"ChatGPT, Don't Tell Me What to Do": Designing AI for Context Analysis in Humanitarian Frontline Negotiations

Zilin Ma
Harvard University
Allston, MA, USA
zilinma@g.harvard.edu

Yiyang Mei
Emory University
Atlanta, GA, USA
yiyang.mei@emory.edu

Claude Bruderlein
Harvard University
Cambridge, USA
cbruderl@hsph.harvard.edu

Krzysztof Gajos
Harvard University
Allston, MA, USA
kgajos@g.harvard.edu

Weiwei Pan
Harvard University
Allston, MA, USA
weiweipan@g.harvard.edu

Abstract

Frontline humanitarian negotiators are increasingly exploring ways to use AI tools in their workflows. However, current AI-tools in negotiation primarily focus on outcomes, neglecting crucial aspects of the negotiation process. Through iterative user-centric design with experienced frontline negotiators (n=32), we found that flexible tools that enable contextualizing cases and exploring options (with associated risks) are more effective than those providing direct recommendations of negotiation strategies. Surprisingly, negotiators demonstrated tolerance for occasional hallucinations and biases of AI. Our findings suggest that the design of AI-assisted negotiation tools should build on practitioners' existing practices, such as weighing different compromises and validating information with peers. This approach leverages negotiators' expertise while enhancing their decision-making capabilities. We call for technologists to learn from and collaborate closely with frontline negotiators, applying these insights to future AI designs and jointly developing professional guidelines for AI use in humanitarian negotiations.

CCS Concepts

• **Human-centered computing → User studies**.

Keywords

Frontline Negotiation, AI, Design Probe

ACM Reference Format:
Zilin Ma, Yiyang Mei, Claude Bruderlein, Krzysztof Gajos, and Weiwei Pan. 2025. "ChatGPT, Don't Tell Me What to Do": Designing AI for Context Analysis in Humanitarian Frontline Negotiations. In *CHIWORK '25: Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25), June 23–25, 2025, Amsterdam, Netherlands.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3729176.3729194

CHIWORK '25, Amsterdam, Netherlands

ACM ISBN 979-8-4007-1384-2/25/06
https://doi.org/10.1145/3729176.3729194

1 Introduction

Humanitarian negotiations in conflict zones, or frontline negotiations, are vital for securing access to crisis-affected populations for aid delivery [24, 42]. Frontline negotiations is a unique type of work that is often adversarial, complex, and high-risk, involving stakeholders with diverse geographic, political, and cultural backgrounds. A major challenge in this workflow is the synthesis of unstructured information—such as interviews, stakeholder meetings, and historical documents—quickly and accurately [e.g. 18, 92]. With the rise in violent conflicts and the increasing demand for frontline negotiations globally [35, 36, 45, 78], practitioners are exploring AI to ease the growing pressure on human experts. These efforts, however, remain informal without collaboration with machine learning or HCI experts. For instance, in 2024, Frontline Associates—a global network of negotiation practitioners—hosted an educational summer program on AI's role in negotiation [37]. Over 100 humanitarian negotiators from organizations like Médecins Sans Frontières, the International Committee of the Red Cross, and the United Nations participated, experimenting with tools like ChatGPT for context analysis and information synthesis.

AI-assisted decision-making raises broader concerns for automation at work. Prior research on AI in professional settings suggests that deskilling [64, 102], shifts in occupational identity [102], and overreliance on automation [34, 75] can significantly alter workplace dynamics. Negotiators, who rely on experiential knowledge and adaptive reasoning, may face challenges integrating AI in ways that preserve their expertise rather than diminishing their agency. Given the high-stakes nature of frontline negotiation, LLM limitations—such as bias, confidentiality risks, and hallucinations [13, 33, 54, 63, 88, 94, 101]—raise concerns about responsible AI use in this line of work.

Given the potential harms of ad-hoc LLM adoption in frontline negotiation, it is essential to assess AI's impact not only on the decision-support needs of workers but also on the profession as a whole. While there is existing research on the application of AI in negotiation, they target negotiation contexts that are not aligned with the constraints and goals of humanitarian negotiation, such as zero-sum settings in contracts negotiation. These current approaches are often goal-centric, focusing on using statistical benchmarks to achieve high accuracy. The emphasis is on work automation rather than assistance; and these works largely seek to "surpass human performance", potentially replacing human negotiators [1, 62, 72].

Frontline negotiation, however, is a line of work that demands human rapport and requires strategy contextualization. Frontline negotiators need to keep abreast of the emotions, cultural backgrounds, and human irrationality involved in the negotiation process [24, 25, 28, 32]. This focus on the human element aligns with the approaches in human-AI collaboration – particularly the"**Process-Oriented Support**" [105]. Process-oriented support is a way to design AI tools that help workers through each step of their job rather than just delivering a final answer. This concept is particularly relevant in complex, high-stakes workplaces, where decision-making involves multiple steps and requires integrating diverse pieces of information [105]. Given the nuance and high stakes involved in frontline negotiation, process-oriented support may be more appropriate for this work than the current goal-centric approach.

We ask the following research questions: **RQ1:** What decision support needs do negotiators have that can be assisted by LLMs? **RQ2:** What concerns do negotiators have over using LLMs in negotiation? **RQ3:** Do current LLM tools assist negotiators' key decision support needs? **RQ4:** What are the LLMs' anticipated impacts on the work of frontline negotiators?

We first conducted a formative study, with 14 frontline negotiators, where we uncovered the major decision-support needs in frontline negotiation, and negotiators' concerns over using LLMs in negotiation. These decision-support needs include: context analysis, compromise ideation, risk analysis and knowledge sharing.

Then we sought to understand how the **actual usage** of LLMs can assist and impact negotiators' work, and if they can address the needs that we discovered in the formative study. Thus, we designed a probe interface, informed by the user needs identified in the formative study, to elicit negotiators' grounded reactions to an AI-powered decision-support tool built on a process-oriented paradigm. As opposed to typical negotiation AI support tools, which emphasize outcomes by automating the work of negotiation, this interface emphasizes AI as a collaborator that supports workers in reaching their own decisions. This way, we presented negotiators with a different kind of a worker-AI relationship, where the users were not mere supervisors of AI decisions. Through this user study asking the participants to use ChatGPT and this probe, we explored how negotiators perceived AI's role in their work, its alignment with their professional expertise, and its potential impact on existing negotiation practices.

We made the following contributions:

- A formative study that discovers the obstacles and needs of frontline negotiators when conducting a negotiation.
- Through process-oriented support, an implementation of a design probe that assisted *existing negotiation practices*, including context analysis, compromise ideation, and risk analysis.
- A qualitative study that involves ChatGPT and the design probe, demonstrating that negotiators prefer a process-oriented support interface that leverages negotiators' existing practices, potentially reducing errors that novice negotiators may make due to lack of negotiation experience.

Finally, our work with the negotiation community points to the need for "worker development" – helping workers build technological fluency to use AI tools effectively and responsibly. The complexity of LLM technologies poses barriers for non-technologically trained workers. We see our paper as a call to action for collaborations between technologists and practitioners to facilitate knowledge about AI capabilities and limitations and support informed guidelines for ethical AI use. Current understanding of AI's impact on the work of frontline negotiation remains limited. Our study showed that existing AI-driven negotiation tools primarily focus on generating outcomes, often without integration into workers' workflows. If not carefully designed, LLMs risk automating relational, ethical, and creative aspects of negotiation, potentially diminishing the overall quality of negotiations. While future model advancements may enable LLMs to enhance emotional intelligence and relationship-building, negotiators are unlikely to accept AI as a replacement for the core of their work. Further research is needed to define the complementary roles of AI and human negotiators, ensuring that AI supports, rather than undermines, the complexities of negotiation practice. Our work suggest that in emotionally charged, high-stakes workplaces where human judgment, relationship-building, and adaptability are essential—such as crisis response, social work, and labor mediation—AI should be designed to support workers' expertise rather than automate complex, context-sensitive decisions.

2 Background and Related Work

2.1 Characterizing Challenges in the Work of Humanitarian Frontline Negotiation

International law [7] requires governments to provide assistance and protection to those living within their jurisdiction. When governments fail to meet this obligation, humanitarian organizations address these violations [25]. Negotiators work to secure assistance and build cooperation among key stakeholders, despite power asymmetries and ideological differences [71, 94].

Despite the importance of their work, the negotiators often come from a position of weakness, as they need to rely on humanitarian principles or international laws that sometimes mean little to the counterparties. Moreover, they often deal with armed groups while unarmed or negotiate with local governments supported by military forces [24, 68].

Due to power asymmetry and ideological differences between the negotiating parties, frontline humanitarian negotiators often struggle to reach principled agreements or accept compromises. Humanitarian negotiations often rule out finding a middle ground, as doing so may involve intolerable concessions or satisfying illegal and immoral interests. Negotiators frequently encounter situations where compromises would lead to outcomes that violate core humanitarian principles [24].

To navigate these challenges, negotiation frameworks such as Island of Agreements, Iceberg CSS, and Red Line/Bottom Line (see Section 5.2.1) help structure key positions and contexts, relying on practitioners' own expertise rather than rigid instructions [24]. Without detailed step-by-step instructions, they emphasize judgment over rigid guidelines. Usually, these tools require the negotiator to read massive amounts of documents and hours of meeting transcripts to successfully produce these summaries. Especially in high-stakes environments, negotiators may struggle to use these frameworks, as the need for quick decisions and adaptability often takes precedence.

2.2 The Impact of LLMs on Knowledge Work

LLMs are language models that can generate human-like text, answer questions, assist in writing [41], help with coding [99], translate languages [47], and tutor in various areas [23, 77]. Recent general LLMs, such as GPT-series [17, 79], LLaMA [96], and PaLM [2], have advanced Natural Language Processing (NLP) capabilities, including text generation and summarization.

In many professions, LLMs appear to be influencing a shift from expertise-based roles to ones that involve overseeing or interpreting AI-generated outputs [102]. For instance, call center agents may be increasingly expected to mediate between AI-generated responses and customers [19], while software engineers are often encouraged to review and refine AI-suggested code rather than writing it from scratch [99]. Although this shift is frequently described as enhancing efficiency, existing studies also suggest that it can introduce additional cognitive demands when integrating AI-generated content [29, 91].

This shift raises concerns about deskilling, as increased reliance on AI-generated insights may diminish professionals' ability to apply their own judgment [64, 102]. In negotiation, AI-support tools could similarly automate key decision-making processes, reducing negotiators' reliance on their own skills. Over time, this could erode practitioners' ability to synthesize complex stakeholder dynamics, evaluate risks, and manage power imbalances, fundamentally altering negotiation as a profession. Additionally, as AI reduces the skill gap between novices and experts, expertise itself may become less valued. Evidence suggests that AI disproportionately benefits workers with lower initial proficiency—consultants at Boston Consulting Group, for example, saw a 43% performance improvement with AI assistance compared to just 17% for their more skilled peers [31]. If AI makes expertise less of a differentiator, frontline negotiation—once defined by deep expertise in managing stakeholder dynamics—could shift toward oversight of AI-generated analyses, diminishing the negotiator's role as an active decision-maker. However, evidence also suggested that with careful design, deskilling is not inevitable [21, 38].

While concerns about deskilling and shifting professional roles are well-founded, discussions on AI's impact often oscillate between fears of widespread job displacement and optimism about AI as a productivity tool. Sensational narratives [3, 26, 65, 70, 87, 98], such as claims that LLMs are "just about to replace everything" [3], have fueled anxieties about societal and labor disruption. A Pew survey found that about a fifth of workers hold "high-exposure jobs", where key activities could be replaced or assisted by AI [61]. Another survey revealed that 37% of U.S. respondents were "more concerned than excited" about AI in daily life, with one in five citing job loss as their primary concern [86]. Speculative claims about artificial general intelligence (AGI), likened to science fiction, envision scenarios like machines dominating workers or forming a new underclass of human [40, 65, 102]. As Woodruff et al. note, these discussions often overlook the perspectives of knowledge workers themselves, compounding uncertainty [102]. Interestingly, workers in high-exposure domains report that AI is more likely to help than harm them personally [61]. Workers anticipate outsourcing mundane tasks like note-taking to AI while retaining full control over their work, countering predictions of widespread workforce automation through generative AI [102]. While there is growing research on the general implications of LLMs, their specific impacts on high-stake domains such as frontline negotiation remain underexplored.

Deploying LLMs in high-stakes contexts poses unique challenges, including technoskepticism[84], hallucination risks[56], privacy concerns like Personally Identifiable Information leakage[60], and limited domain-specific evaluation benchmarks[81]. Ensuring ethical alignment and multilingual support is critical, as shown by efforts like BiMediX for bilingual translation [83] and India's Bhashini for language accessibility [74]. Advanced techniques such as citation aggregation[95] and RALL generation[49] offer potential solutions but remain underexplored in negotiation settings. As LLMs evolve, it is crucial to anticipate challenges, identify benefits, and develop strategies to support these transformative practices.

