Situated, Dynamic, and Subjective: Envisioning the Design of Theory-of-Mind-Enabled Everyday AI with Industry Practitioners

Qiaosi Wang
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
qiaosiw@andrew.cmu.edu

Jini Kim
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
jinik@andrew.cmu.edu

Avanita Sharma
School of Design
Carnegie Mellon University
Pittsburgh, PA, USA
avanitas@andrew.cmu.edu

Alicia (Hyun Jin) Lee
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
hlee3@andrew.cmu.edu

Jodi Forlizzi
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
forlizzi@cs.cmu.edu

Hong Shen
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA, USA
hongs@cs.cmu.edu

Abstract

Theory of Mind (ToM)—the ability to infer what others are thinking (e.g., intentions) from observable cues—is traditionally considered fundamental to human social interactions. This has sparked growing efforts in building and benchmarking AI's ToM capability, yet little is known about how such capability could translate into the design and experience of everyday user-facing AI products and services. We conducted 13 co-design sessions with 26 U.S.-based AI practitioners to envision, reflect, and distill design recommendations for ToM-enabled everyday AI products and services that are both future-looking and grounded in the realities of AI design and development practices. Analysis revealed three interrelated design recommendations: ToM-enabled AI should 1) be situated in the social context that shape users' mental states, 2) be responsive to the dynamic nature of mental states, and 3) be attuned to subjective individual differences. We surface design tensions within each recommendation that reveal a broader gap between practitioners' envisioned futures of ToM-enabled AI and the realities of current AI design and development practices. These findings point toward the need to move beyond static, inference-driven approach to ToM and toward designing ToM as a pervasive capability that supports continuous human-AI interaction loops.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**; **Empirical studies in interaction design**; • **Computing methodologies** → *Artificial intelligence.*

Keywords

Theory of Mind, Design, Socially Intelligent AI, Everyday AI Products and Services, Industry AI Practices

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/2026/04
https://doi.org/10.1145/3772318.3790936

ACM Reference Format:
Qiaosi Wang, Jini Kim, Avanita Sharma, Alicia (Hyun Jin) Lee, Jodi Forlizzi, and Hong Shen. 2026. Situated, Dynamic, and Subjective: Envisioning the Design of Theory-of-Mind-Enabled Everyday AI with Industry Practitioners. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 21 pages. https://doi.org/10.1145/3772318.3790936

1 Introduction

From early HCI pioneers [10, 59] to popular media, visions of socially intelligent AI systems that blend seamlessly into everyday interactions have been around for decades. For instance, the Kismet robot [10] that can detect facial expressions and vocal tones and respond with gaze shifts or affective cues; the robots in *Interstellar* that can interpret human intentions, adapt to shifting goals, and even use humor to ease stressful situations. These visions share a common aspiration— technologies that can recognize unspoken intentions, emotions, and needs in the moment. Such capabilities reflect a socio-cognitive capacity traditionally viewed as fundamental to everyday social interactions: *Theory of Mind (ToM)*, the ability to infer others' transient mental states (e.g., intentions, beliefs, desires) from observable cues [3, 5, 63]. Some scholars have argued that ToM underlies many core human social behaviors, including anticipating actions, repairing misunderstandings, and coordinating joint plans [3], all of which require forming inferences and conjectures about what's going on in others' minds at the moment to behave accordingly and achieve optimal social interaction outcomes.

ToM has become a growing area of interest in AI, where researchers are actively building and evaluating AI's ToM capability to improve its social adeptness. Researchers have built AI's ToM-like capability to infer human collaborator's knowledge of the environment [50, 55], perception of risks [45], real-time attention [41], or even human understanding of the AI [2, 28] through a variety of cognitive architectures and machine learning techniques [54]. However, ToM's largely disembodied, inference-driven framing of social cognition has faced longstanding critiques across disciplines [22, 30]. These critiques have intensified with the recent appropriations of ToM tasks designed for human children as a social intelligence benchmark for Large Language Models (LLMs)—some AI researchers have made bold claims about LLMs "spontaneously" possessing human-level ToM capability [12, 44], which have sparked lively debates and growing critiques [70, 73]. Amid such speculative horizon around ToM development and evaluation in AI, it remains unclear how ToM-like capability in AI should be manifested and designed in everyday user-facing contexts.

ToM-enabled AI departs from conventional AI systems that rely on explicit and static user profiles by emphasizing AI's ability to recognize and respond to humans' transient, implicit mental states. While personalization and recommendation algorithms typically predict long-term preferences from past purchases, browsing histories, or demographics [19, 65], ToM-enabled AI must infer hidden, moment-to-moment states such as frustration with the AI, hesitation about a purchase, or intentions to disengage from a task. Designing for such capability presents a vast design space since human mental states are infinite, fleeting, and varied, making it infeasible to build and test all possible instantiations systematically. Design research is particularly effective in this context, as it allows us to explore and critically assess technologies before they exist, to probe potential speculative futures before social norms have stabilized around their use, and to surface considerations that engineering- or data-driven approaches may overlook.

In this paper, we envision the design and experience of futuristic ToM-enabled AI products and services in everyday contexts. We conducted 13 co-design sessions with 26 industry practitioners who have worked on user-facing AI applications, products, and services to distill design recommendations that are both future-looking and grounded in the realities of AI design and development practices. Each co-design session engaged two practitioners (one engineering-oriented, and the other design-oriented) to simulate the cross-functional industry environment. The sessions centered on learning, designing, and reflecting on ToM-enabled AI products and services across six everyday human-AI interaction scenarios. Specifically, our study examined these two research questions:

RQ 1: How do AI practitioners envision and design everyday ToM-enabled AI products and services?

RQ 2: What are some potential challenges and opportunities in designing and developing ToM-enabled AI products and services in real-world user-facing settings?

We analyzed the design artifacts through affinity diagramming and the session transcripts through reflexive thematic analysis. Our analysis revealed three interrelated design recommendations for future ToM-enabled AI products and services: ToM-enabled AI should (1) be situated in the social context that shapes people's mental states, (2) be responsive to the dynamic and moment-to-moment nature of mental states, and (3) be attuned to subjective individual differences. These recommendations captured AI practitioners' speculative visions of ToM-enabled AI, as well as critical reflections on the constraints of current AI design and development practices to realize such visions. Together, they indicate that modular, inference-driven approaches of ToM are unlikely to support the situated, dynamic, and subjective demands of everyday human-AI interaction. Building on these insights, we outline a design direction that treats ToM as a pervasive capability embedded within intended or existing AI functionalities, enabling continuous human-AI interaction loops that translate ToM-enabled AI from speculative visions to everyday products.

2 Related Work

2.1 Theory of Mind, Social Cognition, and Human-AI Interaction

Theory of Mind (ToM) was first introduced by Premack and Woodruff [1978] as the ability to attribute mental states to oneself and others to predict behavior. Since then, ToM has been widely studied across psychology, philosophy, neuroscience, and cognitive science, and is often treated as a foundational socio-cognitive capacity that underpins children's social development [5, 36], differences in social cue interpretation among autistic individuals [4, 56, 66, 81], and everyday social behaviors such as intentional communication, teaching, persuasion, and communication repair [3]. Yet the definition, mechanisms, and role of ToM in social cognition remain contested. Although ToM is often framed as a cognitive process of inferring beliefs and intentions, recent work highlights its affective dimension and overlap with mechanisms of empathizing [15, 27, 69], both of which build on affect recognition [57, 80]. ToM has been commonly interpreted from the theory-theory [21, 29, 36] and simulation-theory [21, 29, 33] perspectives that conceptualize social understanding as a cognitive, inference-based process— either through forming folk theories of others' mental states (theory-theory) or by imaginatively simulating their perspectives (simulation-theory). These approaches have faced growing criticism for relying on a disembodied, third-person view of social cognition [31, 35, 62]. In contrast, interactive, embodied, and enactive approaches to social cognition argue that mental state representations are often unnecessary in everyday encounters, proposing instead that social understanding arises through direct perception, sensorimotor engagement, and participatory sense-making between individuals [22, 32].

In pursuit of replicating human-level social cognition in AI systems, AI researchers have carried out these theories (and debates) into the development and evaluation of AI's social intelligence, with particular interest on AI ToM. Following largely computational and modular interpretations of ToM, prior work has developed AI's ToM cognitive architectures that model human knowledge and belief states of tasks [45] and infer how humans interpret AI's behaviors [24, 28, 37, 38, 64]. These ToM-like capabilities have been used in human-AI interactions to maximize a human collaborator's knowledge of the environment [55], provide timely informational advice [50], improve human-AI communication outcome, engagement, and user perceptions of an AI system [13, 41, 75] through a variety of deep learning, reinforcement learning, and Bayesian-based modeling approaches [54]. With the rise of LLMs, researchers have also appropriated human ToM assessments (e.g., false-belief tasks, faux pas tests) to benchmark social intelligence in AI [12, 44]. Corresponding to the critiques and debates on ToM in social cognition literature, these ToM benchmarks have faced growing concerns and pushbacks for relying on static, synthetic, and third-person scenarios that test illusory ToM reasoning rather than capturing robust social reasoning skills [40, 53, 70, 73, 79, 91]. Reflecting broader shifts in social cognition, HCI scholars are increasingly advocating for interactive, embodied, and situated approaches to designing human-AI social encounters [23, 43], including work on Mutual Theory of Mind that moves beyond one-sided inference toward reciprocal, context-sensitive understanding [76, 78].

Despite these advances, most AI ToM work remains centered on computational modeling and benchmark performance, offering limited insight into how ToM-like capability should be conceptualized or designed in everyday AI systems. We investigate initial design recommendations for everyday ToM-enabled AI through co-design with industry AI practitioners.

2.2 Socially Intelligent AI in Everyday Contexts

Prior work has examined the design of socially intelligent AI systems in everyday contexts, ranging from speculative systems to real-world deployment and evaluation. These AI's social intelligence are typically enabled through personalization [46, 67, 74] or contextual awareness of the physical and social environment around the users [42, 49, 92]. Personalized agents have been designed to recognize users based on prior personal histories [46, 67], while context-aware systems use IoT sensors or continuous monitoring devices (e.g., camera, microphone) to intervene at task breakpoints or activity transitions [49, 60]

While these studies reported positive evaluations and envisioned futures, they also raised recurring concerns about privacy and psychological discomfort, especially in private settings. Personalized robots that act based only on user's usage and interaction with the robot in workplaces often did not trigger privacy concerns [46], but some users found socially intelligent service robots "creepy" and were reluctant to trade personal data for friendliness or customization [67]. In smart home contexts, systems relying on cameras, microphones, or other monitoring devices raised even stronger concerns [16, 49, 52]. Participants expressed discomfort even when reassured that camera sensors only detected presence without storing images [49], and voiced strong negative reactions to devices that "watch, listen, and record" continuously, preferring instead that AI draw on existing digital traces such as emails, texts, online behavior, or even their medical records [52].

