

We simulate a negotiation between a seller A and buyer B over a used car, uploaded as a single video where the active speaker is displayed. The seller's goals, G_A —avoiding rejection and securing a price above \$5000 (i.e., θ_A = \$5000)—are defined pre-negotiation. The video is manually segmented into monologues (A or B speaking), with transcriptions of $D(t)$ generated at the end of each segment or every 30 seconds for segments exceeding that duration. The system then displays two synchronized streams, mimicking Zoom, with the active speaker (e.g., A) highlighted by a blue border and the other (e.g., B) frozen at their last frame. An example dialogue unfolds as follows:

> **Seller:** "This car is in **excellent condition** with a new set of tires. I'm asking for **\$6000**, but I can be flexible."
> *(The seller opens with a strong pitch, setting the initial price but signaling room for negotiation.)*
> **Buyer:** "That sounds good, but my budget is around **\$4500**."
> *(The buyer responds with a lower counteroffer, testing the seller's flexibility.)*
> **Seller:** "I understand. Would you be willing to go up to **\$5200** if I include a free oil change?"
> *(The seller lowers the price and adds an incentive to make the offer more attractive.)*
> **Buyer:** "That makes it more appealing, but I'd need to think about it."
> *(The buyer shows interest but hesitates, keeping the negotiation open.)*
> **Seller:** "How about we finalize the deal now at **\$5100** with an oil change included?"
> *(The seller pushes to close the deal with a slightly better offer.)*
> **Buyer:** "Alright, I think we have a deal."
> *(The buyer agrees, concluding the negotiation successfully at \$5100 plus an oil change.)*

Throughout, the agent transcribes each exchange in $D(t)$, updates M_A and M_B, and delivers advice via interface buttons to ensure o aligns with G_A and exceeds θ_A.

4.2.3. How the Agent Assists the Seller

The negotiation agent supports A by analyzing B's responses in $D(t)$ and aligning strategies with G_A. After transcribing B's initial \$4500 offer, it updates M_B, noting budget constraints and a positive reaction to incentives. Early advice, accessed via the "Give Advice" button, suggests, "Continue offering small incentives without lowering the price significantly." As B hesitates at \$5200, the agent refines its recommendation to, "Propose a minor price drop if resistance persists," leading to the successful \$5100 deal, exceeding θ_A. This dynamic adjustment showcases the agent's real-time utility in guiding A toward o.

4.3. Experiment 2: Apartment Sale

The second example (Figure 2) illustrates a dialogue between two actors, depicting alternative negotiation scenarios. These generated scenarios provide a potential preview of how the dialogue might unfold in a real-life situation. As shown in the figure, the scenarios, developed across three levels, can result in various outcomes: a successful agreement, rejection by either party, or a situation requiring further negotiation. For the states in the negotiation diagram that require further negotiations, we can generate follow-up diagrams until the end states (successful agreement or rejection) are reached.

4.3.1. Experimental Setting

The simulated discussion with Google Gemini 2.0 model involves two actors: *Marko (Buyer)* and *Martin (Seller)*, regarding the sale of an apartment. The initial asking price is \$500,000. The negotiation is multidimensional and revolves around the following four items:

- Apartment Price: Starting at \$500,000.
- Kitchen Renovation: Whether the kitchen will be renovated before the sale.
- Bathroom Renovation: Whether the bathroom will be renovated before the sale.
- Furniture: Whether the apartment is sold furnished or unfurnished.

5. Discussion

Existing research highlight that optimal performance is achieved when AI and human negotiators collaborate [12, 13]. While both humans and large language models (LLMs) have their respective limitations—humans may struggle with data processing and consistency, whereas LLMs lack contextual awareness and emotional intelligence — their combined strengths lead to significantly improved negotiation outcomes. By integrating AI into the negotiation process, we leverage the speed, data processing power, and strategic recommendations of LLMs while maintaining the human ability to interpret context, emotions, and social cues.

A significant portion of modern business negotiations, price discussions, and strategic conversations now take place over virtual platforms such as Zoom and Google Meet. This shift toward online communication presents an opportunity to enhance negotiations with AI-driven agents that provide real-time assistance without disrupting the natural flow of conversation. Our research demonstrates that an AI-based negotiation assistant can be integrated seamlessly into virtual discussions, offering valuable insights and recommendations while remaining unobtrusive.