2.3 AI-Assisted Decision-Making in the Workplace

Our work adapts insights from human-AI decision-making research, particularly explainable AI (XAI), to the design of AI systems for negotiators and other frontline workers. Recent work on AI-assisted decision-making in a wide range of domains has been shaped by the explainable AI (XAI) paradigm [48]. In this paradigm, the AI offers a decision recommendation and some additional information (an explanation) that is meant to help the human decision maker understand how the decision recommendation was arrived at. It was initially assumed that this paradigm would enable human-AI teams to make better decisions than either people or AIs could make on their own [57, 58, 69]. However, subsequent empirical studies consistently showed that people assisted by XAI made less accurate decisions than the AI systems alone [8, 39, 52], possibly because users do not engage cognitively with AI-generated explanations in XAI settings [20, 38]. Some researchers have also pointed out that the XAI approach requires people to engage in backward reasoning (why is this recommendation a good or bad decision?) rather than a more natural forward reasoning process (given what is known, what is the best decision?) [105]. Yet others have higher level concerns arguing that through the act of offering a decision recommendation, algorithms take some authority away from the human decision makers (in many institutional settings there is an additional effort or liability burden if a human disagrees with the algorithm) without assuming any of the accountability [43].

In response to these concerns, alternative human-AI interaction paradigms have begun to emerge for AI-assisted decision making. Some examples include presenting people with syntheses of relevant information [38], supporting people in developing situational awareness [105], asking people targeted questions [30], or presenting people with systematic compilations of key arguments for and against each of the likely options [73].

At a high level, the conventional XAI approach uses AI to automate a human cognitive task: the AI recommends a decision and the human decision-maker's primary job is to supervise the automation and intervene only in the (presumably) rare situations when the AI did not get things entirely right. The alternative approaches, in turn, leave the final decision-making to the person while using AI to support the human decision-maker in some way: by surfacing relevant knowledge, helping them be ready to make a time-sensitive decision should an emergency arise, identifying key factors to be considered, or supporting systematic consideration of trade-offs.

2.4 AI in Negotiation Work

Reflecting the established trends toward automation in LLM-based applications and in AI-assisted decision-making, current AI research in negotiation largely focuses on goal-centric approaches, often overlooking the process (e.g., relationship-building) that leads to successful negotiation outcomes [105]. For example, some works argue that human emotion and rapport-building can hinder negotiation [27]. These goal-centric approaches often involve developing multi-agent simulations that aim to "surpass human performance" by negotiating on behalf of humans [1, 46, 53, 59, 62, 72].

In narrowly defined, zero-sum negotiations (e.g., contract negotiations), AI systems already attempt to replace human negotiators. For example, CICERO demonstrated "human-level performance" in the strategy game Diplomacy, conducting negotiations without human intervention [72]. Similarly, companies like Walmart, using Pactum's technology, automated supplier negotiations, achieving agreements with 64% of suppliers, far exceeding their 20% target [100].

However, these goal-centric approaches do not align with the realities of frontline negotiations, where emotions, cultural contexts, and irrationalities play crucial roles [24, 68]. Unlike AI systems designed for optimizing predefined objectives, frontline negotiators rely on deeply human skills—such as reading emotional cues, managing ambiguity, and adapting strategies dynamically—to build trust and secure cooperation [24, 25, 94]. If negotiation AI systems continue to prioritize outcome-driven automation, they risk undermining the core professional competencies that define negotiation work. Recent work at CHI WORK highlights these complexities. Studies on programming assistance reveal tensions between human agency and automation [75], while research on LLM-assisted ideation suggests process-oriented support is more effective than outcome-driven optimization [50]. While AI for frontline negotiation remains undeveloped, the current trend of chatbots that either instruct users or negotiate autonomously suggests a risk of misalignment unless research addresses the unique needs of this domain.

The complexities of frontline negotiations — emotional, cultural, and contextual - make it difficult for current AI systems to replicate human performance [10, 14]. This mismatch shows the need for alternative AI designs, such as **process-oriented support** that assist with the decision-making process in negotiations [103, 105]. By focusing on the negotiation process rather than just outcomes, such systems would better align with the needs of frontline negotiators, assisting them in navigating emotional, cultural and contextual factors.

We have built on all of the above insights in this project: we used the formative study to identify the key cognitive tasks that the frontline negotiators engage in and to understand what makes these tasks challenging. In our second study, we developed a system prototype to support the negotiators in some of the challenging aspects of one of the tasks while leaving the negotiators in charge of the key decisions. Throughout the project, we also paid particular attention to factors that make frontline negotiation an engaging and desirable occupation for our participants. In the second study, we strove to design a prototype in a manner that would enhance rather than automate away the aspects that the negotiators identified as core to their profession.

3 Formative Study: Understanding the Decision-Support Needs of Frontline Humanitarian Negotiation

3.1 Overview

As frontline humanitarian negotiations grow increasingly complex, practitioners are exploring how ChatGPT can support their work. One member of our team, an experienced frontline negotiator with decades of critical mission experience, worked with us in generating the research questions. Prior to our collaboration, this negotiator and peers in their community have already found ways in which ChatGPT can be helpful to their work, and are equally concerned that these tools can be harmful. Thus, this study focuses on answering the following three aspects: how negotiators are currently using LLMs, the outcomes they hope to achieve, and the concerns they have about applying these tools to their work. Because the negotiator co-author did not want their own views to disproportionately impact the results of the research, they did not participate in data collection or analysis in either of the two studies.

3.2 Survey

We reached out to humanitarian frontline negotiators through the email list of the Frontline Negotiators Network, inviting them to participate in a voluntary survey. The survey aimed to gather information on the participants' negotiation experience, frequency of computer usage for work, educational background, and the extent of AI usage in their professional activities, particularly in relation to frontline negotiation. The questions can be viewed in Appendix B

Before we finalized the list of negotiators to interview for this study, we collected 30 responses in total, with 14 negotiators scheduling an interview with the first author.

3.3 Semi-structured interviews

We conducted semi-structured interviews with 14 participants. Participants are from 9 different humanitarian organizations (e.g. International NGO Safety Organisation, Doctors Without Boarders), from 12 different countries (See Appendix E). Prior to conducting our interviews, we made sure each participant provided informed consent, during which we emphasized their right to withdraw from the study at any time if they felt uncomfortable. After completing the interviews, participants received a compensation of US $30 for their time. Interviews typically lasted 45 to 60 minutes.

The interviews were structured into three main segments: understanding the participants' negotiation workflow, discussing their experiences with AI tools like ChatGPT, and exploring potential areas where AI could further supported their negotiation processes. During the second stage where we asked the participants' experience with AI tools, we showed the participants a video of ChatGPT generating the Iceberg CSS (see Section 5.2.1 for a description of Iceberg CSS tool) framework to facilitate future-state ideation based on existing LLM capabilities. The prompts and outputs used can be found in the online supplementary files.

Detailed interview guidelines are provided in the Appendix C. Following each interview, conversations were transcribed, anonymized, and analyzed to extract key insights. Given the current state of research on AI-assisted decision-making (Section 2.3) and on LLMs in knowledge work (Section 2.2), we paid particular attention to two broad topics: 1. What are the major cognitive tasks that frontline negotiators engage in? What is challenging about these tasks? What opportunities exist for supporting negotiators' work while leaving them in charge of key decisions? 2. What are the major concerns about using AI in frontline negotiation? What mistakes would be considered catastrophic? What changes to the nature of the profession would be unacceptable to the negotiators?

3.4 Data analysis

The first author independently coded 5 interview transcripts using an open coding technique [22]. This initial analysis revealed overarching benefits, specific advantages for negotiators, and their primary concerns. Subsequently, three members of the research team convened to finalize a codebook for further analysis. The first author then proceeded to analyze the remaining transcripts, continuously refining the codebook based on new insights until data saturation was reached.

4 Results - Interview Study

This section summarizes the main findings of our interviews, focusing on the key cognitive tasks involved in frontline negotiations (and the emerging design opportunities) and concerns regarding the use of LLMs in humanitarian negotiations. We adopt negotiators' use of term "AI" to refer to LLMs.

4.1 Key Negotiation Tasks and Design Opportunities

4.1.1 Support the preparatory stage of negotiation. Negotiators analyze large datasets to understand context, party positions, shared and contested norms, motivations, and risks—ultimately forming a flexible negotiation strategy.

Context analysis: Negotiators often deal with long documents, unstructured texts, and the need to update their preparation in light of new information. *"Providing guidance and support on analysis is where AI can be a game changer. [AI can help finding] what are [different parties'] positions, interest, and needs."* (P10) In fact, some negotiators have already started using ChatGPT to summarize cases for them. However, prompting still poses challenges. Negotiators want more support on how to prompt LLMs: *"I am not a professional user for ChatGPT [...] if there is any official support [on prompting], I will be very happy."* (P2)

Ideating Compromises: Effective negotiation involves understanding potential compromises and their associated risks, which negotiators think that LLMs can assist by proposing alternative plans for negotiators to consider, helping them understand a wider range of options and their associated risks. One participant highlighted this potential: *"I just don't want to let my brain be single minded. [...] Maybe AI could be an eye opener for the other way around and proposing secondary solutions and options."* (P3) Furthermore, one negotiator suggested that LLMs could aid in *"Finding creative solutions and new positions that can satisfy all parties' needs."* (P10)

Risk analysis: In the context of risk assessment, negotiators believe LLMs can help identify information gaps that could lead to unforeseen risks in compromises. One participant noted, *"You could cross check [with AI] what you've got. This is what I've got on this counterpart. Is there anything out there I've missed? That would be very useful."* (P12)

4.1.2 Support knowledge sharing. The second broad task mentioned by the negotiators was synthesizing information about past negotiations to be shared with other negotiators. Knowledge sharing between negotiators includes training new team members and passing on insights from previous negotiations—especially those involving specific parties—to teams who will soon engage with them. P14 shared that in their decades of negotiation experience, they have *"lost 22 colleagues"* to violence and emphasized, *"With so many conflicts happening, negotiators don't have the luxury of 10 or 20 years to learn. I hope AI can help new negotiators build skills without facing hostile environments."*

Another (P10) highlighted the loss of institutional memory, noting, *"We lost incredible knowledge about negotiations from 20 years ago, even with the same people and challenges."* They stressed capturing strategies and common ground, not just outcomes: *"What I care about is the strategy, the common ground. A dataset that proposes past cases from various organizations would be invaluable."* However, need for keeping certain details of past negotiations confidential often limits shared learning: *"Maybe AI can be trusted enough to decontextualize these cases."*

4.2 Concerns Over LLM in Negotiation

4.2.1 Well-known LLM Limitations. Negotiators expressed concerns about well-known limitations of LLMs like ChatGPT, highlighting issues of confidentiality, Western biases, reliability, and the risk of overreliance.

Confidentiality is a major concern, as negotiators often handle sensitive information. Many feel inadequately informed about protecting privacy when using LLMs, and organizational guidelines are often vague. As one negotiator noted, *"My organization has put in place some kind of restrictions on the way we can use AI, but we're still at the beginning stages of understanding the capabilities and potential consequences. The guidance we've received is still very vague and broad."* (P8) Despite initial hesitation, some become more comfortable over time and are willing to *"put transcripts in and ask for summaries of documents"* (P4) but remain cautious about inputting highly sensitive data.

Western bias embedded in LLMs limits their effectiveness in diverse cultural contexts. Negotiators pointed out that these tools often reflect the perspectives of their Western developers, which can be problematic for global field teams. P12 emphasized, *"You really need to go to Nigeria, to Myanmar, and talk to the field teams and get their inputs on all these."* This bias also *"affects locally recruited* *negotiators"* who may struggle with interface designs and language barriers.

Reliability and trust issues arise due to the potential for confusing or inaccurate outputs from LLMs. Without understanding the sources behind the responses, users find it difficult to fully trust the tool. One negotiator questioned, *"What information is the AI basing this on? And why is it choosing those specific documents to answer your question?"* (P12) However, even imperfect results can be acceptable under time constraints or serve as inspiration, challenging ideas and helping to refine thoughts. One negotiator noted, *"Even if AI is not right, we can tell. It's okay because it challenges our ideas."* (P4) They likened AI to an *"aspiring partner"* that stimulates thought.