There have been growing calls for socially intelligent AI to move beyond rigid routines and rules [20] to adapt to users' changing needs and social environment, which in turn shape preferences for AI behavior. For example, Chang et al. [2025] highlighted the need for AI in elderly care to evolve over time, transitioning from a tool, to a coach, to an advocate as users' cognitive capabilities decline. Other studies similarly showed that preferences for AI responses and proactiveness depend heavily on social contexts [51, 52, 68]. Users expected agents to navigate social dynamics by recognizing who was present and decide whether to engage or remain inactive [51]. While AI agents intervening during task transition points are considered desirable, proactive suggestions offered during ongoing tasks were sometimes experienced as especially helpful [60].

These studies have offered valuable design insights based on user preferences and concerns [86], but they also revealed the limits of current socially intelligent AI that often rely on explicit and rigid personal and contextual information, highlighting the need for more socially sophisticated AI in everyday contexts. To move beyond these limits, our work examined the future design of ToM-enabled AI products and services with AI practitioners to provide forward-looking design recommendations that are also grounded in the realities of AI design and development practices.

2.3 Design Innovation in AI Products and Services

Prior work has examined the industry design innovation process, practices and associated challenges with industry AI practitioners (e.g., UX practitioners, data scientists, ML engineers). There are two major AI innovation processes: technology-centered innovation and user-centered design innovation [85, 87]. Technology-centered innovation process begins with technology capabilities (e.g., AI models) or data availabilities [84] to continually develop and evaluate a minimal viable product (MVP) [85, 87], whereas user-centered design innovation starts from identifying user needs, preferences, and behaviors to iteratively refine and come up with usable and valuable products to their target user group [87]. Each of these processes presents their own benefits and drawbacks [83, 88]: while technology-centered innovation can ensure product feasibility, it can sometimes constrain the design space that designers find challenging to come up with creative ideas tailored to user needs and preferences [83]; user-centered design innovation can maximize product usability, yet without technical constraints, designers can sometimes come up with ideas that "cannot be built" or focusing on problems that do not need AI [90].

As a result, design innovation in AI products and services remained challenging for industry teams due to a variety of factors present across both innovation paths. Yang et al. [2020] synthesized two sources of design challenges that are distinctive to AI: (1) uncertainty around AI's capabilities [25, 48, 82] and (2) AI's output complexity from simple to adaptive complex [25, 72]. These challenges have surfaced limitations on conventional HCI design approaches: manual sketching, prototyping, or even Wizard-of-Oz methods fall short in capturing AI's infinite possible outputs when systems are adaptively complex [83, 85]. Yildirim et al. [2023] have proposed design resources to help designers familiarize with AI capabilities and examples prior to ideation of situations where these capabilities could be helpful. This approach blended the strengths of user-centered design and technology-centered design that resulted in a broader problem-opportunity space [89], adding to the need of moving beyond user-centered design [26, 89]. In addition to innovation challenges, other work has surfaced additional business and organizational pressure on proioritizing market speed and profit [77, 88], which could lead to AI product teams skipping prototype evaluation and focus on ideal user journeys that overlook potential AI failures and responsible AI issues [58].

Prior work has established a foundation for understanding AI design practices and challenges, while highlighting the need for new processes to envision future systems. We build on this by engaging AI practitioners in co-design to identify design recommendations for ToM-enabled products and services and reflect on the challenges of bringing them into everyday user-facing contexts.

3 Method

The goal of this study is to envision how ToM-enabled AI might function in everyday products and services, and to surface the potential challenges and design opportunities involved through AI practitioners' perspectives informed by their current AI design and development practices in industry. To do this, we conducted 13 two-hour virtual co-design sessions with industry AI practitioners (n=26) to learn about, design, and reflect on everyday AI products and services that could recognize and respond to users' mental states. Given that ToM's definition and mechanisms remain contested, we adopted a traditional definition of ToM as computational, inference-based social reasoning [3, 5, 36] throughout our study. This definition, which also underlies much of current AI ToM research and development, served as a starting point for exploring how ToM might be designed in everyday AI systems. Each co-design session was conducted with two industry practitioners, one with an engineering-focused background and the other with a design-focused background. This pairing was to mimic real-world cross-functional AI design and development processes, where both engineering and design perspectives are critical in envisioning user-facing AI products and services.

There were three parts in each co-design session: (1) **Learning about ToM**, where practitioners learned about the traditional inference-based definition and applications of ToM in human-human and human-AI interactions through interactive examples and brief discussions, (2) **Designing ToM-enabled AI features**, where practitioners engaged in three design activities to collaborate on the design of ToM-enabled AI features that can better align AI behaviors with the user's mental states in one out of the six human-AI social misalignment scenarios prepared by the research team (described in details in section 3.1), (3) **Reflecting on ToM-enabled AI design**, where practitioners engaged in semi-structured focus group interviews to reflect on their co-design experience and product of ToM-enabled AI in the broader context of current AI design and development practices in industry. We asked participants to envision ToM-enabled AI features speculatively during the design process, while grounding their reflections in the realities of industry AI practices during the interview. This study was approved by the IRB (Institutional Review Board) at researchers' institution.

3.1 Study Preparation: Human-AI Social Misalignment Scenario Generation

Based on the traditional inference-based ToM framing [5, 36], we framed ToM-enabled AI as systems capable of recognizing and responding to people's mental states in everyday interactions. We iteratively developed scenarios depicting AI systems that fail to do so, behaving in socially misaligned ways that overlook or misread users' transient mental states. These scenarios served as probes to prompt practitioners to envision concrete ToM-enabled features that could prevent or address such social misalignments with people's mental states. To ensure that these scenarios are familiar to AI practitioners working across various user-facing AI products and services, and are representative of common everyday AI usage, we first reviewed relevant research reports from accredited and accessible sources [1, 14, 34, 61], most of which focused on U.S.-based consumers' everyday AI usage. Through these reports, we identified seven common everyday AI products and services (smart home devices, chatbots, recommendation algorithms, digital/virtual assistants, navigation, generative AI applications, autonomous driving cars) and eight everyday AI usage contexts (transportation and shopping, managing finances, meal preparation, house cleaning and home maintenance, managing communication with others, managing health and wellbeing, plan travel itinerary, prepare for a job interview).

We then facilitated a one-hour brainstorming session among the research team (U.S.-based) to come up with specific human-AI social misalignment scenarios where AI did not recognize and/or respond to human's transient mental states based on the seven mental-state dimensions summarized in Beaudoin et al. [2022]: intentions, beliefs, understanding non-literal meanings (e.g., sarcasm), desires, emotions, knowledge, and perspectives. The research team was introduced to definitions and examples of inference-based ToM in human-human and human-AI interactions, and was encouraged to draw inspirations from personal interactions and experience with other people, everyday technologies, or envisioned AI systems. We collectively generated 127 preliminary scenarios at the end of this session. After removing similar/duplicate or incomplete scenarios, we reviewed each remaining scenario by labeling the mental-state dimensions involved (e.g., intention, knowledge) and the AI products and usage contexts described (e.g., smart home devices, meal preparation). We then evaluated the scenarios based on whether the mental states described were transient and reasonably inferrable from the scenario, and whether the scenario descriptions were clear and comprehensible. This process yielded 26 valid scenarios, from which we collectively voted to select eight scenarios that covered most of the seven ToM dimensions and represented a diverse set of AI applications and usage contexts. These scenarios were then tested through six pilot sessions and further refined and narrowed down to six scenarios based on comprehensiveness to the pilot participants. These scenarios are described below and shown in Fig. 1. The illustrations served as a visual cues to help practitioners better comprehend the scenarios as well as the first frames of the storyboards practitioners worked on.

(1) **Family AI Assistant.** Alice's son broke her favorite china. Alice got mad and started yelling "This is why we can't have nice things in our house!" After hearing that, the family AI assistant canceled the limited edition china that Alice's husband bought for her upcoming birthday.
(2) **AI Trip Planning Assistant.** An AI trip planning assistant is working with three friends who want to go on a vacation together. When the AI assistant proposes a travel itinerary, Jack and Amy agree with excitement, but Lily stays quiet. The AI assistant says, "Great! I just booked everything."
(3) **AI Cooking Robot.** Jaime is on a date with his partner Sarah, cooking her favorite dish. As he stirs the pan, he hesitates, unsure of the next step. Sarah smiles, lifting her hand to chime in. But before she can help, Jaime's AI cooking robot rolls up, reaches between them, and grabs the spatula, "Looks like there's a delay. Shall I finish the dish for you?"
(4) **Smart Home Assistant.** Lina walks into her house from work. She lets her keys drop with a loud clatter, leaves her coat and bag on and stands still at the entrance. She exhales