A key finding from our test application is that our negotiation agent is highly user-friendly. Through our demonstrations, we observed that the agent does not introduce significant distractions, allowing users to remain engaged in their discussions. Users can maintain their focus on the negotiation at hand

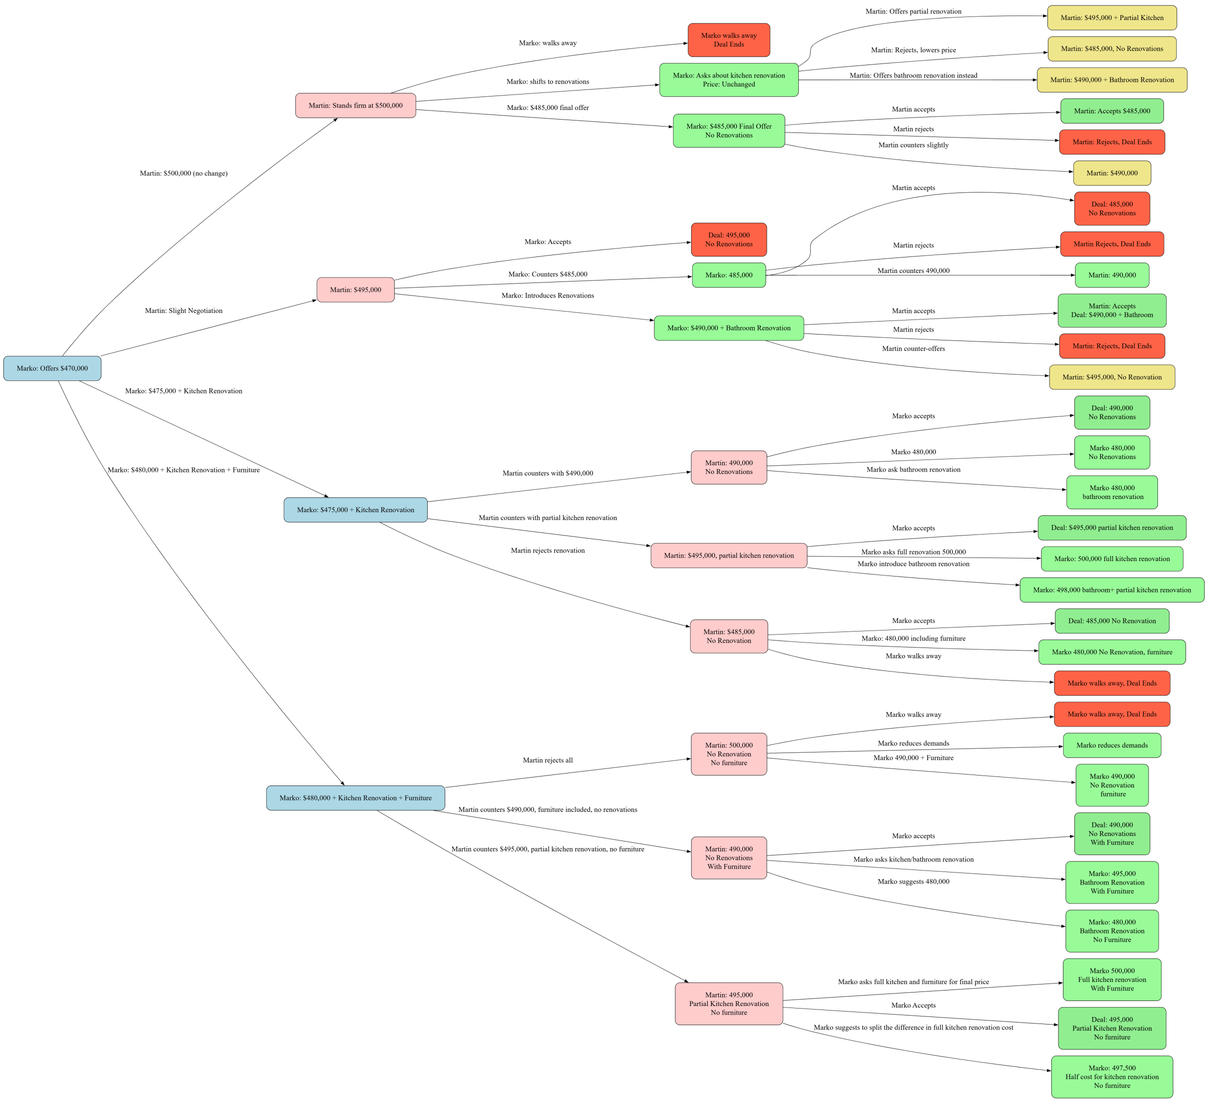

Figure 2: A diagram illustrating alternative scenarios between two agents negotiating the purchase of an apartment. Initial Positions: Agent 1 'Marko': Wants a lower price, renovations, and furniture; Agent 2 'Martin': Wants $500,000, no renovations, and no furniture.

and call upon the AI agent only when they require updated information, strategic advice, or contextual analysis of the ongoing conversation.