Overreliance on LLMs raises concerns about diminishing negotiators' skills. Excessive automation might lead to reduced ability to conduct effective negotiations independently. As one negotiator warned, *"I think it's important to do your own analysis. My worry is lack of engagement in the analysis."* (P4) Another added, *"These tools will shrink our brains because we will start relying on them."* (P3)

4.2.2 Public opinion/ mandators' opinions on AI. Negotiators often operate within the limits set by their mandators and consider the well-being of local communities. A mandator, such as the head of a humanitarian organization or a country, sets the strategic objectives and boundaries for negotiators, for example, by providing guidelines on terms like ensuring aid worker safety and maintaining neutrality during conflict zone negotiations. The opinions of both mandators and the public significantly influence whether and how negotiators would think about using LLMs. One negotiator voiced concerns over using LLMs without *"taking the time to really explain and understand what the perception of local communities is [on AI]."* Similarly, mandators such as governments' opinion on LLMs can impact the usage of LLMs in negotiation. *"[...] I'm not sure governments would be very happy or trustful to know that we're using AI in order to negotiate for them."*(P10)

The doubts may come from lack of understanding of these tools. Negotiators think that more public understanding about the limitations and benefits of these tools will facilitate the conversation of whether to integrate LLMs into negotiators' workflow. *"I feel if there were loads of people who understood it a lot better, they'd feel more comfortable using LLMs in negotiation, because we all have these myths about what this stands for. [...] There should be a whole education on the reality of these tools and what they can offer. I don't think people are very well informed."* (P12)

4.2.3 Practical Limitations of AI in Negotiation. The effectiveness of AI tools like ChatGPT in negotiation relies heavily on the quality and structure of the input. Negotiators noted that significant human effort is required to prepare information, such as converting meeting notes, emails, or phone call records into structured formats for AI to process. Moreover, negotiators often lack structured notes, particularly when handling sensitive or confidential information that isn't formally documented. As P11 explained, *"If this [case file put into ChatGPT] originates from existing notes or meetings ... sometimes we don't have those notes. [...] After a phone call, I might jot down two points, but I wouldn't have a full script or source for this [to input into ChatGPT]."* However, they also thought that AI might encourage negotiators to be more organized: *"It may enable a negotiator who is all over the place to be more structured. And so it can be a positive thing."*

5 Understanding LLM in Negotiation with a Probe

5.1 Overview

To address the concerns identified in the interview study—such as difficulties with prompt engineering, limitations of AI summarization, and fears of overreliance—we designed a second study to observe how negotiators interact with AI assistance in practice. While negotiators were already experimenting with tools like ChatGPT, they emphasized a need for support in specific cognitive tasks: context analysis, compromise ideation, and risk assessment. The interview study relied on self-reports; in contrast, this second study introduced a process-oriented AI probe designed to support human reasoning, not automate negotiation tasks. This approach positions AI as a tool to enhance, not replace, negotiators' decision-making processes.

We first asked participants to interact with the unmodified ChatGPT interface as a baseline, reflecting current AI capabilities. We then introduced a custom-designed probe interface that structured interactions around key negotiation tasks identified in the interview study, allowing negotiators to engage with AI as a facilitator of their own reasoning rather than a decision-maker. The probe structured negotiation preparation, eliminating the need for prompt engineering—a key frustration from the interview study.

Through this study, we aimed to observe:

- How do negotiators use LLM tools during realistic negotiation tasks?
- How do negotiators engage with AI when it is designed to support rather than automate decision-making?
- How does reducing the burden of prompt engineering affect usability and interaction quality?

This study was designed to elicit grounded reactions to an alternative AI-assisted work model—one that prioritizes negotiator agency over automation. By presenting negotiators with two different visions of AI collaborations (using ChatGPT and the probe), we sought to understand how such a shift might reshape expectations and practices in negotiation.

5.2 Probe Design Overview

We used OpenAI's GPT-4o API [80] to design tools for the preparatory stage of negotiation—context analysis, ideating compromises, and risk analysis—rather than providing direct instructions.

The prompts for generating any content in this probe interface were created using the ChainForge interface [4], with input from a co-author with over 30 years of frontline negotiation experience. The full flow of the prompts is shown in figure 1. These prompts were designed and evaluated to ensure the most suitable outputs.

5.2.1 Context analysis with negotiation tools Island of Agreements and Iceberg CSS. The probe first collects a case file input and identifies negotiating parties. It then summarizes the case and populates two context analysis templates that some negotiators rely on in their preparations: Island of Agreements and Iceberg CSS [24]. The Island of Agreements template is used to capture (1) agreed facts, (2)

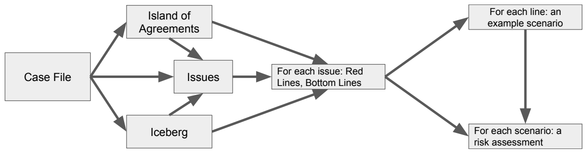

Figure 1: Workflow of the negotiation interface: Starting from a case file, the system generates Island of Agreements and Iceberg context analysis. It identifies issues, determines red lines and bottom lines for each issue, and provides example scenarios with associated risk assessments. The arrows represents the order and what results from the prior generation the system takes to generate the next.

contested facts, (3) shared norms, and (4) conflicting norms between the negotiation parties. The agreed facts and the shared norms comprise the actual Island of Agreements and highlight some of the common ground between the negotiating parties. The Iceberg CSS is constructed by first capturing the stated position of the counterpart — this is the visible part of the "iceberg". The negotiators then work backwards to try to understand the invisible parts: the reasoning of the counterpart and, ultimately, their motives and values. The Iceberg CSS can also be used in the bottom-up direction to clarify the motives, values, reasoning and stated position of the party represented by the negotiators. Putting the two icebergs side by side can also allow the negotiators to identify additional common ground between the negotiating parties.

Then the interface sends the Island of Agreements and Iceberg CSS as prompts to generate negotiation *issues* to be addressed, as many negotiation usually contain multiple issues to be discussed and settled. Once these are generated, users can view them on a page. They can then validate and modify the content in the provided text boxes. An example overview of the context analysis interface is in Figure 2.

5.2.2 Find Zone of Possible Agreements with Bottom Lines and Red Lines. After summarizing the context using IoA and Iceberg CSS, the interface utilizes these outputs along with identified issues to generate the red lines and bottom lines tool. The red lines and bottom lines tool defines negotiation scenarios within the boundaries of the mandate. The interface first identifies key "lines" for both parties: 1. Ideal Outcome (Point A): The best-case scenario that fully achieves the user's organization's objectives without considering the counterpart's needs. 2. Bottom Line (Point B): The minimum acceptable outcome where further concessions start to diminish benefits but remain tolerable. 3. Red Line (Point C): A non-negotiable boundary tied to critical legal, organizational, or reputational risks. The same process is conducted for the counterpart, identifying their ideal outcome, bottom line, and red line, denoted as A', B', and C', respectively (see Figure 3 for a conceptual illustration).

The interface generates a spectrum for each negotiation issue, delineating possible scenarios based on predefined negotiation lines (Figure 4). Each spectrum consists of the following points, from left to right: (1)The user's red line is violated. (2) The user's bottom line is violated. (3) Neither party's bottom line or red line is violated. (4) The counterpart's bottom line is violated. (5) The counterpart's red line is violated. Moving from right to left on the spectrum represents increasing risks for the negotiator's party, with compromises shifting toward unfavorable outcomes.

The Zone of Possible Agreement (ZOPA) is identified within scenarios 2, 3, and 4, where neither party exceeds their red line. Scenario 3, located at the center of the spectrum, represents the most balanced outcome as neither party needs to cross their bottom line. The interface visualizes these zones, helping negotiators strategically identify viable agreements while assessing risks.

5.2.3 Generate Risk Assessment. The user can click on any of the scenario boxes to indicate that they are considering picking the outcome for that specific component of negotiation (Figure 4). The user can choose to let the interface generate a table of risk assessment, indicating the short term risks and long term risks, mitigation strategies, and risks after mitigation. These assessments are generated by categories: Security of Field Teams, Relationship with Counterpart, Leverage of Counterpart, Impact on other Organizations/ Actors, Beneficiaries/ Communities, and Reputation. These standards were adapted from the practices of P1 and their organization from the interview study. The users can modify these columns directly after they are generated.

5.3 Methods

5.3.1 Participants. We recruited 18 negotiators different from the interview study from the same survey data collection. Before we finalized the list of negotiators to conduct the second study, we collected 61 responses. We provided each participant with $30 Amazon gift card to thank them for their time. Again, participants are from 13 humanitarian organizations, from 15 different countries (See Appendix E).

5.3.2 Study Design. This was a qualitative within-subjects study that followed the baseline-intervention design (similar to, e.g., [97]) with the order of conditions fixed such that all participants interacted with the ChatGPT condition first followed by the probe interface. The baseline-intervention design is appropriate when the baseline condition is similar to a practice that participants are already familiar with. Our recruitment survey showed that all participants in the probe study have used ChatGPT or other LLM-based chatbots, 9 out of 18 participants have at least used LLMs a few times a week, and 8 out of the 18 participants have at least used LLMs a few times a week for negotiation specifically. (See Appendix E)

The baseline-intervention design is used when the intervention condition may have lasting spillover effects on how people approach the task, while the baseline condition is unlikely to have such an effect. Our probe interface can have a spill over effect by explicitly focusing on specific parts of the preparation process, which can lead to participants "getting hints" on how they could have prompted ChatGPT if they were to use the probe interface first before using the ChatGPT. For example, after using the probe, the participants might start with generating Island of Agreements while using ChatGPT, which, in our observation, is uncommon for participants using the ChatGPT first.

Each study session lasted approximately 60 minutes or over, and was conducted remotely via Zoom. The study consisted of three parts:

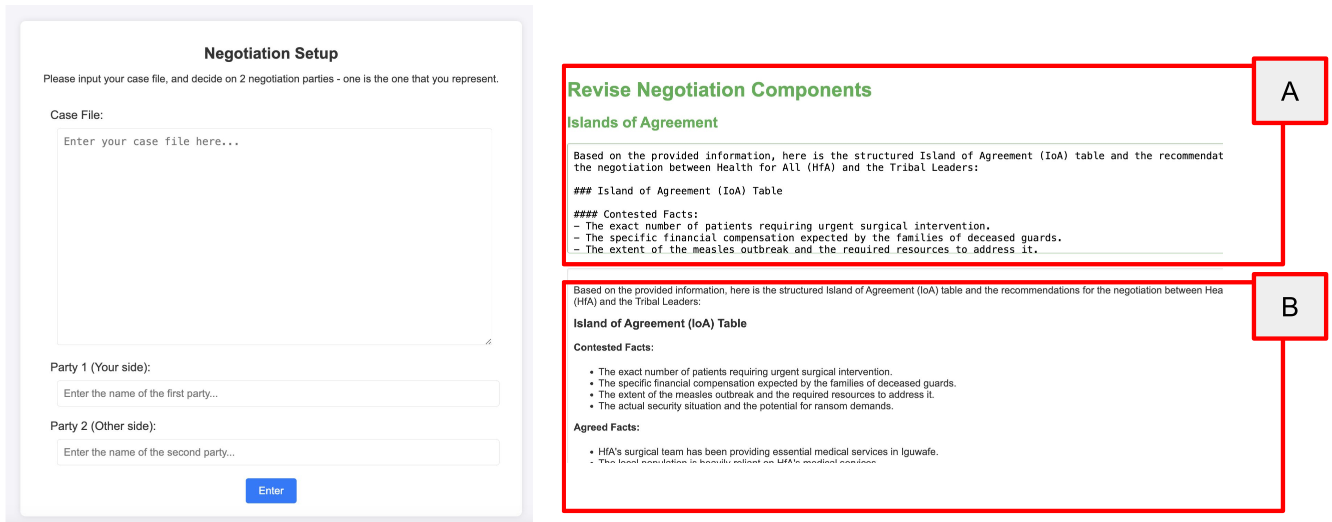

Figure 2: At the left, the user can input the case file, and the negotiation parties. Then the probe will generate Islands of Agreement, Iceberg CSS, and Components. The user can choose to inspect the generated content (for example in box B) in markdown format, and modify in a text box (for example in box A).

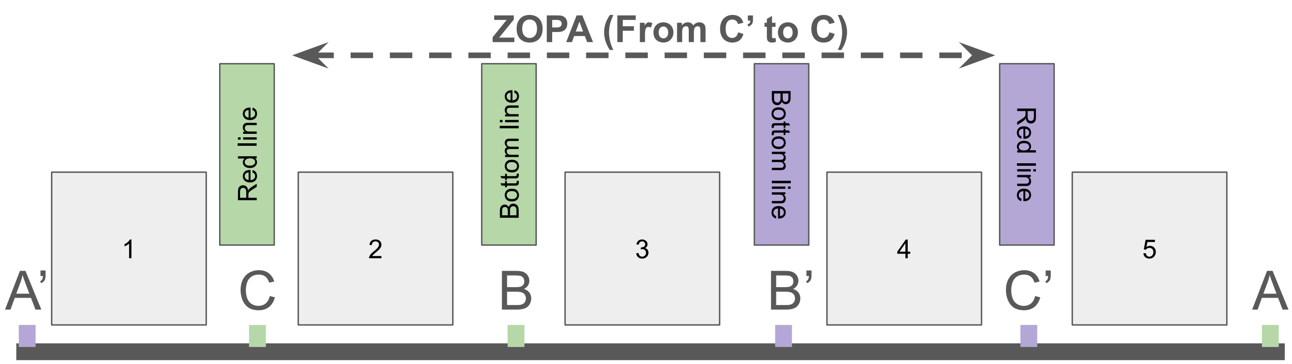

Figure 3: Spectrum of negotiation lines and outcomes: This is a conceptual illustration of red lines and bottom line tool, mapping the identified red lines (A), bottom lines (B), and ideal outcomes (C) for the users' party, and A', B' and C' are corresponding lines for the counterpart.Between each line lies 5 possible scenarios where (1)The user's red line is violated. (2) The user's bottom line is violated. (3) Neither party's bottom line or red line is violated. (4) The counterpart's bottom line is violated. (5) The counterpart's red line is violated. Zone of Possible Agreement (ZOPA): The range of scenarios (2, 3, and 4) where neither party's red line is violated. Scenario 3 is the most balanced and viable agreement. The top section illustrates the implementation of the red lines and bottom lines for *"continuation of HfA operations in Igwafe"* in the interface, detailing red lines, bottom lines, and potential compromises for both parties. Scenarios closer to the center represent balanced outcomes with minimized risks. The bottom section is the conceptual illustration of red lines and bottom lines tool for comparison.