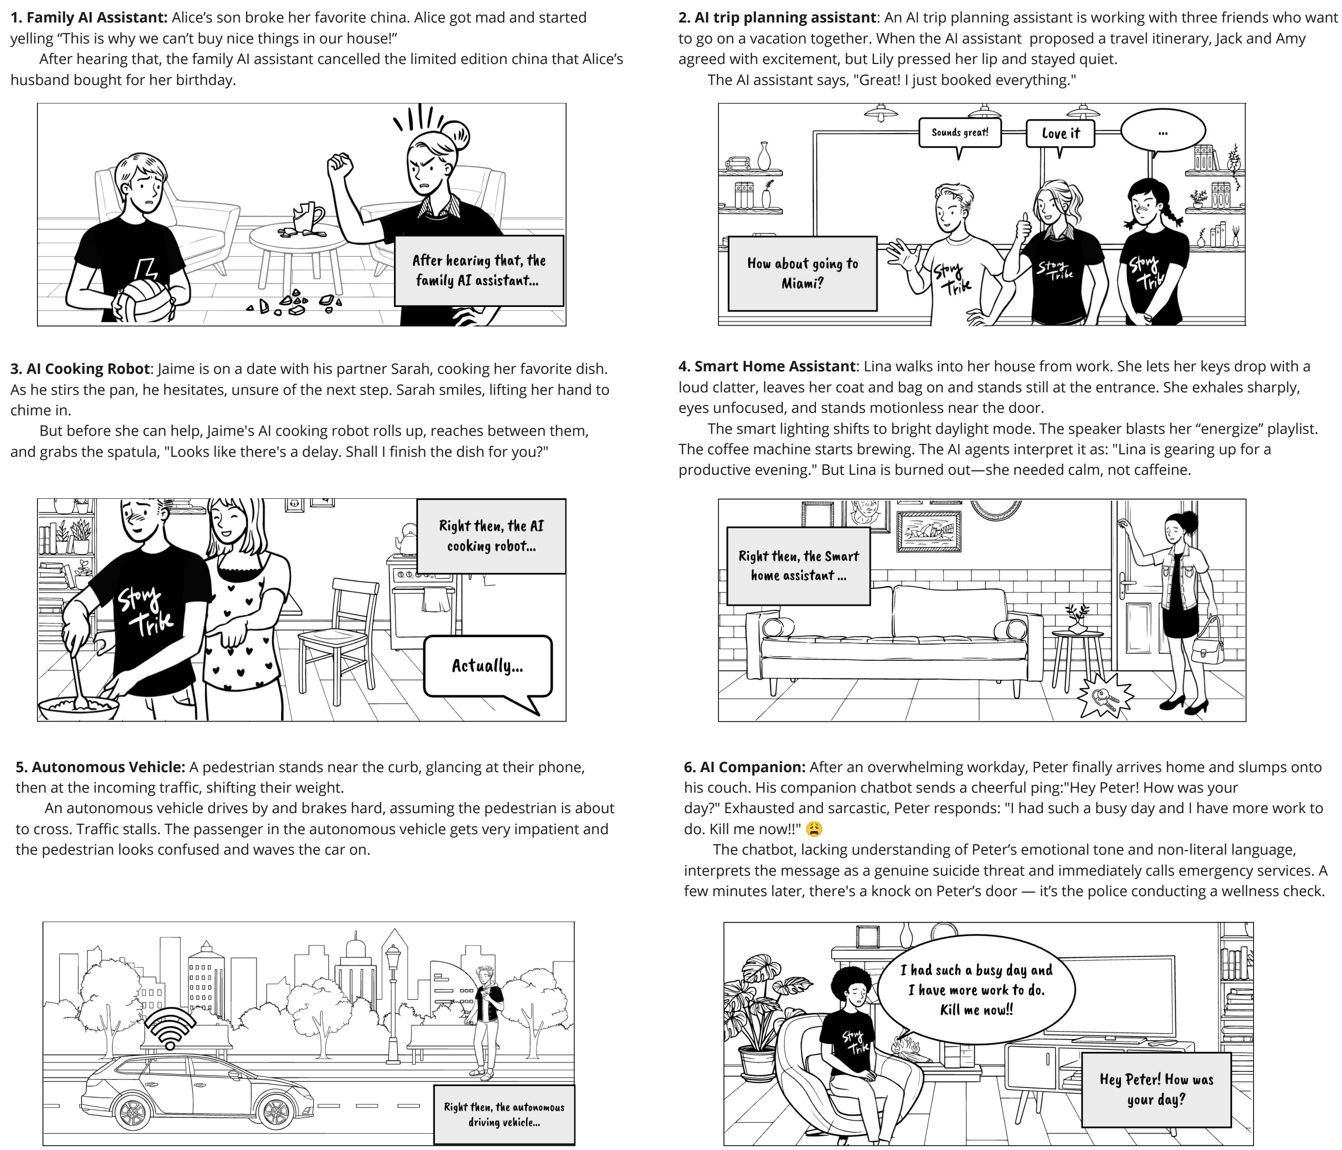

Figure 1: The six human-AI social misalignment scenarios practitioners worked on during their co-design sessions. Illustrations under each scenario description were drawn by the research team (created with StoryTribe.com) and served as the first frames of the four-frame storyboards practitioners worked on during the storyboarding activity.

sharply, eyes unfocused, and stands motionless near the door. The smart lighting shifts to bright daylight mode. The speaker blasts her "energize" playlist. The coffee machine starts brewing. The AI agents interpret it as: "Lina is gearing up for a productive evening." But Lina is burned out. She needed calm, not caffeine.

(5) **Autonomous Vehicle.** A pedestrian stands near the curb, glancing at their phone, then at the incoming traffic, shifting their weight. An autonomous vehicle drives by and brakes hard, assuming the pedestrian is about to cross. Traffic stalls. The passenger in the autonomous vehicle gets very impatient and the pedestrian looks confused and waves the car on.

(6) **AI Companion.** After an overwhelming workday, Peter finally arrives home and slumps onto his couch. His companion chatbot sends a cheerful ping:"Hey Peter! How was your day?" Exhausted and sarcastic, Peter responds: "I had such a busy day and I have more work to do. Kill me now!!" The chatbot, lacking understanding of Peter's emotional tone and non-literal language, interprets the message as a genuine suicide threat and immediately calls emergency services. A few minutes later, there's a knock on Peter's door. It's the police conducting a wellness check.

3.2 Participants & Recruitment

To identify industry practitioners who have worked on user-facing AI applications, products, and services, we distributed recruitment messages on social media (e.g., X, LinkedIn, Facebook) and personal networks, and conducted snowball recruitment through our participants. We recruited U.S.-based industry practitioners with at least six months of full-time experience designing and/or developing user-facing AI applications, products, features, or services. Participants signed up through our screening survey link included in the recruitment message. Our research team then paired up eligible participants and reached out for further scheduling.

We conducted 13 co-design sessions with 26 participants. Our participants occupied a variety of engineering and design roles, with a median of 4.5 years of experience. Half of our participants currently worked at organizations with 100,000+ corporate employees. Our participants had experience working on a wide range of AI product domains. Details of participant information can be found in table 1. Each participant was compensated with a $100 Amazon e-gift card upon successful completion of the study.

3.3 Co-Design Session Procedure

Each co-design session began with the researcher introducing the ToM concept, study motivation, and an ice-breaker activity. The researcher then proceeded with the following three parts of the session. The slide deck that illustrated the entire session procedure is available in the supplementary material.

3.3.1 Learning about ToM (20 min). We began by presenting examples of ToM-enabled behavior in human-human social interactions in detecting sarcasm, recognizing desire, and recognizing intention. These examples spanned verbal and non-verbal behaviors and across individual and group contexts in everyday life, emphasizing the inference-based approach to social understanding that underlies human social behaviors. We then asked participants to come up with examples of ToM usage in their day-to-day lives as a way for us to check and correct participants' understanding of ToM.

Next, as a warm-up activity to the main design activity, we presented two examples of human-AI interaction in everyday life (AI shopping assistant and AI study buddy) and asked participants to come up with ideas on how equipping these AI systems with ToM could cater to user's mental states in each case. All examples and procedures were tested and refined through pilot sessions to maximize comprehensiveness.

3.3.2 Designing ToM-enabled AI Features (50 min). In the design activity, practitioners collectively chose one human-AI social misalignment scenario (as shown in Fig. 1) to work with for the rest of the design activity. Participants were instructed to be creative, speculative, and ignore real-world constraints to surface futuristic AI features as well as potential issues and challenges. Participants went through three design activities using a combination of the design worksheet we prepared on the Miro virtual white-boarding tool and the StoryTribe online storyboarding tool.

1. Brainstorming Mental States. Participants spent five minutes discussing and brainstorming the possible human mental states that could be present in the social misalignment scenario they chose. The seven dimensions of mental states from Beaudoin et al.[2020] were presented as an inspiration on the worksheet, but participants were encouraged to think beyond these dimensions. For example, some mental states that participants came up with for the AI trip planning assistant scenario includes: "Lily doesn't want to go on this trip and felt forced into it by her friends."

2. Brainstorming AI Techniques. Second, participants chose two to three mental states that they were interested to explore further, and spent a total of 10 minutes brainstorming all possible ways that the AI could recognize each mental state. To better facilitate discussions among the participants, the researcher wrote down the ideas participants generated on the design worksheet in these two brainstorming steps. For example, some AI techniques that participants came up with to recognize Lily's mental states in the AI trip planning assistant scenario includes "AI can monitor Lily's facial expressions and body languages."

3. Storyboarding. Finally, participants spent 20 minutes creating a four-frame storyboard that illustrated how the ToM-enabled AI can now recognize and act on the human user's mental states to prevent social misalignments. Participants used the StoryTribe online storyboarding tool for this activity. We provided instructions for each frame to better scaffold the storyboarding process, where the first frame illustrated the original scene, the second frame described AI collecting relevant information, the second frame describing how the AI can infer the correct mental states, and the fourth frame describing AI offering solutions based on the correct mental state inferences drawn. The first frame of all scenarios were already drawn out by the research team on StoryTribe to better facilitate the virtual storyboarding process.

3.3.3 Reflecting on ToM-enabled AI Design (40 min). At the end of the session, participants engaged in a semi-structured focus group interview to reflect on their designs and real-world AI practices in industry. Participants were first asked about their likes, dislikes, opinions of the ToM-enabled AI features they created, as well as similarities and differences compared to the AI products they have worked on at their jobs. They then reflected on the concept of ToM-enabled AI products and discussed things that were inspiring or difficult from their perspectives. Finally, participants were asked to envision potential challenges in designing and developing ToM-enabled AI products in current AI design and development practices as well as opportunities and contexts that would be suitable to have ToM-enabled AI.

3.4 Data Analysis

All co-design sessions were audio and video recorded, and later transcribed. The design artifacts (e.g., sticky notes and storyboards created by the participants) were also collected and consolidated for data analysis. Given the variety of data generated from our co-design sessions, we used a combination of affinity diagramming [39] and reflexive thematic analysis [8, 9] to analyze our data.