Despite these advantages, adoption remains a challenge. Even if AI assistance improves negotiation outcomes, individuals may be reluctant to integrate it into their workflow unless there is a compelling need or a clear pain point that justifies its use. Resistance to change, concerns over trust in AI-generated advice, and the learning curve associated with new technologies may all impact adoption rates.

6. Future Work

Persona-based testing will be a key focus of future research to evaluate the negotiation agent's performance. We plan to compare scenarios such as persona vs. persona, persona vs. persona with the negotiation agent, and persona vs. negotiation agent alone to assess how the agent influences outcomes across controlled settings. For personas we could use ones created for project TWON.

Real-world comparisons will further validate the agent's impact by pitting users with the negotiation agent against those without it. These experiments will aim to quantify improvements in negotiation success, user confidence, and economic outcomes, providing concrete evidence of the agent's value in practical applications.

Future research will also focus on enhancing the negotiation agent's adoption and effectiveness. To encourage user uptake, we aim to improve the interpretability of AI recommendations, demonstrate tangible value through real-world case studies, and better align the agent with human decision-making processes. Additionally, we plan to expand testing by exploring diverse negotiation scenarios to identify where the agent performs best, while gathering human feedback through user critiques and recommendations to refine its functionality.

The current proof-of-concept application has limitations, including its reliance on manual video segmentation, restriction to single-speaker display videos, and lack of integration with streaming platforms like Zoom or Google Meet. Future work will address these by developing automated segmentation tools, supporting multi-display inputs for multiple participants, and pursuing partnerships with Zoom and Google Meet to integrate the agent for public use. Staying abreast of advancements in transcription and data analysis tools will also ensure the agent evolves with the rapidly progressing AI landscape.

Declaration on Generative AI

During the preparation of this work, the authors used Grok 3 for rephrasing paragraphs in the introduction and discussion sections and Gemini for editing code. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

[1] Statista, Revenue of the Software as a Service (SaaS) Market in the United States from 2018 to 2028, Online Database, 2023. URL: https://www.statista.com/outlook/tmo/software/software-as-a-service/saas/united-states, provides baseline SaaS revenue data for loss estimates.

[2] Statista, FinTech Market Revenue in the United States from 2018 to 2028, Online Database, 2023. URL: https://www.statista.com/outlook/fmo/fintech/united-states, provides baseline FinTech revenue for loss estimates.

[3] R. . S. A. I. International Health, IHRSA 2023 Fitness Industry Report: Key Statistics and Trends, Industry Report, 2023. URL: https://www.ihrsa.org/publications, provides membership churn data for fitness loss estimates.

[4] E. Brynjolfsson, A. McAfee, The Business of Artificial Intelligence: What It Can—and Cannot—Do for Your Organization, Harvard Business Review (2017). URL: https://hbr.org/2017/07/the-business-of-artificial-intelligence, highlights AI's limits in soft skills domains.

[5] M.-H. Huang, R. Rust, V. Maksimovic, The Feeling Economy: Managing in the Next Generation of Artificial Intelligence, MIS Quarterly Executive 18 (2019) 251–268. Demonstrates AI's challenges with emotional nuance in service contexts.

[6] Statworx, Leveraging AI for Strategic Negotiation Analytics, Industry Report, 2024. URL: https://www.statworx.com/en/resources/.

[7] SirionLabs, AI-Powered Contract Management: Enhancing Negotiation Outcomes, White Paper, 2023. URL: https://www.sirionlabs.com/resources/white-papers/, highlights focus on contract automation, not real-time advice.

[8] Lindy, AI Assistants for Everyday Negotiation: Scheduling and Beyond, Product Overview, 2024. URL: https://www.lindy.ai/resources, notes basic task support, not dynamic negotiation.

[9] Gong, Gong.io: Conversation Intelligence Platform, Company Website, 2024. URL: https://www.gong.io, emphasizes post-call analysis, not real-time support.

[10] Chorus, Chorus.ai: Call Recording and Analysis, Company Website, 2024. URL: https://www.chorus.ai, highlights post-event insights, not during calls.

[11] P. R. Daugherty, H. J. Wilson, Human + Machine: Reimagining Work in the Age of AI, Harvard Business Review Press, Boston, MA, 2018. Shows hybrid AI-human superiority in soft skills tasks like sales.

[12] J. Weisz, et al., AI as a Collaborative Partner: Enhancing Productivity in Software Engineering, in: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, ACM, New York, NY, 2021, pp. 1–13. doi:10.1145/3411764.3445298.
[13] M. Ciaschi, et al., Exploring the Role of AI in Assessing Soft Skills: A Collaborative Approach, 2024.