Part 1 (20 minutes): Participants were given one of two anonymized real-world case files. These case files can be found on the online supplementary material. They were asked to use ChatGPT while sharing their screen over Zoom to develop a negotiation plan from a humanitarian perspective. After 20 minutes or when they felt ready, participants drafted a preliminary negotiation script categorized into three tiers:

- Issues and terms that are low-cost and high-benefit for both parties, aimed at building relationships and setting a positive tone.
- Issues and terms that involve some cost and benefit for both parties, intended to establish a fair basis for distribution.
- Complex issues and terms that are high-cost and central to the conflict, potentially contentious and best addressed later to avoid confrontation.

These three tiers reflect the actual preparation work that negotiators typically undertake before a negotiation, following the use of the Red Line/Bottom Line tool [24]. We included this task to prompt participants to engage with a realistic aspect of negotiation preparation during the interview.

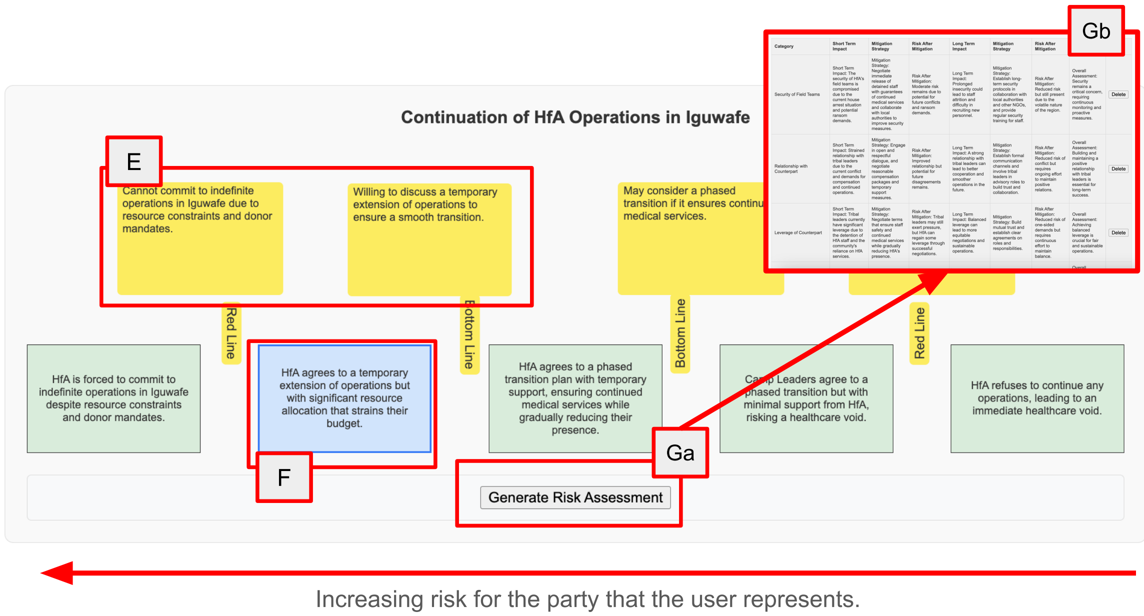

Figure 4: Overview of the red line and bottom line interface. Panel E shows the red line and bottom line of the negotiation party that the user represents. The two boxes on the right show the red line and the bottom line of the counterparty. Panel F shows one of the scenarios or outcomes that the negotiator decides on. The outcomes are listed such that, from the right to the left, there is an increasing amount of risk of negotiation failure as the outcome crosses the bottom line and the red line of the party that the user represents. After selecting the scenario, the user can generate a risk assessment for this selection through button Ga, and show a risk matrix in Gb.

Part 2 (20 minutes): Participants were then provided with the other case file, which they had not yet seen. They were given access to the probe interface and again asked to share their screen over Zoom. The task was to perform the same analysis (3 tiers) using the probe interface. Participants were encouraged to think aloud and share their immediate thoughts during the first two parts of the study. The order of using ChatGPT and the probe interface was not varied, as starting with the probe might have influenced participants to approach the analysis differently (e.g., by breaking the task into smaller steps) than they would when using ChatGPT.

Part 3 (20 minutes): In the final part, participants were interviewed about their experiences using ChatGPT and the probe for case analysis, the added value of each system, any issues encountered, suggested improvements, and other AI features they would like to see. A complete set of interview guidelines can be found in Appendix D.

5.4 Data Analysis

As in the interview study, all sessions were audio recorded and transcribed. To analyze the data, we employed thematic analysis with open coding [16]. The first author initially coded the transcripts and identified a preliminary set of themes. Three members of the research team then convened to review the themes and associated data segments, resolve any discrepancies, and refine the theme definitions. The first author subsequently re-coded the data using the finalized themes, and the research team validated the final set of themes during a follow-up meeting.

6 Results

6.1 Common Practices of Negotiations

This section will discuss the common practices related to frontline negotiation that emerged from the interviews.

6.1.1 Frontline negotiation is complex and context-dependent. Frontline negotiators think that humanitarian negotiation differs from other forms of negotiation, such as business or transactional negotiations, due to its complex context. As P5 explains: *"In the humanitarian world, there is such a myriad of breadth of complexity of environments, social environments that we operate in. Business, well, of course, there are different approaches, but primarily it's very transactional in terms of a negotiation."* For example, P2 explained how nonprofits in Japan are perceived differently from US and Europe: *"Nonprofits are not trusted in Japan. So if ChatGPT said something along lines of 'you should seek public support for your project'. All of us in the room would go. 'That's a valid answer, maybe for United States, or maybe Europe. But for Japan, based on our past-experience, our own understanding of where we stand within society, that would not be a valid strategy.' "* This complexity requires practitioners to continuously adjust and adapt their negotiations styles in accordance with changing situations (P6).

6.1.2 Frontline negotiation requires empathy and flexibility. Being Empathetic and flexible are important qualities for reaching an agreement in negotiation. Yet,*"the human ability to be flexible, to adapt and react [...], and understand the human element and the human dimension of a negotiation"* (P3) are missing in current AI tools. This human touch extends beyond mere strategy to building trust and relationships, sometimes requiring negotiators to patiently invest time and effort in creating a foundation of trust. P5 illustrated: *"I've sat in various operations where I had to drink tea for weeks on end with a counterpart and not get to the nitty-gritty until they were ready to do that because they felt that they could trust you."* In other cases, it simply requires the negotiators to act quickly in response to rapidly changing dynamics: *"Things might get tense. You know. I'm thinking of a negotiation in my case where suddenly gunshots are fired during the negotiation. That changes the dynamic very quickly. Because now I'm really not in favor of the local guards being in charge of security because they're firing guns in the air."*

6.1.3 Verifying information in negotiation is iterative. Negotiation is an interactive process involving continuous information exchange and analysis between negotiators, counterparties, and stakeholders. Its dynamic nature requires ongoing updates as new information emerges and gaps are identified. P1 illustrated, *"Negotiation is not a give, put money, get answered process [like a vending machine]. It's an interactive process, so it means that we collect information, we work on it, and then we analyze it."*

The iterative nature of negotiation affects how information is verified. Negotiators reported verifying information at different levels, from individual checks to group consensus, ensuring accuracy and preventing complications. P3 described their verification process: *"Every time there's been a doubt we verified it. As far as I know, nothing that has been produced as part of my job has led to, you know, mistakes, complications, or diplomatic tensions, because we got something wrong. If we saw something strange, we verified it, and then we kicked it in. If it was correct and we moved it. And we couldn't move without [verification]. "* (P3) Achieving consensus iteratively can be challenging, especially when cultural factors and individual skepticism complicate the process. P6 highlighted this difficulty, particularly in contexts where unanimous agreement is valued: *"I'm in Japan, where consensus building is a huge issue, that you could have 10 people in a room, and if 9 people agree, that's still not consensus. So, in other words, if there's only one person who's providing an analysis, there could be a sense of 'is that really true?'"*

6.1.4 Negotiators believe that their work practices can reduce the impacts of AI biases and mistakes on outcomes: Negotiators believe that biases are unavoidable, and managing them is essential in the negotiation process. As P4 noted, impartiality is often an *"ideal perspective"*: *"Impartiality is idealistic—nothing is truly neutral. We are never independent, never neutral, never impartial. Humans compromise on principles."*

Mistakes, whether human or AI-generated, are viewed as natural parts of negotiation, and opportunities for deeper discussion about disagreements and differing perspectives. P6 explained, *"No, not at all. [...] I go in with my own assumptions and biases. [...] If I'm aware of those, I can adjust when speaking to the other party. [...] Disputes over facts don't bother me as much. [...] What would concern me is if we both believed we agreed on a fact, but in reality, we didn't."*

While AI mistakes can introduce ambiguities that complicate human interactions, negotiators, as P1 explained: *"When I use ChatGPT to analyze something, I always do another layer of triangulation, human intelligence, talk with some other people to verify if this thing, which I thought or which we thought, is okay to go forward."* P6 provided a similar approach: *"It gave me a direction, but I would need to verify on the ground ... I'd probably find the people that's closest to me and run it by them to find out what their perspective is."* P1 stressed the importance of using information from trusted sources, stating, *"If I use ChatGPT, I would instruct it to reference only ICRC-approved materials. I'd only trust the output if it aligns with our official strategy and tools."*

Experienced negotiators might have more reliable expertise to mitigate risks, while novices may over-rely on AI outputs without questioning them. As P2 explained, *"If I didn't have a strong background in due diligence, I might be overly confident in my assessment. I'm more concerned about novice negotiators using ChatGPT and assuming it's correct. Experienced humanitarian negotiators have enough buffers around them to avoid major mistakes, but novices could misstep."*

Finally, negotiators emphasized that successful outcomes depend more on relational dynamics and trust than on factual accuracy alone, such as building rapport with local stakeholders, understanding cultural nuances, and fostering mutual respect to create a foundation for productive negotiations. P3 pointed out, *"It's important to meet with local stakeholders to build rapport and understanding. [...] Having that kind of information also tells me the impact of going with a certain line could have consequences on the relational aspect of the negotiation."*

6.1.5 Negotiators Feel the Need to Take Full Responsibility When Using AI tools. Negotiators stressed the importance of maintaining full responsibility when using AI tools like ChatGPT, unlike situations with human counterparts where institutional power may influence decisions. P2 explained, *"ChatGPT is just a tool. If I choose to trust it and things go wrong, the responsibility is still mine."*With human input, responsibility may shift due to institutional power dynamics or professional networks. In comparison, taking positions from a human can differ because a person may hold institutional power (such as *"a manager who has institutional demand power"* (P1)) or be part of the negotiator's professional network (for example, *"a colleague who is also in the negotiation"* (P1)), so that negotiators are able to transfer some of the responsibilities to them.

6.2 Limitations of LLM-powered applications (for both ChatGPT and Probe Interface)

6.2.1 Efficiency isn't all that negotiators care about. While negotiators *felt* that both ChatGPT and the interface increase productivity, negotiators emphasized that they don't address the core aspects of negotiation. P4 described using ChatGPT as *"cheating"* for tasks negotiators should already handle, while P1 noted that efficiency cannot replace human expertise: *"The Interface saves time, but if it didn't exist, I'd still need to manage these scenarios on my own. I can also ask a manager or a more experienced colleague."* P7 agreed, stating that such tools can only supplement, not replace, the knowledge and experience negotiators rely on: *"It can only be in addition to already existing systems, knowledge, experience."*

Negotiators also stressed that these tools are most effective during preparation, with their benefits diminishing during actual negotiations. P1 explained, *"If someone has this strategy [generated by the Interface], but as cold as it could be, lacks negotiation skills, this is just a document and this is just information. But how to use information is important with the complete package of the negotiation."*

6.2.2 Automating the Human Elements Away. Negotiators expressed concerns that current AI tools fail to capture key aspects of human interaction. As P5 explained, *"It's often the soft stuff [...] like body language and tone of voice"* that plays a significant role during negotiations. This limitation means essential non-verbal cues, important for understanding and influencing counterparties, are lost. P4 echoed this, warning that over-reliance on AI could *"remove some of the human components from the whole negotiation."*

Beyond non-verbal cues, negotiators worried that widespread AI use might result in *"less genuine interactions,"* as P5 noted, reducing creative problem-solving and fostering a *"sterilization of the negotiation process."* P16 added that negotiators might struggle to find innovative solutions to resolve deadlocks if AI oversteps its supportive role.