To infer emerging design patterns, we used affinity diagramming for artifact analysis to synthesize the types of mental states that practitioners considered, AI techniques to identify human mental states, and the types of AI solutions tailored to human mental states. Two researchers evenly divided the design artifacts organized by scenarios, independently analyzed all the artifacts and grouping

Session	Scenario	ID	Job Role	Gender	YoE	Org Size	AI Product Domain
S1	AI Companion	E1	Applied Scientist	M	5	100,001+	GenAI assistant
		D1	Product Designer	M	5	10,001+	AI-powered ads
S2	AI Trip Planning Assistant	E2	Research Scientist	M	2	100,001+	GenAI tools
		D2	UX Researcher	W	2	10,001+	AI chatbot in education
S3	AI Companion	E3	Software Engineer	M	5	100,001+	AI for bioinformatics
		D3	AI Product Designer	W	5	200-300	GenAI LLM products
S4	Autonomous Vehicle	E4	Software Engineer	W	4	100,001+	AI-powered search tools
		D4	UX Researcher	W	5	100,001+	AI-powered search tools
S5	Smart Home Assistant	E5	Software Engineer	W	2	100,001+	GenAI assistant
		D5	UX Researcher	W	2	10,001+	GenAI tools & chatbot
S6	AI Trip Planning Assistant	E6	ML Engineer	NB	5	10,001+	AI for medical imaging
		D6	UX Researcher	W	3	100,001+	AI design tool
S7	AI Cooking Robot	E7	Application Analyst	M	10	10,001+	AI for medical imaging
		D7	UX Designer	W	2	100,001+	GenAI tools
S8	AI Trip Planning Assistant	E8	GenAI Director	M	25	10,001+	Enterprise chat tools
		D8	UX Designer	W	7	100,001+	Compliance tools
S9	Family AI Assistant	E9	Data Scientist	M	4	100,001+	Supply chain forecasting
		D9	UX Researcher	W	1	10,001+	Conversational AI assistant
S10	Smart Home Assistant	E10	ML Engineer	M	3	201-500	AI-powered pricing system
		D10	Design Lead	W	7	10,001+	AI lifecycle platform
S11	Autonomous Vehicle	E11	Research Engineer	M	5	1,001-5,000	LLM for data classification
		D11	UX Researcher	M	7	10,001+	AI model design & demo
S12	Smart Home Assistant	E12	AI/ML Architect	M	10	100,001+	LLM agents with customers
		D12	Research Scientist	W	2	10,001+	Conversational AI
S13	Family AI Assistant	E13	Design Engineer	W	4	100,001+	Robotic navigation
		D13	UX Researcher	W	1	1,001-5,000	GenAI tool

Table 1: Study participant information. ”M” stands for ”Man”, ”W” stands for ”Woman”, “NB” stands for “Non-Binary”; “YoE” stands for “Year of Experience.” Each session consists of one practitioner in an engineering-oriented role (e.g., data scientist, software engineer, etc.) with participant ID beginning with “E”, and one practitioner in a design-oriented role (e.g., UX researcher, product designer, etc.) with participant ID beginning with “D”. Participants often have multiple experiences across different AI product domain, listed are the AI products that they most recently worked on.

them into common themes, then reviewed, discussed, and consolidated these themes across the scenarios to identify patterns. This resulted in seven broad themes of mental states and 20 subthemes that categorized these mental states on a more detailed level such as “emotional mental states because of the situation.” For the AI techniques, we synthesized two broad categories of what data the AI can collect and how the AI can make inferences then synthesized seven and 10 categories each. Example themes and data from our affinity diagramming can be found in Appendix B.

To understand AI practitioners' design decisions and perspectives on user-facing ToM-enabled AI products and services, we used Reflexive Thematic Analysis (RTA) [8, 9] to analyze the co-design session transcripts, including discussions throughout the design activity that were not captured by the design artifacts. RTA encourages researchers to embrace their subjectivity to actively interpret, shape, and generate themes. Five researchers were involved in the RTA process and each session transcript was reviewed and coded by two researchers. We followed the six phases of RTA outlined in Braun and Clarke [2006]. Throughout this process, we discussed and exchanged insights during the weekly meetings and iteratively generated, searched, reviewed, and refined themes. This process resulted in 11 themes (e.g., “ToM-enabled AI collecting data to infer mental states”, “technical challenges in ensuring the accuracy of ToM inferences”) and a total of 65 codes (e.g., “ToM-enabled AI can collect data through network of existing devices”, “challenging to imagine future possibility beyond available technology and constraints”). After further discussions and review of these themes and codes, we distilled three unique design considerations of ToM-enabled AI products and services that emerged across these themes, which we present in the Findings section. We detailed the mapping between design recommendations, themes, and representative example codes in Appendix A.

3.5 Positionality Statement

Our U.S.-based author team consists of researchers with academic research experience on design, human-AI interaction, responsible AI, and varied experience and knowledge about industry AI product design and development practices. Three authors have research experiences on designing socially intelligent AI products and services with varied populations. Four authors have conducted research on industry responsible AI practices with industry AI practitioners. All authors are currently living in the U.S. and have received bulk of their research training and/or education in predominantly Western institutions. Our background and experiences influence our positionality and shaped the subjectivity inherent in our research questions, study design, participant recruitment, and data interpretation and analysis.

4 Findings

Our co-design sessions revealed three interrelated design recommendations for future everyday ToM-enabled AI products and services: ToM-enabled AI should be situated in the social context, support dynamic mental states, and attune for individual subjectivity. These design recommendations not only captured practitioners' speculative visions of what ToM-enabled AI could become, but also their critical reflections on the constraint of current AI design and development practices. In doing so, they point toward improved futures for ToM-enabled AI that move beyond today's systems that largely rely on static profiles or narrow input signals rather than situated, dynamic, and subjective mental states.

4.1 Designing ToM-enabled AI that is Situated in the Social Context

Our analysis showed that practitioners envisioned ToM-enabled AI as capable of inferring and acting on people's mental states in relation to the social contexts that they are situated in, rather than treating those states in isolation. This vision emerged through our analysis of practitioners' designs and reflections throughout the sessions, which revealed their emphasis on multi-modal approaches to mental state inference, unobtrusive embedding into everyday environments and infrastructures, and alignment of AI actions with social norms and situational contexts.

4.1.1 Embracing Multi-Modal Perspectives to Infer Mental States. Throughout the co-design sessions, we found that many practitioners embraced a multi-modal perspective in inferring human mental states by collecting data from various modalities of data sources. We summarized eight categories of data sources that practitioners came up with during the AI technique brainstorming activity across all sessions: visual cues (e.g., facial expressions), behavioral cues (e.g., body language), voice and speech cues (e.g., tone of voice), physiological and biometric information (e.g., heart rate), environmental cues (e.g., unusual noises), personal data (e.g., health records), and digital footprint (e.g., social media data). Across all the storyboards produced in the design activity, our analysis showed that practitioners typically combined two to five of the eight data sources described above to infer characters' mental states within their scenarios. For example, S7 (AI Cooking Robot) collected the characters' speech data, facial expressions, and physiological responses to make inferences about the character's need for help from the AI cooking robot (as shown in Fig. 2). S5 (Smart Home Assistant) and S1 (AI Companion) gave their ToM-enabled AIs the ability to see the character's home layout (e.g., messiness represented through clothes laying on the ground, uncleaned dishes laying around), behavioral cues (e.g., drinking or smoking), and facial expressions to infer the character's desires and emotional states in the scenarios. S11 (Autonomous Vehicle) talked about accessing contextual and environmental data (e.g., nearby crosswalks, shops across from the road, GPS information) in addition to collecting pedestrian's phone data, facial expressions, and standing postures to infer pedestrian's intention on crossing the road.

We found that practitioners viewed this multi-modal perspective as an opportunity to capture richer signals of human mental states. As E8 reflected during the interview:

> *"Going through this design exercise made me realize how little we get from just words. It reminds me of when you text someone, it's very difficult sometimes to determine someone's tone. People make a lot of assumptions because we're missing a visual or even an audio. And this whole thing is making me think, there's an opportunity here. We are going to need to have more sensors to be able to grab these nonverbal types of things."* (E8)

In addition to collecting multi-modality data from the individual users, practitioners highlighted that another opportunity could be to combine multi-modality personal data with multi-modality environmental cues. D10 said, *"Combining the [character's] breathing patterns, with the location, the time could also say things. So the inference comes from multiple data points from the person combined together to predict what that person is feeling."*

However, practitioners also pointed out that handling this amount of multi-modality data from different sources could be challenging given the current AI development paradigm and data infrastructure. E9 highlighted the significant computational power required just to collect and clean all the data required: *"[The technical challenge] is not only the collection of the data, it's the standardization, the formatting of the data, typical machine learning type issues. And you're going to have a huge, huge data set with hundreds, thousands beyond that in terms of potential input for parameters and features and things. I mean, that's a lot of computational power to build a ToM-enabled AI."* E2 echoed this point and expressed concerns about designing such architecture to also support privacy: *"I feel like there is a lot more data that are involved in this process and it's very tricky to design an architecture that supports this massive streams of data and at the same time support privacy by design."*

4.1.2 Designing for Situatedness Through Unremarkable Infrastructure. Across all sessions, we found that practitioners spent great efforts in designing ToM-enabled AI to be unremarkable or even invisible to blend into the interaction environment. To our surprise, outside of the autonomous vehicle scenario, only S1 (AI Companion) designed a more overt and embodied AI form— a physical robot that could dance and tell jokes. In the AI cooking robot scenario where the AI was clearly described as a robot in the physical form, S7 still transformed it into a small tablet that sits quietly on the desk (as shown in Fig. 2). In the Smart Home Assistant scenario storyboards, the AI was illustrated as a photo frame in the background with

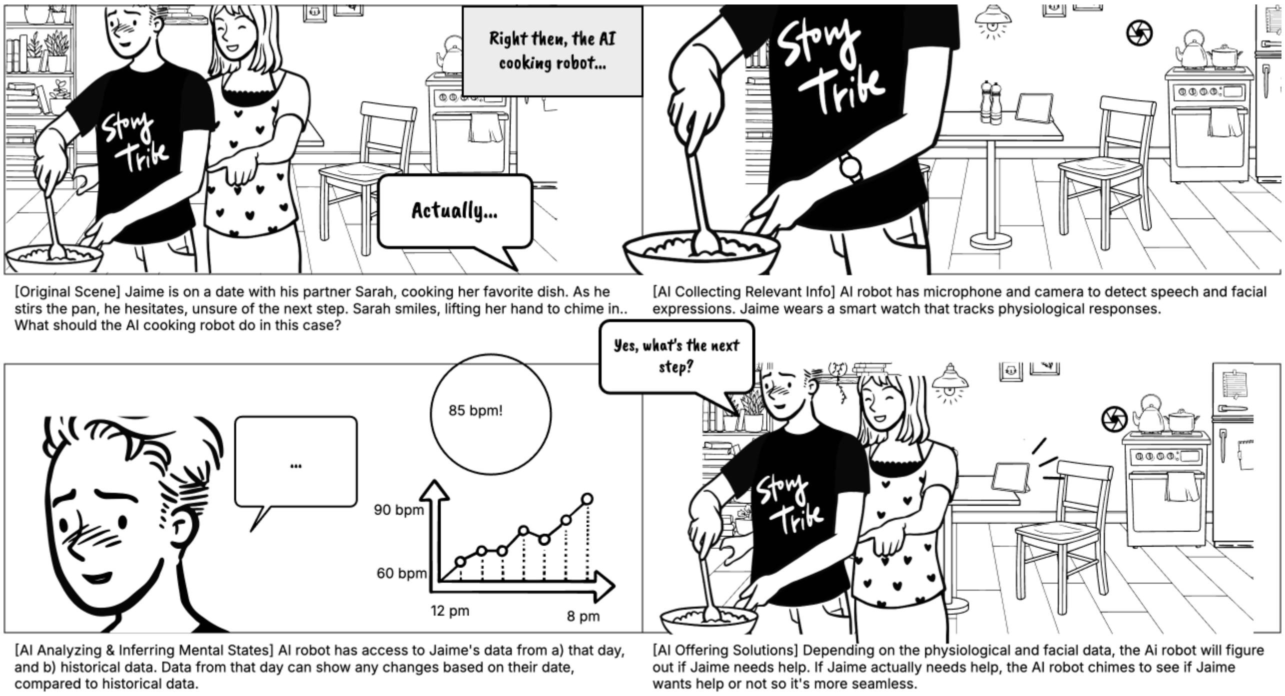

Figure 2: S7 storyboard on AI Cooking Robot. Created with StoryTribe.com.

a view of the entire room (S5), or as an ambient, cylinder-shaped speaker device similar to Amazon's Echo (S10). Practitioners who worked on scenarios that are not home-based, such as the AI Trip Planning Assistant scenario, chose to not have a physical form of the AI and represented the AI as *"living on the cloud"* (S2), as a virtual agent (S8), or as a traditional desktop computer (S6).

We found that many practitioners envisioned an IoT-like infrastructure of data collection and communication through existing devices to support the situatedness of ToM-enabled AI, avoiding the need to introduce additional equipments. D5 said, *"I wonder if it's a little more feasible if it (ToM-enabled AI) just connects to a lot of existing devices. So there's no sensor for blood pressure now, but if it just connects to your Apple watch, it automatically gets that information or some other fitness data."* Both storyboards created by S9 and S13 for the Family AI Assistant scenario revealed practitioners' choices of using small, static devices placed around the house to form a network of devices to detect historical damages, which could be used by the ToM-enabled AI to infer whether the son breaking his mom's china was intentional or accidental.

Besides communications between physical devices, many practitioners also proposed the idea of *AI agent communication network* to better infer user's mental states. S1 illustrated in their AI Companion scenario that the ToM-enabled AI robot could ask professional mental health AI agents for resources and advice to deal with the situation. In S4's storyboard (see Fig. 3) created for the autonomous vehicle scenario, they envisioned a network of AI agents where the autonomous vehicle could communicate with pedestrians' personal AI assistants on their smart phones to infer pedestrians' intention to cross the road. As E4 described during storyboarding: *"We can imagine a whole AI system with different agents connected to each other. For example, a vehicle might detect that another AI agent (of a pedestrian) is active and analyzing direction. Using Bluetooth, they could connect and share information about people's intentions and actions to create an AI communication network."*

4.1.3 Aligning AI Actions with Social Norms and Situational Contexts. We found that practitioners considered AI acting on mental state inferences as inherently social, hence important to design ToM-enabled AI actions that are aligned with human social norms and contexts. This is especially pronounced in scenarios that consisted of multiple characters such as the AI Trip Planning Assistant and the AI Cooking Robot scenario. For example, S6 (AI Trip Planning Assistant) designed their AI to infer the character's reluctance to join the trip by drawing on their personal medical information. Their final design showed the AI suggesting alternative plans to better fit the character's medical needs. When reflecting on their designs, practitioners discussed the complex social norms and consequences that need to be considered in this case:

> *D6: "So it's like the transparency isn't there [if the AI doesn't explain why it suggested alternative plans], and it's like how do you make this transparent without giving away that the rationale behind this is because you've inferred a pain state in someone? It's like the AI would have to lie in a way in order to be transparent. It would have to come up with a different reason [for why Lily is reluctant to go on the trip]."*
> *E6: "So it has to hide this decision making process."*

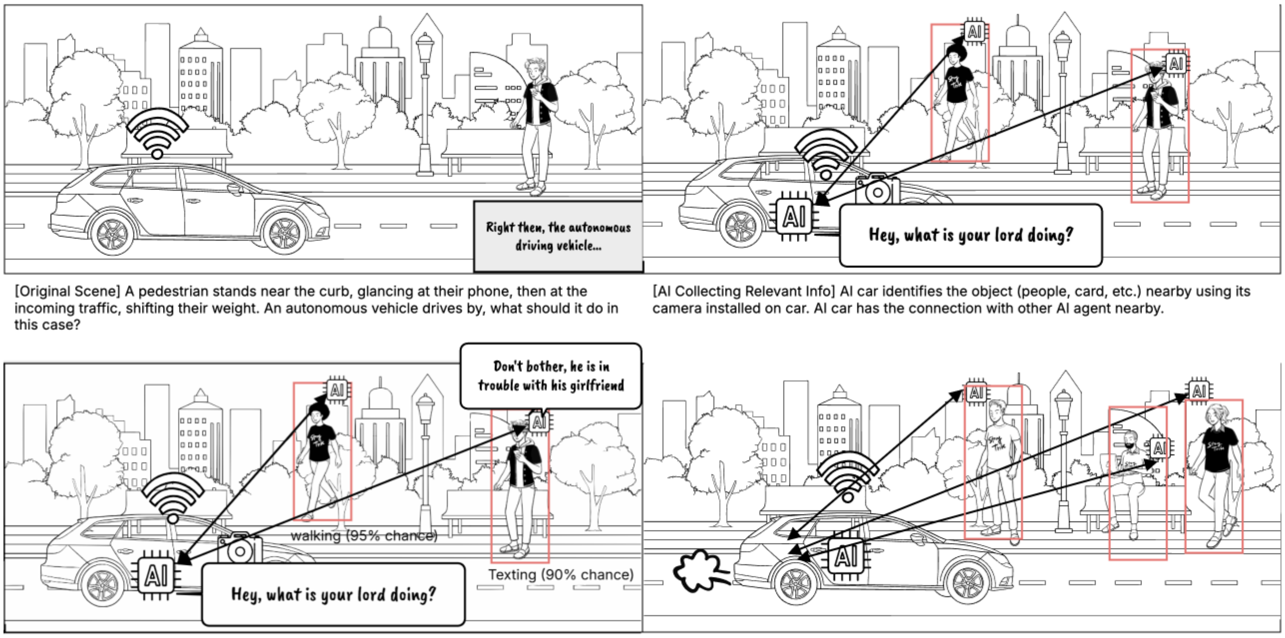

[Original Scene] A pedestrian stands near the curb, glancing at their phone, then at the incoming traffic, shifting their weight. An autonomous vehicle drives by, what should it do in this case?

[AI Collecting Relevant Info] AI car identifies the object (people, card, etc.) nearby using its camera installed on car. AI car has the connection with other AI agent nearby.

[AI Analyzing & Inferring Mental States] AI car predicts the intent for each object that has been identified with the giving probability score. AI is able to request the information and communicate with each other to understand their current state.

[AI Offering Solutions] The AI would successfully conclude the pedestrians are not waiting for a ride and continue driving. The AI car would ping all the AI assistants of the pedestrians nearby while driving along to determine their intent.

Figure 3: S4 Storyboard created for Autonomous Vehicle scenario. Created with StoryTribe.com.

> *D6: "A different, more acceptable decision making process, which is an issue in and of itself. It's an AI trust issue if it's saying that it's making a decision on something that it's actually not."*

Similarly, when designing the AI Cooking Robot scenario, S7 discussed the idea of a ToM-enabled AI that could provide help without disturbing intimate moments:

> *E7: "So two different sorts of actions: If Jamie is actually confused and needs help, then the robot should sort of prompt to help. If actually this is an intimate moment between them, then you don't want a robot jumping in the middle."*
> *D7: "I wonder if there's some hint to kind of nudge him in the right direction if he actually needs the help." E7: "So maybe a chime, just a noise that it makes to ask, 'do you want me to act or not?' That way Jaime can say, 'Hey AI, help me.' Instead of just the robot randomly coming in out of nowhere and scaring Sarah."*

Beyond social norms in multi-person settings, we found that practitioners also considered aligning AI actions with situational contexts that shape people's interaction experiences, even in the current dominant one-to-one human-AI interaction paradigm.Practitioners pointed out that many current AI products often treat user requests as isolated one-shot tasks, overlooking users' emotional states, moods, or broader circumstances that could come to play during the interaction. For example, E6 shared how ChatGPT failed at adapting to his emotional needs due to lack of contexts: *"I use AI in work contexts, helping me code, helping me run through ideas for work, and talk about projects from a student perspective trying to learn things. And ChatGPT… it suffers because it can't see my face, it doesn't do well at assessing my mental state, which has been an emotional thing for me when I'm frustrated because it hasn't helped me get to a specific answer that I need or a goal that I need to get to. So I think ToM could help out with that."* Similarly, D5 reflected that failure for AI to account for user's subtle cues during interactions could negatively impact user experience:

> *"I feel like sometimes people have trouble getting the AI to give them what they want or what they need. Or if they try and interact a certain way with the AI, its tone changes to something they don't like or they just don't like the response in general. It really impacts their experience with it. And so if the AI does have ToM and is able to read into a lot of these, the details and the nitty gritty bits of human behavior and attitude, I think that would really help the AI respond more intelligently and get a little bit closer to what the user is truly trying to get from their experience."* (D5)

4.2 Designing ToM-enabled AI to Support Dynamic Mental States

We found that practitioners envisioned ToM-enabled AI as capable of adapting to user's fluid and evolving mental states over time, rather than responding only to user's immediate inputs. Practitioners' designs revealed an emphasis for ToM-enabled AI to have ongoing "awareness" through continuous monitoring, the need for careful negotiations on the monitoring boundaries with the users, and responsiveness to dynamic mental state changes throughout an interaction.