However, these limitations highlight the distinct roles AI tools and human negotiators play. While AI can assist by providing context and aiding strategy development, it cannot replace the human elements that build trust and empathy. As P4 explained, *"Humanitarian negotiations require proximity with the person, empathy, those kinds of things. And this kind of tool might not be able to capture that."* P3 echoed this sentiment, *"[Human emotions and alike] wouldn't be something that you would be able to necessarily input in there... But it's not designed to do that... It is not supposed to capture the diplomatic aspect of the negotiation or the human aspect of the negotiation"*

6.2.3 "Overconfidence" of LLM tools. ChatGPT always appears to be confident, which makes it hard for negotiators to identify uncertainties: *"Well, it will answer with great amount of certainty. And so that's a caution for making sure that you also feel confident about what it's the information that it's giving back to you so that you, you would need to know enough about the situation to go 'ehhh that doesn't sound quite right,' "* (P12)

What makes LLMs harder to trust is that they do not currently proactively identify or inquire about missing information. Unlike human negotiators who ask clarifying questions, LLMs provide answers based solely on the given input without flagging gaps. P1 pointed out that if something is not mentioned in the document, *"the generator will not mention it as a suggestion."* P1 also shared how they got the LLM to query missed information by breaking tasks into smaller parts and providing explicit instructions, but noted that this requires extra effort.

P10 proposed a potential solution to address this challenge by shifting the interaction with LLMs away from purely text-based outputs, which often convey an unwarranted tone of confidence. They suggested using alternative modalities, such as visual representations, to make uncertainties or logical errors more apparent: *"The problem with artificial intelligence [...] is that it's so good at generating very well-sounding content that it's difficult to identify mistakes. [...] If we work more with something like conceptual maps, maybe it would be easier."*

6.3 Findings Emerged Through Probe Interface: Design Negotiation Tools by Supporting Negotiation Practice

6.3.1 Show the Process rather than Recommendations. Consistent with the interview study findings, many participants struggled to intuitively use ChatGPT or leverage prompt engineering effectively. More than half (n=10) initially attempted to obtain fixed recommendations by inputting entire case files in a single query. However, when using the probe interface, negotiators recognized the value of breaking tasks into steps and preferred having options and flexibility over receiving fixed recommendations—a tendency that emerges when ChatGPT is prompted in a one-shot manner, often yielding static responses rather than iterative guidance. P8 explained, *"I'd rather the AI poses those questions, what compromises are you comfortable with? And here are the different options, and then it's up to you to choose which one... when it says prioritize and avoid, that's up to me to decide. That's not up to the AI to decide. In certain situations, I might take more risk because there's more lives at stake or the situation is by far worse, but sometimes I might take less risk if the situation is not as bad or if there's a possibility that there might be another solution that might come up."*

On the countrary, **the probe interface's focus on process** rather than fixed strategies allowed negotiators to explore options and develop context-specific strategies. P3 emphasized, "[The interface]'s set up to help you organize your ideas, give you a comprehensive view of everything, and not to give you every solution or guarantee you a solution. I think it's excellent." Others appreciated that this approach led to less biased decisions. "What your program did was it actually kind of removed my initial sensory inferences that may have clouded my assessment." (P6) This observation echoed with other negotiators who prefered that AI tools only give the users options rather than recommendations: "I rather these tools give us just different options with different risks associated with it, and then just say, okay, user, it's up to you what you want to do, what choice you want to make." (P8). Additionally, to the negotiators, our interface that showed the process aided in the ideation process, helping teams brainstorm and outline potential solutions. As P2 explained, "The benefit of this model is that a lot of those things are in my head, more at the subconscious level. But to act on them and make them useful, I've got to bring them to the conscious level. When I see it here, I can say, 'Oh, yeah, I see how that aligns with what I'm thinking.' "

6.3.2 Process-Oriented AI May Bridge Experience Gaps. Our probe interface seems to help level the playing field between negotiators of varying experience levels. P5 pointed out that if only the experts can generate the necessary frameworks like Island of Agreements then *"the success of a negotiation would only depend on the expert negotiators' [...] With this tool, teams can work better, faster, in shorter time, and they don't need the huge experience in negotiation."* and the probe interface can help less experienced negotiators generate visualizations and analyses that typically require expert knowledge and more years of experience.

However, expert negotiators argue that they don't require many guardrails but instead need flexibility from AI tools to tailor outputs to their specific needs and adjust strategies dynamically. When asked how the probe interface could be improved, P11 suggested to allow interacting with a chatbot in the interface directly to modify the content: *"For example, if I can make some changes here directly, even without altering the original documents, that would be helpful."*

6.3.3 Support Negotiator's Collaborative Practice.

In contrast to ChatGPT, participants said that the probe interface facilitates teamwork by making it easier to present, discuss, and revise cases collectively. By generating structured outputs, the interface could potentially simplify projecting information to the team and engaging in productive discussions. As P1 noted, *"It's easier, not just for the individual, but for teamwork. When I prepare a case and put it on the Interface, it's easier for whoever uses it to project it to the team and then have a discussion."* Echoed P5: *"it may help us not just with the negotiation preparation itself, but it's going to help us considerably be able to communicate in a very digestible manner from the different levels of people involved in a complex negotiation."*

This interface can leverage negotiators' existing information validation practices to cross-check LLM-generated outputs. P2 described how they would use the interface to validate information with their peers: *"The added value is, if I have a negotiation team, and collectively, we write up the case study, upload it to the interface, and it gives us island of agreements, we can see where we agree or disagree on facts and norms. It's helpful, especially if not everyone has authored the case study. Running it through the interface provides a simplified way of looking at the data."* P2 further explained how the interface could help categorize information by relevance and accuracy: *"I'd categorize information into three groups: one, confirming what I know; two, definitely incorrect or incomplete; and three, things I need to verify. Even the 'let me check that out' category adds value, as it becomes something I need to confirm during the negotiation and assess its importance."*

The participants thought that the interface not only facilitated collaboration but also helped negotiators clarify boundaries and align with stakeholders. P11 discusses how the interface could potentially help in planning and strategy development, allowing negotiators to clarify what can and cannot be done, which is important for aligning with the mandate and gaining consensus among stakeholders. *"I would use that in the planning, specifically in planning to elaborate the strategy and the planning of the different steps, and also for sharing with the mandate subject. [...] So you will clarify the mandate subject, what I can do or not do, and to clarify the mandate more clearly also."* (P11) This negotiator mentioned that they would use the Interface for strategic planning, particularly when clarifying key boundaries like the red line and bottom line, while using ChatGPT for strategies that are more granular: (P11) *"For example, this work on the red line and bottom line is something that concerns them directly. While, for example, 'which tactics I use for the negotiation' is more delegated to me. I would use the normal ChatGPT we saw in these days, more tactically."* (P11)

7 Discussion

7.1 RQ1: What needs do negotiators have that can be assisted by LLMs?

In the formative study (interview study), we discovered two potential processes that the negotiators try to or envision using LLMs for. First, negotiators would use LLMs for negotiation preparation, including context analysis, ideating compromises, and risk analysis. A number of negotiators have already explored how to use ChatGPT for these tasks on their own. Second, negotiators hoped that LLMs could assist in preserving institutional memory and in the training of new negotiators without exposing them to hostile environments by anonymizing old cases so they could be shared between organizations.

7.2 RQ2: What concerns do negotiators have about using LLMs in negotiation?

7.2.1 General Concerns about LLMs.

Frontline negotiators brought up some concerns that had been previously highlighted in other application domains. These concerns included confidentiality, hallucination and bias, which were frequently discussed in concerns of AI in knowledge work [13, 93, 102]. We briefly discuss how these conventional concerns about LLMs can specifically affect the work of frontline negotiation.

Confidentiality: Confidentiality breaches in frontline negotiation can undermine effectiveness and endanger team safety, as sensitive details about strategies or stakeholders carry high stakes. It is not clear to many frontline negotiators whether the information they share with an LLM can be accessed by others, leading to a false sense of security when using tools like ChatGPT. Over time, growing comfort with such tools can result in complacency, increasing the risk of disclosing sensitive information.

Hallucination and Bias: Although negotiators have concerns over hallucination and bias, they employ strategies like fact-checking, leveraging field experience to mitigate risks of LLMs' hallucinations and biases. Due to their years of experience, experienced negotiators are less concerned about these risks, viewing them as manageable with their expertise. Additionally, under time pressure, having some results from LLMs is better than having nothing at all. However, participants hypothesized that novice negotiators may struggle to identify and counteract biases, increasing the risk of over-reliance on AI outputs. These concerns highlight the importance of focused training and support, especially for less experienced practitioners.

7.2.2 Concerns Specific to the Work of Frontline Negotiation.

Frontline negotiators have also brought up several concerns that were specific to their profession and that are not commonly discussed elsewhere.

Mandators' Influence: Negotiators take two key factors into account when it comes to mandators: how mandators perceive the use of AI by negotiators and whether AI can enhance the presentation of information to mandators. When adopting AI tools, negotiators must consider the perspectives and directives of their stakeholders, such as the countries or organizations they represent. These mandates significantly influence the feasibility of integrating AI into the negotiation processes. Moreover, negotiators appreciated that our probe facilitated clear communication and justification of their strategies to mandators, emphasizing the value of an AI tool that supports this function—an aspect not explicitly identified as a cognitive support need. Such tools should therefore enable transparent documentation and presentation of processes and outcomes to align with stakeholder expectations. For instance, the probe interface provides outputs like red lines and bottom lines—key strategic components that are often prioritized by mandators.

LLMs do not raise questions about missing information: Negotiators noted that LLMs (both ChatGPT and our probe interface) currently cannot actively identify or inquire about missing information. This limitation is frustrating, while recognizing and addressing gaps is crucial in negotiations. Unlike LLMs, human negotiators instinctively ask clarifying questions when sensing missing details. Designs of negotiation tools need to offer ways for the users to probe for missing information.

Integrating LLMs into negotiators' work flow may generate more cost than benefits: Frontline negotiators weigh the benefits of LLMs against their challenges, sometimes finding them impractical for integration into negotiation practices. As participants noted, tasks like creating Island of Agreements with ChatGPT require gathering data that are not currently available in a digital format, such as stakeholder meetings or local archives, which is labor-intensive and time-consuming. The lack of structured note-taking practices among negotiators would also demand significant workflow changes, increasing human effort. Furthermore, the resource demands of inputting, verifying, and editing LLM-generated outputs can outweigh their benefits, making their use less advantageous and, at times, counterproductive. Our findings align with prior research on AI adoption in high-stakes decision-making. Addressing the motivations, values, professional identities, and established norms shaping people's work is essential for successful technology deployment in such settings [12, 44].

7.3 RQ3: Do current LLM tools assist negotiators' key decision support needs?

Our research has identified that, in contrast to our probe interface that provides process-oriented support, interface like ChatGPT encourages the users to obtain outcome oriented support. Thus, current LLMs accessed via chat interfaces may fail to inherently support the essential cognitive processes that negotiators employ during negotiation preparation.

During our study, negotiators instinctively inputted entire case files into these systems when interacting with ChatGPT, anticipating comprehensive results. Both prior studies [104] and our observations indicate that users struggle to decompose their tasks into smaller, manageable steps. The chat-based interaction mode — posing a single question and expecting a flawless answer — parallels a goal-centric AI approach focused on delivering recommendations and solutions rather than the **process-centric approach** focused on supporting users and empowering them to make their own informed decisions.

When comparing the goal centric approach to the established practices of negotiators that we discovered, we realized that the chat-based interaction style may be inadequate for supporting humanitarian negotiation processes. In frontline negotiations, the preparatory phase involves organizing information, identifying priorities, and aligning strategies with organizational mandates. These tasks are inherently iterative and context-dependent, requiring a balance of facts, norms, and potential risks. Unlike goal-centric AI, which offers direct recommendations or steps, negotiators may require process-oriented tools that support their cognitive processes, facilitate collaboration, provide the flexibility to adapt strategies dynamically, and ultimately grant negotiators the agency to make their own decisions. Consequently, our results also indicate that existing AI negotiation tools like Nibble and Practum [76, 100] prioritize transactional outcomes and may not fully meet the needs of humanitarian negotiations. These systems often provide a single recommendation or handle negotiations for users, restricting the consideration of alternatives and adaptability to changing circumstances.

In contrast to chat-based interactions, our probe design demonstrated how LLM-powered tools and technological tools in general could support negotiators in their key cognitive processes while leaving them in charge of the key parts of the process. This approach offers two main benefits: it eliminates challenges associated with prompt engineering, and it allows negotiators to leverage the probe design to mitigate inherent limitations of LLMs, such as hallucinations and biases, by facilitating processes like information validation and cross-checking results with peers.

Moreover, the probe interface has the potential to support collaborative workflows, allowing multiple stakeholders to coordinate, align strategies, and synthesize diverse inputs—essential aspects of humanitarian negotiation. By providing structured outputs and interaction models tailored to negotiation processes, the probe can enable teams to maintain a coherent strategy across different members. By contrast, chat-based LLM interfaces like ChatGPT tend to be designed for individual use, lacking mechanisms to support these collaborative efforts. They do not facilitate managing multiple viewpoints, integrating team feedback, or ensuring alignment in group decision-making. Without structured collaboration features, negotiators may struggle to use LLMs in team-based negotiation settings.

7.4 RQ4: What are the LLMs' anticipated impacts on the work of frontline negotiators?

7.4.1 LLMs may automate away the human elements of negotiation. LLMs are currently incapable of replicating the human elements of negotiation, and even if they could in the future, negotiators strongly resist the idea of replacing these critical human aspects. Our findings highlight the unique contributions of frontline negotiators, including building long-term rapport with counterparts, upholding humanitarian principles, exercising duty of care, balancing trade-offs between risks and human lives, and generating creative solutions in deadlock situations. These aspects not only drive successful negotiations but are also often viewed by negotiators as the defining qualities of their profession—qualities they believe AI should never attempt to replace.