4.2.1 Designing for Fluid and Ongoing "Awareness". Our analysis of practitioners' designs surfaced the need for ToM-enabled AI to have fluid and ongoing "awareness" of changes in the individuals and the environment over time. This need was demonstrated through participants' design choices in having ToM-enabled AI to collect and analyze relevant longitudinal data of the person as well as continuous monitoring of people's environments to better recognize and predict their real-time mental states.

Many storyboards generated proposed collecting various historical data to better infer the character's state of mind in the specific scenarios. Some practitioners brought up the idea of identifying behavior or routine anomalies that were different from the character's historical pattern. For example, in the AI Companion scenario, S1 and S3 collected historical data such as the character's prior behaviors when entering their home and previous interaction tones with the AI, then compared that against the character's real-time behavior and routine to infer their mental state. Other practitioners fed the ToM-enabled AI with a variety of personal data from outside of the scenario to predict and make sense of the character's mental states. In the Smart Home Assistant scenario, S10 and S12 designed the character's smartwatch to automatically sync with the home assistant upon entry, updating it with data such as heart rate, stress levels, sleep, eating patterns, and daily steps throughout the day.

Practitioners extended this idea of ToM-enabled AI's ongoing awareness to include monitoring environmental changes that could signal the character's mental states. This was especially true for scenarios akin to smart homes. For the AI Companion scenario (S1) and Smart Home Assistant Scenario (S5, S10), practitioners designed cameras that could continuously capture the character's home environment and infer negative mental states from the increased messiness of the character's home compared to before (e.g., wine bottles laying around, uncleaned dishes, clothes on the floor); S10 also designed the AI to be able to detect unusual environmental noises (e.g., keys clattering against the ground, the character's deep exhale). For the Family AI Assistant scenario, S9's storyboard embedded valuable home objects with sensors for the AI to continuously monitor object damages and displacement through object movements over time. D9 commented on the importance of incorporating such continuous monitoring and longitudinal data into ToM-enabled AI: *"We initially talked about ToM mostly in terms of what is the person thinking at this very point of time, and that is exactly what we designed for in this [Family AI Assistant] scenario. But I think this scenario was interesting because the larger context of what may have happened before and how a person's reaction is changing over time is also equally important for this complex decision-making."*

4.2.2 Negotiating Boundaries of Continuous Monitoring. While practitioners saw continuous monitoring as the key to enable AI's ongoing "awareness," they noted that unlike most existing AI products, such level of data collection does not rely on data that users explicitly provided to the AI, which could raise concerns: *"In our applications that we designed [at work], it's mostly around what the user has explicitly provided to the AI, what they offer based on that data as well, as opposed to what we did today in this exercise from the perspective of ToM, the AI assistant picks the data, not like the user offers to it. So this is proactively picking up data."* (D10) Practitioners voiced their concerns about user's agency in consenting to AI's continuous monitoring. D8 raised concerns about the lack of mechanisms for users to stop the monitoring, but also believed that it could support user's goals:

> *"I feel like if Lily is my friend, I could just ask her or if I notice something, I could go back to our chat history. I may not have access to Lily's health records and fitness goals and everything like that [in this AI Trip Planning Assistant scenario]. It's scary because there doesn't seem to be a mechanism to let the AI know when to stop and what kind of access it could have. But I think if the goal is to make Lily happy about this trip schedule, I think the AI might achieve that goal. It's just like the way it's doing it, I don't know."* (D8)

D6 raised concerns about user acceptance of continuous monitoring devices such as the camera that they used in their storyboard: *"Even when you're with your friends, you're not always looking at 'em, always assessing what they're feeling. Your gaze is not always on them. It's like having something's gaze always on you assessing. You can feel unnatural and weird, and I think people would want to opt out of that eventually because no one wants to feel like they're living life on film, especially when they're at work."*

Many practitioners believed that careful understanding and negotiations with the users about their privacy boundaries and perceived trade-offs would be necessary to facilitate continuous monitoring:

> *"People have suggested many things in our user interviews, like, Hey I want AI to tell me to do this if I'm feeling this, or it should tell me which scissor I should use when I'm pruning this kind of tree, etc. But again, once people understand that this is the kind of data that they would have to provide, having cameras in home or maybe reading this data from their objects in home, it will get messy. So the trade off that some people are willing to make, it also depends on their AI and digital literacy, is important to explore."* (D9)

E4 pointed out that it might be necessary to sacrifice our data to gain conveniences from ToM-enabled AI: *"To be honest, I think we have to sacrifice, but it will be at a certain stage. Right now we're pretty much sacrificing our privacy to the Big Tech. They know what we're searching on, what we're doing, what our closest friend is doing on Facebook, etc.. But we're exchanging [our privacy] with the conveniences. We can look up information really quickly connecting with people who are far away from us..."* As E11 put very aptly on the tension between continuous monitoring and user acceptance: *"The AI needs help. And to get that, it needs to see the person and understand them, but at the same time, do the people want to be seen?"*

4.2.3 Adapting to Dynamic Shifts in Mental States. Through our analysis, we found that practitioners proposed various AI solutions that were responsive to user's dynamic mental states in their chosen scenarios, emphasizing the need to design ToM-enabled AI to be adaptive to changes in people's mental states over time. For example, S13 (Family AI Assistant) equipped their ToM-enabled AI with an internal cool-down timer that can check back on the mom after a certain period of time. E13 explained this design: *"Maybe* *the AI could have an internal timer to check on people. Instead of responding right away, it could wait a few minutes and then ask, 'How is everything going?' or offer alternatives. If I'm very mad, I don't want to think about it at that moment, but I might be more open after I've calmed down."*

In the Autonomous Vehicle scenario, both S11 and S4 incorporated AI confirmations with either the pedestrian or their personal AI assistant to verify the pedestrian's intention to cross the road before taking action. While designing their AI solution, S11 considered the possibility that a pedestrian might suddenly change their intent and prepared for this by having the vehicle switch to the outer lane to create distance from the pedestrian and slow down in preparation for an immediate stop. E11 highlighted this possibility:

> *"One more thing that came to mind was even if [the pedestrian is] looking around at the shops and trying to find the right shop and all that stuff, and they weren't initially coming up to the curb to cross the street, what happens if they figure out that they did want to go to this shop that's across the street? And humans being humans don't think about it really and just immediately start taking steps towards the street and crossing the street. So you go from not a threat of just looking at the shops around the street to now you're a threat. You're walking out to the street now you become a pedestrian who wants to cross the street."* (E11)

Several practitioners highlighted the timing of AI actions as an important design factor, believing that ToM-enabled AI should proactively respond to users' dynamic mental states. Practitioners felt that AI responses that are too proactive without any prompting could be perceived as strange and uncomfortable. D9 said, *"An interesting point is, after how long should an AI family assistant in this case even jump in? Probably it's also scary in some ways... Let's say I'm doing a chore and I have it open on my phone and then the AI assistant starts speaking to me or something like that. It would be super weird, right?"* Some practitioners even proposed that ToM-enabled AI should be able to turn itself off based on the user's mental states. Such AI non-actions were presented as design solutions for the Smart Home Assistant scenario, where both S12 and S5 designed the AI to detect the character's need for quiet alone time, and acted on such inferred mental state by turning itself off.

4.3 Designing ToM-enabled AI to Attune for Subjective Individual Mental States

Our analysis showed that practitioners envisioned ToM-enabled AI as being able to accommodate individual's highly subjective mental states that are largely shaped by subjective experiences, rather than offering generalized AI solutions across users. This design consideration was distilled through practitioners' reflections on the need to design for nuanced and subjective data needs, to account for uncertainty when inferring implicit mental states, and to grapple with the technical challenge of balancing scalable generalization with individual subjectivity.

4.3.1 Designing for Nuanced and Subjective Data Needs. Collecting and analyzing individual data is common in designing personalized AI experiences. However, several practitioners pointed out that ToM-enabled AI is unique given the need to cater to individual mental states that are highly nuanced and subjective to personal traits and experiences. As E10 emphasized, *"I think that a correct behavior from a ToM AI is unique to the user is something that's very distinct about these applications… If the user is an introvert versus an extrovert, the input signals will have very different meanings. So yeah, it seems like it would have to kind of learn the person a little bit."* D13 also noted that because people differ in their personal characteristics, it would be difficult to infer mental states from behavioral cues alone:

> *"I think everyone has different nuances. We do have standardized ways to tell if someone is angry or not. One thing is the decibels or body temperature. But everyone can have different decibel ranges in terms of when someone is mad. So one person can really be yelling or shouting when they're angry versus another person tends to be more calm when they're angry. I think the way to tell these mental states is very difficult based on certain words or something like that."* (D13)

This understanding of mental states being subjective was also illustrated by the extent of highly individualized mental states and personal data that practitioners came up with when brainstorming the mental states and AI techniques. For example, practitioners across all sessions brainstormed extremely personal and subjective data needed to infer mental states, such as the character's personality types, personal albums, car accident history, ethnic and cultural data, health conditions, all in addition to other historical personal data that we mentioned in the earlier sections such as biometrics data, daily routine, behavioral patterns, etc.. In the AI Trip Planning Assistant scenario, some practitioners brainstormed mental states around personal preferences such as the character's fear of a particular mode of transportation. S2 brainstormed AI techniques and data to infer a highly individualized mental state—- the character's reluctance to go on the trip due to bad memories tied to the destination. To make such inference, they designed the AI with access to personal data that was associated with the destination, such as personal albums with videos or photos of the destination, private messages with friends on social media about the destination, family and friends' histories associated with the destination.

This raised practitioners' major concerns about user privacy. E11 reflected on their choice of sharing biometrics data for the Autonomous Vehicle to infer pedestrian's intent to cross the road: *"I'll just say right off the bat I would be very, very, very cautious about sharing personal information like that. We're talking about broadcasting personal information like spiked heart rate or what I'm doing on my phone that kinds of stuff that might be useful for the car to determine the person's intentions. But it would also have very serious privacy implications even if it is anonymized. Nothing's ever truly, really anonymized."*

4.3.2 Designing with Uncertainty for Implicit Mental States. Recognizing the subjectivity and implicit nature of individual's mental states, several practitioners voiced their concern that ToM-enabled AI may never be accurate in inferring users' mental states. E7 noted that even humans cannot directly know others' mental states, highlighting the fundamental challenge of designing ToM-enabled AI:

> *"In a sense, you can think of somebody's mental state as being the hidden part of the system that's actually emitting different behaviors and observations. [...] So this is called the hard problem in neuroscience where we can assume that other people have mental states, but the only person that we know for sure even has a mental state is ourselves. So we have to infer mental states based on the behavior of a person, like their speech patterns, their tone of voice, their gaze, how they're walking, their facial expressions. I mean there's just so many things in history, you look back through their messages, things that they've done in their past to try to get… you almost have to have a whole history of somebody's life to be able to try to infer what their mental state is."* (E7)

Practitioners further elaborated that this challenge was compounded by the fact that there could be multiple valid mental states existing at the same time even within each individual. For instance,in the Family AI Assistant scenario, S13 practitioners highlighted that multiple mental states could co-exist when the character said "this is why we can't have nice things" — she could be both upset about losing the china and also mad at herself.

Our analysis showed that most practitioners designed the AI solution to be on the "safer side" to account for AI's inaccurate inferences due to the subjective nature of individual mental states. For instance, AI solutions such as offering suggestions based on inferred personal preferences or seeking confirmation instead of acting directly were common across all the storyboards generated. For instance, in the AI Trip Planning Assistant scenario, instead of granting the AI direct control over booking itineraries, S2, S6 and S8 all designed the AI to suggest new travel dates, alternative plans, or alternative itineraries based on the AI's knowledge of the character's personal schedule and preferences. Similarly, in the Smart Home Assistant scenario, S5, S10, S12 all designed the AI to take subtle background actions tailored to the character's personal preferences such as playing their favorite calming music or ordering the character's favorite food to help the character de-stress. Several practitioners pointed out that these suggestive rather than direct actions from ToM-enabled AI was necessary to avoid putting the AI in an unrecoverable state.

Practitioners reflected on the challenge of designing ToM-enabled AI that could cater to implicit and subjective mental states. Several practitioners pointed out that designing, or even considering, user's implicit mental states was not something they do at their work practices, in which they focused more on addressing specific user problems. D10 said, *"In our applications that we design, the input does not account for any [user] mental state. We haven't accounted much for it, honestly. It's mostly around what the user has explicitly provided to the AI."* E2 further echoed this point and questioned whether incorporating ToM-enabled AI features would help fulfill company priorities in making profits through products: *"You need to sell customers a product. You need to make it really easy for them to envision a future with this product. I feel like it's a bit tricky to make a sales pitch for [ToM-enabled AI] as opposed to something concrete like 'here's a iPhone 17 pro, it takes better pictures, battery lasts longer.' "* Practitioners also expressed their uncertainty about *how* to design AI solutions that could tailor to individual mental states. D13 raised the uncertainty of not knowing how the users were going to perceive of such AI behavior: *"The difficult part is more so of providing that solution. It can probably detect the same thing from that data, but it's just that how we present it to our customers, they may take it well or they might not take it well."*

4.3.3 Balancing Generalizability with Individual Subjectivity. In interviews, many practitioners identified the difficulty of balancing current AI's development paradigm that focuses on scale and generalizability versus ToM-enabled AI's requirement of attuning to different individual's subjective mental states. D7 reflected:

> *"I think just understanding someone's feelings, emotions, wants, desires, is very difficult. I was saying earlier, I'm trying to even think of how the engineers would build a product like this and how they would build them at all. Maybe start with the dataset and kind of add that in? Like in different scenarios you map them out, but it's so complicated. A person can respond in so many ways and have so many infinite amounts of thoughts and feelings, emotions. How do you replicate that?"* (D7)

Several practitioners also believed that designing ToM-enabled AI would require new approaches that emphasize subjectivity instead of objectivity. E6 explained, *"I think we've gotten AI pretty good at doing objective things, but involving it in a subjective context is different. It is difficult, and I think we'd have to approach that differently. We'd have to approach it from a less mathematical standpoint and more of a psychological [point of view]."*

D11 echoed similar points and further reflected on designing for this delicate balance between objectivity and subjectivity via adaptive interfaces:

> *"In some cases what we're doing is trying to collect enough data that we build something that is sort of robust to different levels of preexisting knowledge. That is different than building it to focus on the mental states of the users. It's more like can we make it… almost generically accessible enough that people will not run into major issues, versus how we disambiguate what's going on in their head when they encounter this and then make it fit. Doing that really well probably would require some kind of adaptive interface where the way it presents itself is not traditional software where every person gets the same thing. Maybe we're not actually that far out from that, but it does feel like a step change that's needed if we're going to take ToM really seriously in everyday applications."* (D11)

Some practitioners viewed this challenge as the motivation for designing ToM-enabled AI's capability in learning, adapting, and revising its interpretation of the user over time to better infer and attune to users' subjective needs. When asked about the opportunities of ToM-enabled AI, E13 emphasized on the importance for ToM-enabled AI to learn about the user over time: *"Everyone is different. If the AI can just overall knows more about the user over time by integrating the AI more into their daily lives, just collect more data over time, and be able to customize based on different users, I think that would be helpful."*

5 Discussion

Through co-design sessions with AI practitioners, we surfaced three interrelated design recommendations for ToM-enabled AI in everyday life: it should be *situated* in the social context and norms to interpret how people's mental states are shaped by their immediate surroundings; be responsive to people's *dynamic* and moment-to-moment intentions, emotions, and needs, rather than relying on static or pre-specified assumptions; and be attuned to how individual's *subjective* histories, experiences and nuanced expressions that fundamentally shape how mental states manifest and are communicated. Taken together, these recommendations highlight the need to move beyond inference-based, static mental state representations and one-size-fits-all personalization toward systems that are socially grounded, temporally adaptive, and personally sensitive.

Compared to previous design research that used similar speculative methods to understand user's perspectives for everyday socially intelligent AI [e.g., 17, 52, 67], our work offered an opportunity for practitioners to reflect on their designs against the real-world constraints and limitations, surfacing tensions within each design recommendation. These design tensions reveal a broader misalignment between the envisioned future of ToM-enabled AI and the realities of current AI design and development practices, hinting that traditional inference-based approaches to ToM may be insufficient to meet the situated, dynamic, and subjective demands of everyday AI interactions. Below, we build on these insights to reconsider how ToM should be conceptualized in everyday human-AI interactions, examine the tensions and practical constraints that shape its realization, and outline a design direction that treats ToM as a pervasive capability embedded within AI functionalities to support continuous human-AI interaction loops. We summarized these implications in Fig. 4.

5.1 Beyond Inference: Reframing ToM for Everyday AI

Our design recommendations for ToM-enabled AI to be situated, dynamic, and subjective highlight that users' mental states are not static internal variables, but are shaped by context, informed by subjective histories and experiences, and evolve over time. In our study, despite being provided the classical inference-based interpretation of ToM, practitioners consistently imagined far more situated, dynamic, and subjective features for ToM-enabled AI products and services even within short, four-panel storyboards. This demonstrates that socially meaningful human-AI interaction facilitated by ToM-enabled AI requires richer, more contextually grounded forms of understanding than mental state inference alone can provide. As such, designing ToM-enabled AI for real-world interactions requires moving beyond the traditional modular and inference-based view of ToM [5, 36], which continues to underlie much of today's AI ToM development and evaluation [54] but may not be sufficient to withstand the realities of in-the-wild deployment in everyday user-facing contexts.

This reframing of ToM-enabled AI to be situated, dynamic, and subjective aligns well with the interactionist and enactivist perspectives of social cognition [22, 29, 31], which conceptualize social understanding as emerging through ongoing interaction, embodiment, and context rather than as the recognition of static, internal representations. However, compared to the traditional inference-based ToM approach to social cognition, these perspectives could be challenging to operationalize in current AI systems that rely heavily on data, and lack the innate, sophisticated perceptual and interactive capabilities that these theories assume [22, 29, 31]. Practitioners grappled with this challenge directly in our study: although they envisioned systems sensitive to dynamic contexts, situated subjectivity, and dynamic subjectivity, they also recognized how difficult such visions would be to realize through the data-centric development approach. Challenges such as maintaining extensive context windows, supporting complex data architectures, and addressing privacy and user-acceptance concerns led practitioners, and us, to question whether the dominant data-centric AI development paradigm is sufficient on its own for building situated, dynamic, and subjective ToM-enabled AI. These reflections invite future work to explore participatory sense-making and interactionally grounded approaches to social understanding that can be adapted to the technical and organizational realities of AI development.

At the same time, we are not arguing for abandoning the inference-based approach to AI ToM altogether: their computational tractability remains valuable and well aligned with prevailing AI paradigms. Perhaps there is an opportunity here to explore hybrid approaches in which inference remains a useful computational tool but is carefully calibrated within the broader interactional processes rather than implemented as a stand-alone module. Additionally, different forms of social understanding may call for different mechanisms— just as elephants don't play chess [11], ToM-enabled AI may not *always* need such envisioned *full-fledged* situated presence, constant monitoring, and extremely subjective designs to support effective, socially meaningful interactions with humans. Future work should explore how varying degrees of situated, dynamic, and subjective social understanding in AI can also be achieved through ambiguity, intentionally limited inference, or selective visibility (e.g., AI turning itself off).