Participants hypothesized that widespread reliance on naive applications of LLMs could lead to "sterile" negotiations, lacking the human touch and diminishing creativity. This concern is supported by prior research, which identified the phenomenon of "design fixation" [55, 90], where reliance on mundane examples or suggestions reduced novelty and diversity of solutions. Recent studies on LLMs and creativity echo this concern, showing that while LLMs generate high-quality ideas, they lack novelty compared to human-generated ones [15]. Quantitatively, the variance in LLM-generated ideas is lower than in human-generated ideas [15], potentially leading to more uniform solutions across negotiations. Furthermore, prior work on word suggestions and assistive writing tools demonstrates that these technologies influence users to adopt more uniform tones and sentiments aligned with the models they use [5, 6, 85].

In the future, LLMs may become more creative and adept at relational tasks, such as building connections with counterparts [66, 67], interpreting body language [11], and managing other diplomatic and human-centered aspects of negotiation, particularly as more advanced multi-modal models are developed. However, our study indicates that negotiators are unlikely to welcome this shift. Echoing findings from previous work [102], negotiators value a clear division of labor between AI and humans. They are open to AI handling analytical tasks, such as context analysis, while reserving relational and ethical dimensions for human expertise. This division is not just considered practical but essential for preserving the integrity and creativity of negotiation practices.

7.4.2 LLMs may both support and strain the gap between novice and expert workers. Our results also suggest that the efficiency with which negotiators utilize LLMs—whether through a vanilla chat interface or our custom probe interface— might be influenced by their years of experience. This finding adds nuance to, and in some cases challenges, the prevailing narrative that technology will transform knowledge work by bridging the gap between inexperienced workers and experts [19, 31]. Contrary to the notion that expertise will become obsolete or that labor will be replaced [3, 3, 26, 65, 70, 87, 98], our findings suggest that without technologies that specifically amplify and support expert-designed and tested practices, novice negotiators may not achieve expert-level proficiency, and expert negotiators will continue to thrive in humanitarian negotiation.

In contrast to ChatGPT's lack of guidance in breaking down tasks, participants noted that the probe interface could support novice negotiators by generating frameworks typically crafted by experts. By addressing areas where negotiators struggle the most, the interface has the potential to help less experienced individuals close the skill gap. However, they emphasized that without fundamental negotiation skills, the tools, albeit ChatGPT or probe interface, merely provides information. Additionally, expert negotiators might have different needs. For example, when using probe interface, some negotiators also indicated that, although they liked the structures provided, they still need flexibility to freely prompt LLMs. This finding echos He et al [51], which suggests that users' automation preferences may shift based on task complexity and familiarity. Expert negotiators may require more adaptable interfaces compared to the current version of the probe interface, which designated the interaction flow.

Thus, to effectively support workers with varying experience levels in knowledge work, it is important to consult with expert users and design systems that teach and assist novices with the necessary skills and expertise. Relying solely on generally helpful technologies like ChatGPT may be insufficient for empowering negotiators with different skill levels, especially novices, to achieve their best performance.

7.5 AI in Negotiation: Implications for Organizations that Seek to Adapt AI in Negotiation

Our results highlight that humanitarian organizations are still in the early stages of understanding and managing the impacts of LLMs on negotiation. Per our interviews, current humanitarian organizations' LLM guidelines are vague (Section 4.2.2), focusing only on usage restrictions while neglecting broader risks. Similarly, prior work cautions a critical gap [9]: existing guidelines at workplace for AI usage rarely address how to integrate AI tools into workflows effectively or distinguish tasks for automation versus those requiring human oversight. This lack of clarity may leave negotiators at risk of unknowingly sharing sensitive data with third-party LLM providers, or overrelying on automation for parts of the task where close human attention and accountability is needed.

We have three sets of recommendations for what negotiators, as a profession, may consider going forward. First, develop professional guidelines for the safe and appropriate use of existing LLM-based technologies. Our results suggest that negotiators expressed concerns LLM's limitations (RQ2). Thus, these guidelines should focus on reducing cognitive burdens for users rather than shifting responsibility onto them simply because new technologies are available. For example, guidelines could introduce heuristics to navigate limitations like hallucinations and biases [9], focusing on practical methods that help users detect errors and apply AI tools effectively. The negotiators may also want to collaborate with technologists to understand hidden limitations of existing technologies (e.g., threats to confidentiality) and to use that knowledge to set (and explain) appropriate bounds for the use of these technologies.

Second, the negotiators should invest in workshops that foster ongoing learning and collaborative discussions about the future role of LLMs in humanitarian work. As our results have demonstrated, the chat interface was rarely the right tool the negotiators — how else (if at all) could the underlying technology be helpful? Here, the negotiators should lead activities to identify the values, success criteria and unmet needs of their profession to ensure future tools genuinely augment human expertise while maintaining ethical and effective practices. Frontline Associates, a professional organization for frontline humanitarian negotiators, have already implemented workshops [37] to help negotiators reflect on the limitations of current tools and identify areas where AI support could be most impactful.

Third, engage with technologists and HCI professionals to design the most effective ways to appropriately support negotiators. Once the negotiators have established the values, the success criteria and the potential needs of their profession, a collaboration with technologists and HCI professionals may lead to the development of truly useful supporting technologies. While mentioning the need to include HCI professionals in these activities will seem obvious to the readers of this publication, we observe that many "AI for social good" partnerships include only domain and technology experts but no design experts (e.g., [82, 89]). It is often assumed that the deep understanding of the problem, which the domain experts bring, is sufficient to identify the most effective interventions. We argue that HCI professionals bring an additional and necessary set of skills to help decide which part of the process to intervene in and how to design the intervention in a manner that incorporates concerns of all relevant stakeholders and that anticipates likely indirect consequences of the intervention.

7.6 Limitations and Future Work

Our research conducted qualitative interviews with novel design probes with anonymized cases to understand the negotiators' concerns, hopes and challenges of using LLMs in humanitarian negotiation. However, our work is a simulation of what the negotiators would do given a possible form of LLM-powered tool. In our future work, we will build realistic experiments, involving negotiators using LLM powered tools to conduct simulated negotiations, strengthening the study's conclusion.

Acknowledgments

This work was supported in part by the ONR under Agreement No. N00014-24-1-2726 and by the Cooperative AI Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of ONR or the Cooperative AI Foundation.

References

[1] Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2024. Cooperation, competition, and maliciousness: Llm-stakeholders interactive negotiation. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
[2] Google AI. 2023. PaLM 2: A New Generation of Language Models. *Google AI Blog* (2023). https://ai.googleblog.com/2023/05/announcing-palm-2.html
[3] Stephen J Andriole. 2024. The big miss: AI will replace just about everything. *Communications of the Association for Information Systems* 55, 1 (2024), 29.
[4] Ian Arawjo, Chelse Swoopes, Priyan Vaithilingam, Martin Wattenberg, and Elena L. Glassman. 2024. ChainForge: A Visual Toolkit for Prompt Engineering and LLM Hypothesis Testing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 304, 18 pages. https://doi.org/10.1145/3613904.3642016
[5] Kenneth C Arnold, Krysta Chauncey, and Krzysztof Z Gajos. 2018. Sentiment Bias in Predictive Text Recommendations Results in Biased Writing.. In *Graphics interface*. 42–49.
[6] Kenneth C Arnold, Krysta Chauncey, and Krzysztof Z Gajos. 2020. Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 128–138.
[7] United Nations. General Assembly. 1949. *Universal declaration of human rights*. Vol. 3381. Department of State, United States of America.
[8] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel S. Weld. 2021. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of CHI '21*.
[9] Kristian González Barman, Nathan Wood, and Pawel Pawlowski. 2024. Beyond transparency and explainability: on the need for adequate and contextualized user guidelines for LLM use. *Ethics and Information Technology* 26, 3 (2024), 47.
[10] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, and Seth D. Pollak. 2019. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest* 20, 1 (2019), 1–68. https://doi.org/10.1177/1529100619832930 _eprint: https://doi.org/10.1177/1529100619832930.
[11] Dhaya Sindhu Battina and Lakshmisri Surya. 2021. Innovative study of an AI voice-based smart Device to assist deaf people in understanding and responding to their body language. *SSRN Electronic Journal* 9 (2021), 816–822.
[12] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M. Vardoulakis. 2020. A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–12. https://doi.org/10.1145/3313831.3376718
[13] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 610–623. https://doi.org/10.1145/3442188.3445922
[14] Emmanuel G. Blanchard and Phaedra Mohammed. 2024. On Cultural Intelligence in LLM-Based Chatbots: Implications for Artificial Intelligence in Education. In *Artificial Intelligence in Education*, Andrew M. Olney, Irene-Angelica Chounta, Zitao Liu, Olga C. Santos, and Ig Ibert Bittencourt (Eds.). Springer Nature Switzerland, Cham, 439–453.
[15] Léonard Boussioux, Jacqueline N Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R Lakhani. 2024. The crowdless future? Generative AI and creative problem-solving. *Organization Science* 35, 5 (2024), 1589–1607.
[16] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (Aug. 2019), 589–597. https://doi.org/10.1080/2159676X.2019.1628806
[17] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
[18] Claude Bruderlein. 2023. Challenges and Dilemmas in Frontline Negotiations. *Frontline Negotiations* (2023). https://blogs.icrc.org/law-and-policy/2018/01/04/challenges-dilemmas-in-frontline-\negotiations-interview-with-claude-\bruderlein
[19] Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. 2023. *Generative AI at work*. Technical Report. National Bureau of Economic Research.
[20] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 188 (April 2021), 21 pages. https://doi.org/10.1145/3449287
[21] Zana Buçinca, Siddharth Swaroop, Amanda E Paluch, Finale Doshi-Velez, and Krzysztof Z Gajos. 2024. Contrastive explanations that anticipate human misconceptions can improve human decision-making skills. *arXiv preprint arXiv:2410.04253* (2024).
[22] Philip Burnard. 1991. A method of analysing interview transcripts in qualitative research. *Nurse Education Today* 11, 6 (1991), 461–466. https://doi.org/10.1016/0260-6917(91)90009-Y
[23] Chen Cao. 2023. Scaffolding CS1 Courses with a Large Language Model-Powered Intelligent Tutoring System. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*. 229–232.
[24] Centre of Competence on Humanitarian Negotiation and Claude Bruderlein. 2019. *CCHN Field Manual on Frontline Humanitarian Negotiation* (2.0 ed.). 106 Route de Ferney, 1202 Geneva, Switzerland. Available online at http://www.frontline-negotiations.org.
[25] ASHLEY JONATHAN CLEMENTS. 2021. *HUMANITARIAN NEGOTIATIONS WITH ARMED GROUPS: the frontlines of diplomacy*. ROUTLEDGE, S.l. OCLC: 1282597929.
[26] Gareth Corfield and Matthew Field. 2022. Meet ChatGPT, the scarily intelligent robot who can do your job better than you. *The Telegraph* (December 2022).
[27] Tim Cummins and Keld Jensen. 2024. Friend or foe? Artificial intelligence (AI) and negotiation. *Journal of Strategic Contracting and Negotiation* (June 2024), 20555636241256852. https://doi.org/10.1177/20555636241256852
[28] Jared R Curhan, Jennifer R Overbeck, Yeri Cho, Teng Zhang, and Yu Yang. 2022. Silence is golden: Extended silence, deliberative mindset, and value creation in negotiation. *Journal of applied psychology* 107, 1 (2022), 78.
[29] Hai Dang, Sven Goller, Florian Lehmann, and Daniel Buschek. 2023. Choice over control: How users write with large language models using diegetic and non-diegetic prompting. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
[30] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2023. Don't Just Tell Me, Ask Me: AI Systems That Intelligently Frame Explanations as Questions Improve Human Logical Discernment Accuracy over Causal AI Explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 352, 13 pages. https://doi.org/10.1145/3544548.3580672
[31] Fabrizio Dell'Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. 2023. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *SSRN Electronic Journal* (2023). https://doi.org/10.2139/ssrn.4573321
[32] Hillary Anger Elfenbein, Jared R Curhan, Noah Eisenkraft, and Ashley D Brown. 2017. Who makes an effective negotiator? A personality-theoretic approach to a longstanding question. *A Personality-Theoretic Approach to a Longstanding Question (October 9, 2017)* (2017).
[33] Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature* 630, 8017 (2024), 625–630.
[34] Molly Q Feldman and Carolyn Jane Anderson. 2024. Non-Expert Programmers in the Generative AI Future. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne,

United Kingdom) *(CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, Article 15, 19 pages. https://doi.org/10.1145/3663384.3663393

[35] Center for Strategic and International Studies. 2024. 2024 Global Forecast: Conflict Zones. (2024). https://features.csis.org/global-forecast-conflict-zones/ Accessed: 2024-05-22.

[36] Agence France-Presse. 2024. Global conflicts herald 'dangerous decade': Military think tank. *Courthouse News Service* (2024). https://www.courthousenews.com/global-conflicts-herald-dangerous-decade

[37] Frontline Associates. 2024. Calendar of Activities: FA Summer Program. https://www.frontline-associates.org/activities/educational-program/summer-program/calendar Accessed: 2024-07-19.

[38] Krzysztof Z. Gajos and Lena Mamykina. 2022. Do People Engage Cognitively with AI? Impact of AI Assistance on Incidental Learning. In *Proceedings of the 27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) *(IUI '22)*. Association for Computing Machinery, New York, NY, USA, 794–806. https://doi.org/10.1145/3490099.3511138

[39] Susanne Gaube, Harini Suresh, Martina Raue, Alexander Merritt, Seth J Berkowitz, Eva Lermer, Joseph F Coughlin, John V Guttag, Errol Colak, and Marzyeh Ghassemi. 2021. Do as AI say: susceptibility in deployment of clinical decision-aids. *NPJ digital medicine* 4, 1 (2021), 31.

[40] Timnit Gebru and Émile P Torres. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday* (2024).

[41] Katy Ilonka Gero, Tao Long, and Lydia B Chilton. 2023. Social dynamics of AI support in creative writing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.

[42] Global Interagency Security Forum (GISF). 2020. Humanitarian Access Score Report 2020: Global Synthesis. https://gisf.ngo/resource/humanitarian-access-score-report-2020-global-synthesis/

[43] Ben Green. 2021. The Flaws of Policies Requiring Human Oversight of Government Algorithms. *Available at SSRN* (2021).

[44] Trisha Greenhalgh, Joe Wherton, Chrysanthi Papoutsi, Jenni Lynch, Gemma Hughes, Christine A'Court, Sue Hinder, Rob Procter, and Sara Shaw. 2018. Analysing the role of complexity in explaining the fortunes of technology programmes: empirical application of the NASSS framework. *BMC Medicine* 16, 1 (Dec. 2018), 66. https://doi.org/10.1186/s12916-018-1050-6

[45] International Crisis Group. 2024. Watch List 2024. (2024). https://www.cfr.org/report/conflicts-watch-2024 Accessed: 2024-05-22.

[46] Zhenyu Guan, Xiangyu Kong, Fangwei Zhong, and Yizhou Wang. 2024. Richelieu: Self-Evolving LLM-Based Agents for AI Diplomacy. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. https://openreview.net/forum?id=7Jb4NJS8Yk

[47] Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, and Yoshua Bengio. 2017. On integrating a language model into neural machine translation. *Computer Speech & Language* 45 (2017), 137–148.

[48] David Gunning and David Aha. 2019. DARPA's explainable artificial intelligence (XAI) program. *AI magazine* 40, 2 (2019), 44–58.

[49] Yue Guo, Wei Qiu, Gondy Leroy, Sheng Wang, and Trevor Cohen. 2024. Retrieval augmentation of large language models for lay language generation. *Journal of Biomedical Informatics* 149 (2024), 104580.

[50] Jessica He, Stephanie Houde, Gabriel E. Gonzalez, Darío Andrés Silva Moran, Steven I. Ross, Michael Muller, and Justin D. Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne, United Kingdom) *(CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, Article 9, 14 pages. https://doi.org/10.1145/3663384.3663398

[51] Jessica He, David Piorkowski, Michael Muller, Kristina Brimijoin, Stephanie Houde, and Justin Weisz. 2023. Rebalancing Worker Initiative and AI Initiative in Future Work: Four Task Dimensions. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Oldenburg, Germany) *(CHIWORK '23)*. Association for Computing Machinery, New York, NY, USA, Article 3, 16 pages. https://doi.org/10.1145/3596671.3598572

[52] Maia Jacobs, Melanie F. Pradier, Thomas H. McCoy Jr, Roy H. Perlis, Finale Doshi-Velez, and Krzysztof Z. Gajos. 2021. How machine-learning recommendations influence clinician treatment selections: the example of the antidepressant selection. *Translational Psychiatry* 11 (2021). https://doi.org/10.1038/s41398-021-01224-x

[53] Kokil Jaidka, Hansin Ahuja, and Lynnette Hui Xian Ng. 2024. It Takes Two to Negotiate: Modeling Social Exchange in Online Multiplayer Games. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 85 (apr 2024), 22 pages. https://doi.org/10.1145/3637362

[54] Rishab Jain and Aditya Jain. 2024. Generative AI in Writing Research Papers: A New Type of Algorithmic Bias and Uncertainty in Scholarly Work. In *Intelligent Systems Conference*. Springer, 656–669.

[55] David G. Jansson and Steven M. Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11. https://doi.org/10.1016/0142-694X(91)90003-F

[56] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.

[57] Ece Kamar. 2016. Directions in Hybrid Intelligence: Complementing AI Systems with Human Intelligence.. In *IJCAI*. 4070–4073.

[58] Ece Kamar, Severin Hacker, and Eric Horvitz. 2012. Combining human and machine intelligence in large-scale crowdsourcing. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 467–474.

[59] Simon Keizer, Markus Guhe, Heriberto Cuayáhuitl, Ioannis Efstathiou, Klaus-Peter Engelbrecht, Mihai Sorin Dobre, Alex Lascarides, and Oliver Lemon. 2017. Evaluating Persuasion Strategies and Deep Reinforcement Learning methods for Negotiation Dialogue agents. In *Conference of the European Chapter of the Association for Computational Linguistics*. https://api.semanticscholar.org/CorpusID:216804296

[60] Siwon Kim, Sangdoo Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. ProPILE: Probing Privacy Leakage in Large Language Models. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 20750–20762. https://proceedings.neurips.cc/paper_files/paper/2023/file/420678bb4c8251ab30e765bc27c3b047-Paper-Conference.pdf

[61] Rakesh Kochhar. 2023. *Which U.S. Workers Are More Exposed to AI on Their Jobs?* https://www.pewresearch.org/social-trends/2023/07/26/which-u-s-workers-are-more-exposed-to-ai-on-their-jobs/ About a fifth of all workers have high-exposure jobs; women, Asian, college-educated and higher-paid workers are more exposed. But those in the most exposed industries are more likely to say AI will help more than hurt them personally..

[62] János Kramár, Tom Eccles, Ian Gemp, Andrea Tacchetti, Kevin R. McKee, Mateusz Malinowski, Thore Graepel, and Yoram Bachrach. 2022. Negotiation and honesty in artificial intelligence methods for the board game of Diplomacy. *Nature Communications* 13, 1 (Dec. 2022), 7214. https://doi.org/10.1038/s41467-022-34473-5

[63] Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic Uncertainty: Linguistic Invariances for Uncertainty Estimation in Natural Language Generation. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=VD-AYtP0dve

[64] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User Experience Design Professionals' Perceptions of Generative Artificial Intelligence. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 381, 18 pages. https://doi.org/10.1145/3613904.3642114

[65] Samantha Lock. 2022. What is AI chatbot phenomenon Chat-GPT and could it replace humans? *The Guardian* (December 2022). https://www.theguardian.com/technology/2022/dec/05/what-is-ai-chatbotphenomenon-chatgpt-and-could-it-replace-humans

[66] Zilin Ma, Yiyang Mei, Yinru Long, Zhaoyuan Su, and Krzysztof Z. Gajos. 2024. Evaluating the Experience of LGBTQ+ People Using Large Language Model Based Chatbots for Mental Health Support. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 872, 15 pages. https://doi.org/10.1145/3613904.3642482

[67] Zilin Ma, Yiyang Mei, and Zhaoyuan Su. 2023. Understanding the Benefits and Challenges of Using Large Language Model-based Conversational Agents for Mental Well-being Support. *AMIA ... Annual Symposium proceedings. AMIA Symposium* 2023 (2023), 1105–1114.

[68] D. Mancini-Griffoli and A. Picot. 2004. *Humanitarian Negotiation: A Handbook for Securing Access, Assistance and Protection for Civilians in Armed Conflict*. Centre for Humanitarian Dialogue. https://books.google.com/books?id=LHO5AAAAIAAJ

[69] Aniek F Markus, Jan A Kors, and Peter R Rijnbeek. 2021. The role of explainability in creating trustworthy artificial intelligence for health care: a comprehensive survey of the terminology, design choices, and evaluation strategies. *Journal of biomedical informatics* 113 (2021), 103655.

[70] Bernard Marr. 2024. *What Jobs Will AI Replace First?* https://www.forbes.com/sites/bernardmarr/2024/06/17/what-jobs-will-ai-replace-first/

[71] Bilyana Martinovski. 2010. Emotion in Negotiation. In *Handbook of Group Decision and Negotiation*, Melvin F. Shakun, D. Marc Kilgour, and Colin Eden (Eds.). Vol. 4. Springer Netherlands, Dordrecht, 65–86. https://doi.org/10.1007/978-90-481-9097-3_5 Series Title: Advances in Group Decision and Negotiation.

[72] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyan Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. Human-level play in the game of *Diplomacy* by combining language models

with strategic reasoning. *Science* 378, 6624 (Dec. 2022), 1067–1074. https://doi.org/10.1126/science.ade9097

[73] Tim Miller. 2023. Explainable AI is Dead, Long Live Explainable AI! Hypothesis-Driven Decision Support Using Evaluative AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 333–342. https://doi.org/10.1145/3593013.3594001

[74] Ministry of Electronics and Information Technology, Government of India. 2022. Bhashini: National Language Translation Mission. https://bhashini.gov.in/ Accessed: 2024-12-04.

[75] Michael Muller, Justin D. Weisz, Stephanie Houde, and Steven I. Ross. 2024. Drinking Chai with Your (AI) Programming Partner: Value Tensions in the Tokenization of Future Human-AI Collaborative Work. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne, United Kingdom) *(CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, Article 7, 15 pages. https://doi.org/10.1145/3663384.3663390

[76] Nibble Technology. 2021. The World's Most Experienced AI Negotiation Platform. https://www.nibbletechnology.com/ Accessed on September 9, 2024.

[77] B Nye, Dillon Mee, and Mark G Core. 2023. Generative large language models for dialog-based tutoring: An early consideration of opportunities and concerns. In *AIED Workshops*.

[78] Council on Foreign Relations. 2024. Conflicts to Watch in 2024. (2024). https://www.cfr.org/report/conflicts-watch-2024 Accessed: 2024-05-22.

[79] OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774* (2023).

[80] OpenAI. 2024. GPT-4 Technical Documentation. https://platform.openai.com/docs/models/gpt-4. Accessed: 2024-09-09.

[81] Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2023. Med-HALT: Medical Domain Hallucination Test for Large Language Models. In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*. 314–334.

[82] Andrew Perrault, Fei Fang, Arunesh Sinha, and Milind Tambe. 2020. Artificial Intelligence for Social Impact: Learning and Planning in the Data-to-Deployment Pipeline. *AI Mag.* 41, 4 (Dec. 2020), 3–16. https://doi.org/10.1609/aimag.v41i4.5296

[83] Sara Pieri, Sahal Shaji Mullappilly, Fahad Shahbaz Khan, Rao Muhammad Anwer, Salman Khan, Timothy Baldwin, and Hisham Cholakkal. 2024. BiMediX: Bilingual Medical Mixture of Experts LLM. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 16984–17002. https://doi.org/10.18653/v1/2024.findings-emnlp.989

[84] Sofie Pilemalm. 2022. Barriers to digitalized co-production: the case of volunteer first responders. In *19th International Conference on Information Systems for Crisis Response and Management, Tarbes, France, May 22-25, 2022*. 782–790.

[85] Ritika Poddar, Rashmi Sinha, Mor Naaman, and Maurice Jakesch. 2023. AI Writing Assistants Influence Topic Choice in Self-Presentation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–6.

[86] Lee Rainie, Cary Funk, Monica Anderson, and Alec Tyson. 2022. *AI and human enhancement: Americans' openness is tempered by a range of concerns.* Technical Report. Pew Research Center.

[87] Daria-Brianna Reabciuc, Ana Călugăreanu, and Eduard Balamatiuc. 2023. How AI may replace jobs in the future. In *Conferința tehnico-științifică a studenților, masteranzilor și doctoranzilor*, Vol. 2. 28–31.

[88] Michael Joseph Ryan, William B. Held, and Diyi Yang. 2024. Unintended Impacts of LLM Alignment on Global Representation. In *Annual Meeting of the Association for Computational Linguistics*. https://api.semanticscholar.org/CorpusID:267897555

[89] Zheyuan Ryan Shi, Claire Wang, and Fei Fang. 2020. Artificial Intelligence for Social Good: A Survey. arXiv:2001.01818 [cs.CY] https://arxiv.org/abs/2001.01818

[90] Pao Siangliulue, Kenneth C. Arnold, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Toward Collaborative Ideation at Scale: Leveraging Ideas from Others to Generate More Creative and Diverse Ideas. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) *(CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 937–945. https://doi.org/10.1145/2675133.2675239

[91] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L Glassman. 2023. Where to hide a stolen elephant: Leaps in creative writing with multimodal machine intelligence. *ACM Transactions on Computer-Human Interaction* 30, 5 (2023), 1–57.

[92] Guhan Subramanian. 2023. Ask A Negotiation Expert: Learning From Humanitarian Negotiations Amid International Conflict. *Program on Negotiation at Harvard Law School* (2023). https://www.pon.harvard.edu/daily/negotiation-skills-daily/ask-a-negotiation-expert-learning-from-\humanitarian-negotiations-nb/

[93] Hyewon Suh, Aayushi Dangol, Hedda Meadan, Carol A. Miller, and Julie A. Kientz. 2024. Opportunities and Challenges for AI-Based Support for Speech-Language Pathologists. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. ACM, Newcastle upon Tyne United Kingdom, 1–14. https://doi.org/10.1145/3663384.3663387

[94] Rebecca Sutton and Emily Paddon Rhoads. 2022. Empathy in frontline humanitarian negotiations: a relational approach to engagement. *Journal of International Humanitarian Action* 7, 1 (Dec. 2022), 23. https://doi.org/10.1186/s41018-022-00131-0

[95] Chen Tang, Shun Wang, Tomas Goldsack, and Chenghua Lin. 2023. Improving Biomedical Abstractive Summarisation with Knowledge Aggregation from Citation Papers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 606–618.

[96] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).

[97] Ha Trinh, Koji Yatani, and Darren Edge. 2014. PitchPerfect: integrated rehearsal environment for structured presentation preparation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 1571–1580. https://doi.org/10.1145/2556288.2557286

[98] Abhinav Trivedi, Er. Kanwaldeep Kaur, Chahil Choudhary, Kunal, and Priyashi Barnwal. 2023. Should AI Technologies Replace the Human Jobs?. In *2023 2nd International Conference for Innovation in Technology (INOCON)*. 1–6. https://doi.org/10.1109/INOCON57975.2023.10101202

[99] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 332, 7 pages. https://doi.org/10.1145/3491101.3519665

[100] Remko Van Hoek, Michael DeWitt, Mary Lacity, and Travis Johnson. 2022. How Walmart Automated Supplier Negotiations. *Harvard Business Review* (8 11 2022). https://hbr.org/2022/11/how-walmart-automated-supplier-negotiations Accessed on September 9, 2024.

[101] Skyler Wang, Ned Cooper, and Margaret Eby. 2024. From human-centered to social-centered artificial intelligence: Assessing ChatGPT's impact through disruptive events. *Big Data & Society* 11, 4 (2024), 20539517241290220. https://doi.org/10.1177/20539517241290220 _eprint: https://doi.org/10.1177/20539517241290220.

[102] Allison Woodruff, Renee Shelby, Patrick Gage Kelley, Steven Rousso-Schindler, Jamila Smith-Loud, and Lauren Wilcox. 2024. How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 641, 26 pages. https://doi.org/10.1145/3613904.3642700

[103] Qian Yang, Yuexing Hao, Kexin Quan, Stephen Yang, Yiran Zhao, Volodymyr Kuleshov, and Fei Wang. 2023. Harnessing Biomedical Literature to Calibrate Clinicians' Trust in AI Decision Support Systems. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–14. https://doi.org/10.1145/3544548.3581393

[104] J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 437, 21 pages. https://doi.org/10.1145/3544548.3581388

[105] Zelun Tony Zhang, Sebastian S. Feger, Lucas Dullenkopf, Rulu Liao, Lou Süsslin, Yuanting Liu, and Andreas Butz. 2024. Beyond Recommendations: From Backward to Forward AI Support of Pilots' Decision-Making Process. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 485 (Nov. 2024), 32 pages. https://doi.org/10.1145/3687024

A Negotiation Tools

A.1 Island of Agreements

Island of Agreemens This concept refers to the areas of common ground or shared understanding between parties in a negotiation, despite their overall disagreements. It is based on the paradox that for any disagreement to exist, there must also be some level of agreement or shared perspective. The Islands of Agreement serve as a starting point for dialogue and building trust.

Example A.1 (Access to IDP Camps). In a negotiation between Food Without Borders (FWB) and the Governor of a district regarding access to IDP camps:

- **Agreed Facts:**
  - There are displaced persons from Country A in the no man's land.
  - People are blocked in the no man's land, in a difficult situation in terms of shelter and nutrition.
  - There is little prospect of improvement without immediate access to the displaced.
- **Convergent Norms:**
  - There is a legitimate border between Country A and Country B. B has the right to defend the integrity of its territory and prevent illegal entry.
  - We should not allow people to die from starvation.
  - People have a right to flee armed violence.
- **Contested Facts:**
  - The exact number of displaced persons in the area.
  - The severity of the situation and who is in most urgent need.
  - The presence of armed elements among the civilians.
- **Divergent Norms:**
  - Whether humanitarian organizations have a right of access to people in need under international law.
  - Whether people have a right to enter Country B simply because they flee armed violence.
  - The priority of government security concerns versus humanitarian needs.

A.2 Iceberg and Common Shared Space (CSS)

Iceberg Model This model is used to analyze the position of both the humanitarian organization and the counterpart in a negotiation. It consists of three levels:

(1) **Position (WHAT):** The visible stance or demand at the top of the iceberg.
(2) **Reasoning (HOW):** The logic and interests supporting the position.
(3) **Values and Motives (WHY):** The underlying principles, needs, and drivers at the base of the iceberg.

Common Shared Space (CSS) This refers to the area of potential agreement between parties in a negotiation, where their interests, reasoning, and values overlap or can be reconciled.

Example A.2 (Health for All (HfA) Negotiation). HfA is negotiating with tribal leaders over a hospital closure and staff detention.

HfA's Iceberg:

- **Position (WHAT):** Immediate release of staff and evacuation from District A.
- **Reasoning (HOW):** Ensure staff safety, scale down surgical activities, hand over hospital to third party.
- **Values (WHY):** Humanitarian principles, duty of care, professional health standards.

Tribal Leaders' Iceberg:

- **Position (WHAT):** Keep hospital fully operational under HfA or equivalent.
- **Reasoning (HOW):** Maintain employment for guards, compensate families of injured/deceased guards.
- **Values (WHY):** Community welfare, tribal loyalty, economic stability.

Common Shared Space (CSS):

- Shared concern for community health.
- Recognition of guards' service and sacrifices.
- Desire to maintain reputation and relationships.
- Need for evidence-based decision-making on health needs.
- Importance of addressing emergency medical needs.
- Necessity of community consultation in healthcare planning.

A.3 Red lines and Bottom lines

Red lines Red lines are defined as the outer limits of the possible areas of agreement. They set the parameters within which parties to the negotiation must remain while attempting to maximize their shared benefit. Red lines are generally specified in the mandate given to the negotiator and informed by applicable laws and institutional policies. Red lines cannot be crossed, as doing so would have significant consequences regarding the validity and legality of the agreement and may impact the legitimacy of the negotiator and organization. The negotiator is not allowed to set or revise red lines.

Bottom lines Bottom lines are a tactical tool used by the negotiator to set limits to the conversation when options under consideration show rising risks and diminishing benefits. Bottom lines are under the control of the negotiator as a means to suspend or postpone consideration of additional options below a certain threshold. Before considering options beyond the bottom line, the negotiator may consult again with their hierarchy. The results of this consultation may impact the location of the bottom line.

Example A.3 (Food Without Borders Negotiating Access to IDP Camp). Food Without Borders (FWB) is negotiating access to an IDP camp to distribute food rations. The negotiation involves several key points:

Ideal outcome (Point A): All food rations are distributed only to the affected IDP population based on their nutritional needs. FWB can hire and pay in cash the day laborers of their choice to assist in its work in the IDP camp.

Bottom line (Point B): Food rations should be limited to IDPs but are not necessarily dependent on their individual nutritional needs. FWB could consider including family members of local guards in need as part of the food distribution process, even though they are not recognized as formal IDPs. Direct distribution to the local guards, however, is not permitted.

Red line (Point C): FWB can only distribute food rations to the IDP population and other people in need. It cannot use the food rations as a means of payment for laborers. It further cannot provide any direct assistance to armed personnel.

In this example:

- The bottom line (B) represents the point where FWB negotiators feel the compromises are reaching their limit, but they can still make decisions within their mandate.

- The red line (C) is set by FWB's institutional policy prohibiting the use of food rations as currency or compensation for labor, which the negotiator cannot cross without referring back to higher authorities.

B Survey questions

(1) **How many years of negotiation experience do you have?**
(2) **How often do you use a computer for work?**
  - Once a week or less
  - A few times a week
  - A couple of hours most days
  - Many hours on most days
(3) **What is the highest level of education you have received or are pursuing?**
  - Pre-high school
  - High school
  - College
  - Masters or professional degree
  - PhD
(4) **How often do you use AI?**
  - Once a week or less
  - A few times a week
  - A couple of hours most days
  - Many hours on most days
(5) **If you said yes to the question before, how often do you use AI for work relating to frontline negotiation?**
  - Once a week or less
  - A few times a week
  - A couple of hours most days
  - Many hours on most days

If you are interested in a 1-hour interview to help us learn more about AI in negotiations, please share your email and schedule an interview here:

C Interview Questions: Formative Study

C.1 Part 1: Semi-Structured Interviews with Negotiators

The following questions focus on the negotiation process:

- Tell me about how you prepare a case.
- What was the most difficult part of such a process?
- How do you collaborate with colleagues?

C.2 Part 2: Demo of ChatGPT Generating Iceberg CSS

Demo of an automated system that creates negotiation visualizations. Questions asked after the demo:

- What did you like about the tool?
- Do you think you can use this tool in your workflow?
- Do you have any concerns about using this tool?
- Do you have anything else that you want to mention?

D Interview Questions: Probe Study

Each question is asked for both systems (ChatGPT and probe interface) when time permits.

D.1 1. Impressions of the System

- What was your impression of this system?

D.2 2. Added Value of the System

- What was the greatest added value of this system?

D.3 3. Problems and Improvements

- What did you find problematic or in need of improvement?

D.4 4. Desired AI Features

- What other AI features would you like to see in these systems?

D.5 5. Early-Stage Technology Feedback

- Where else do you think AI could be added to your process?
- Do you see the value of this technology in your work?

D.6 6. Future Development Expectations

- How disappointed would you be if we never developed anything AI-related with you through this collaboration?
- How disappointed would you be if no one ever made technology like this real? Why?

E Participant Demographics

To protect participants' identities, we avoid linking specific details such as years of experience, organizations, and countries of operation that might allow their employers to infer who they are. Instead, we provide a general list of countries of operation and current organizations associated with participants.

List of countries of operation:

- Study 1: Bangladesh, Brazil, Cameroon, Canada, India, Iraq, Kenya, Moldova, Spain, Switzerland, UAE, Vietnam
- Study 2: Bangladesh, Berlin, France, India, Italy, Jamaica, Japan, Lebanon, Malta, Nigeria, Portugal, Switzerland, UAE, UK, USA.

List of organizations:

- Study 1: International NGO Safety Organisation, Humanitarian Association of Dynamic Youths, Doctors Without Borders, International Organization for Migration, United Nations Office for the Coordination of Humanitarian Affairs, International Committee of the Red Cross, World Food Program, United Nations (1 removed because the organization is too small and one can the participant with this information)
- Study 2: Danish Refugee Council, International Committee of the Red Cross, International Federation of Red Cross and Red Crescent Societies, International NGO Safety Organisation, Doctors Without Borders, Norwegian Refugee Council, United Nations, United Nations High Commissioner for Refugees, World Food Program, Center for International Peace Operations. (3 removed because the organizations are too small and one can the participant with this information)

Study	**Participant**	Have you used ChatGPT, Claude 2, or similar LLM-based chatbots?	**How often do you use AI?**	**If you said yes to the question before, how often do you use them for work relating to frontline negotiation?**	**Years of experience**
1	P1	Yes	Many hours on most days	A few times a week	9
1	P2	Yes	Many hours on most days	A few times a week	6
1	P3	Yes	Many hours on most days	A few times a week	8
1	P4	Yes	Once a week or less	Once a week or less	6
1	P5	Yes	Many hours on most days	Once a week or less	9
1	P6	No	–	–	15
1	P7	Yes	Many hours on most days	A few times a week	15
1	P8	No	–	–	10
1	P9	Yes	Many hours on most days	A few times a week	10
1	P10	Yes	Once a week or less	Once a week or less	22
1	P11	Yes	Many hours on most days	A few times a week	10
1	P12	Yes	Many hours on most days	A few times a week	15
1	P13	Yes	Once a week or less	Once a week or less	15
1	P14	Yes	Many hours on most days	Many hours on most days	37
2	P1	Yes	A couple of hours most days	A few times a week	10
2	P2	Yes	Many hours on most days	A few times a week	25+
2	P3	Yes	Once a week or less	Once a week or less	7
2	P4	Yes	A couple of hours most days	A couple of hours most days	7
2	P5	Yes	Once a week or less	Once a week or less	25+
2	P6	Yes	Once a week or less	Once a week or less	25+
2	P7	Yes	A few times a week	Once a week or less	2
2	P8	Yes	A few times a week	Once a week or less	30+
2	P9	Yes	Once a week or less	Never	10
2	P10	Yes	A couple of hours most days	Once a week or less	2
2	P11	Yes	Once a week or less	Once a week or less	20
2	P12	Yes	A couple of hours most days	A few times a week	10
2	P13	Yes	Once a week or less	Once a week or less	9
2	P14	Yes	A couple of hours most days	A few times a week	2
2	P15	Yes	A couple of hours most days	A few times a week	1
2	P16	Yes	A couple of hours most days	A few times a week	10
2	P17	Yes	A couple of hours most days	A few times a week	14
2	P18	Yes	Once a week or less	Once a week or less	14