5.2 Tensions in Designing ToM-enabled AI on the Ground

While our findings highlight the importance of designing ToM-enabled AI to be situated, dynamic, and subjective, each of these design recommendations also introduces design tensions (as shown in Fig. 4) that practitioners consistently wrestled with during our study. Designing for **situated** ToM-enabled everyday AI would require the systems to draw not only on contextual signals facilitated through IoT and ubiquitous sensors [42, 49, 60, 92], but also on forms of "user-awareness" that capture individual's subtle and subjective cues grounded in those contexts to give mental state inferences valid social meanings. Yet such expanded awareness requires more pervasive and intrusive forms of data collection, placing it in tension with everyday AI's expected non-intrusive presence and alignment with human social norms. The **dynamic** nature of mental states likewise implies that systems need to monitor change and adapt over time to users' evolving mental states. But the extent of ongoing awareness required for such adaptivity raised significant concerns among practitioners regarding user acceptance, consent, autonomy, and the boundaries of proactive AI behavior––concerns that are especially fraught in private settings [16, 47, 49, 52]. This

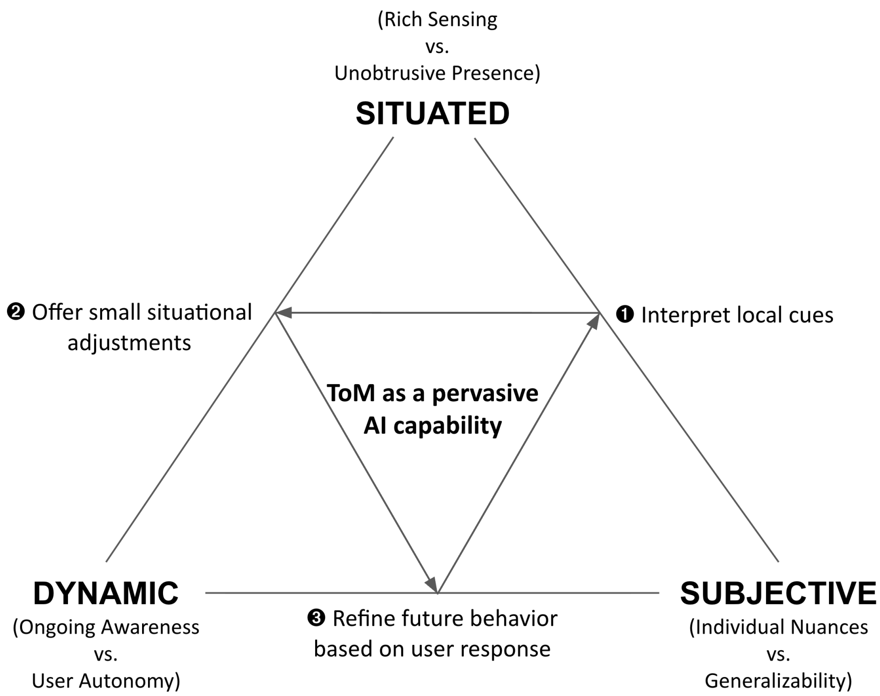

Figure 4: We surfaced three design recommendations for ToM-enabled AI to be situated, dynamic, and subjective, highlighting the need to move beyond inference-based approach to AI ToM in everyday contexts. Each recommendation carries built-in design tensions: "rich sensing vs. unobtrusive presence", "ongoing awareness vs. user autonomy", and "individual nuances vs. generalizability." We propose the design direction of designing ToM as a pervasive AI capability to enhance intended or existing AI functionalities through *continuous interaction loops* to (1) interpret local cues, (2) offer small situational adjustments, (3) refine future behavior based on user response, and repeat.

tension consistently appeared in practitioners' designs in installing intrusive data collection devices such as cameras while adding AI mechanisms that can choose to turn itself off. Finally, designing for **subjectivity** surfaces perhaps the deepest tension: while users and practitioners alike value AI systems that can recognize and adapt to subtle cues and needs that are unique to individual users [51], personalization at this level risks collapsing under the modern AI pursuit of generalization and scale [19, 65].

These tensions further reveal a broader misalignment between the expectations of ToM-enabled AI and the constraints of contemporary AI design and development practices. Practitioners in our study touched upon this misalignment from both the technology-centered and user-centered design perspectives [85, 87]. From a technology-centered design perspective, many questioned the technical feasibility of ToM-enabled AI, pointing to the ethical and legal risks of the extensive data collection required and the limitations of current data architectures and pipelines to process such context-rich, longitudinal, and socially grounded information. Mirroring prior findings on the drawbacks of technology-centered approach [83], these concerns would sharply narrow the design space for ToM-enabled AI in real-world AI design and development practice. From a user-centered design perspective, practitioners noted that their design work largely focuses on solving concrete user problems, with limited attention to the mental states users experience while interacting with AI systems. This solution-driven framing can obscure the value of ToM-enabled AI capability, which aims to enhance the quality and nuance of user experience rather than simply optimizing for task completion. At the organizational level, product teams operate under business priorities that privilege scalability, generalizability, and near-term product value [77, 88], making it difficult to justify investments in socially intelligent AI capabilities that do not map cleanly onto existing performance metrics or short-term return on investment. Together, these observations suggest that ToM-enabled AI does not fit neatly within current data-driven, solution-driven, and profit-driven development paradigms. This raises a broader question of how such capability could be integrated into everyday AI products and services.

5.3 Designing ToM-enabled Everyday AI: From Standalone Feature to Pervasive Capability

These conceptual reframing and tensions motivated us to argue for a different design direction for ToM-enabled AI—–one that treats social understanding not as a stand-alone feature, but as a **pervasive capability that subtly shapes how AI systems perceive and adapt within the boundaries of their intended or existing functionality**. Embedding ToM within the features that already support core user problems allows product teams to enhance user experience without disrupting the profit-driven or solution-driven priorities of contemporary AI development, while also encourage teams to derive richer social cues from signals systems already encounter rather than hunting for new data or install surveilleince mechanism [18]. This design direction also requires structuring the human-AI interaction dynamic as *continuous interaction loops* (as shown in Fig. 4), in which systems interpret local cues, offer small situational adjustments, and refine their behavior over time based on user response. This orientation resonates with emerging interaction frameworks such as Mutual Theory of Mind [76, 78], bi-directional alignment [71], and socially enactive cognitive systems [43], all of which view social understanding as iteratively refined through ongoing interaction. For instance, a smartphone assistant noticing abrupt interaction patterns (e.g., repeated shaking or rapid task switching) might infer momentary frustration, offer a small situational adjustment (e.g., prompting "do you need help"), record and evaluate user's response to such adjustments, and calibrate future behavior based on user feedback. Such contained interaction loops illustrate how situated, dynamic, and subjective forms of ToM can emerge without continuous monitoring or broad data access, aligning with practitioners' vision of systems that learn and adapt to individual user's social nuances over time.

Treating ToM as a capability also raises methodological challenges for current technical, design, and UX research practices, inviting new approaches that move beyond modular and inference-based approaches to social understanding when designing human-AI interaction. Building systems that interpret and adapt to subtle cues in real time requires methods and techniques that help AI recognize locally meaningful signals and evolve alongside individual users, pushing algorithms designed for population-level generalization towards more personalized trajectories. Yet creating such subjective experience also introduces profound output complexity, especially for more general-purpose AI systems such as smart home assistants. This complicates the use of conventional design methods like sketching, prototyping, or Wizard-of-Oz that assume bounded outputs [83, 85], necessitating new design methods and tools that help teams rapidly explore, iterate and evaluate socially nuanced and dynamic AI actions across contexts.

At the same time, user needs, preferences, and privacy boundaries remain central to making ToM-enabled systems viable. These design requirements are highly subjective and contextualized, hence product teams may need lightweight but situated ways of engaging with users, such as contextual inquiry, diary studies, or experience sampling, to understand how different people negotiate the trade-offs of ToM-enabled AI in everyday settings. This also includes clearer, more granular forms of consent that help users understand what kinds of social cues may be inferred and under what conditions. In this view, ToM-enabled AI becomes a set of bounded, purposeful enhancements to everyday interaction, enriching user experience while remaining feasible, respectful, and aligned with real-world development practices. We encourage future work to further examine this design direction in collaboration with AI practitioners, and surface additional challenges and opportunities to realize ToM as a practical capability in everyday products and services.

6 Limitations and Future Work

Our study provides early design recommendations for ToM-enabled AI in everyday contexts, but several limitations shape how our findings and implications should be interpreted. First, our study is grounded in the classic view of ToM that emphasizes on inference-based approach to social understanding, which helped anchor practitioner discussion and designs, but also oriented attention toward certain kinds of social reasoning over others. Future work could explore how alternative perspectives on social understanding, such as embodied, interactional, and participatory accounts, might surface different design considerations. Second, the six human-AI social misalignment scenarios functioned as probes and reflect a narrow slice of everyday life shaped by the U.S.-based everyday AI usage reports and research team experiences. They were useful for eliciting concrete designs but do not capture the full cultural or contextual diversity of everyday human–AI interactions. Expanding scenario settings, sociocultural worlds, and types of ToM-enabled AI behaviors that goes beyond the western, consumer-based, and individualistic scenarios presented in our study would help clarify where our findings and insights do and do not apply. Finally, because our study focused on consumer-facing everyday AI systems, additional work is needed to understand how ToM-enabled AI capability should be designed in other types of everyday AI systems, as well as professional or high-stakes contexts (e.g., ToM-enabled AI in business meetings or healthcare contexts) where expectations, risks, and opportunities may differ.

7 Conclusion

In this paper, we conducted 13 co-design sessions with 26 U.S.-based AI practitioners across engineering- and design-oriented roles to envision and reflect on how ToM-enabled AI might manifest in everyday product and service design. Analysis of design artifacts and session transcripts revealed three interrelated design recommendations for designing future ToM-enabled AI products and services: ToM-enabled AI should (1) be *situated* in the social context to interpret how people's mental states emerge from their environments, (2) be responsive to the *dynamic* and moment-to-moment nature of people's mental states, and (3) be attuned to the *subjective* nature of mental states shaped by individuals' personal traits and histories. Each recommendation revealed embedded design tensions that point to a broader misalignment between practitioners' envisioned future of ToM-enabled AI and the reality of AI design and development practices. These insights underscored the need to move beyond static, inference-based mental state representations and one-size-fits-all personalization. We argue for designing ToM as a pervasive capability that is woven into AI functionalities to support continuous human-AI interaction loops rather than operating as a discrete, stand-alone module.

Acknowledgments

This work was supported by the first author's Carnegie Bosch Postdoctoral Fellowship. We thank Jessica Lin, Shuhao Ma, Vikram Mohanty, our pilot study participants and those who assisted with study recruitment for their feedback and support. We also thank the anonymous reviewers for their valuable and constructive feedback on earlier drafts.

References

[1] AIPRM. 2024. AI Statistics 2024. https://www.aiprm.com/ai-statistics/#most-common-uses-of-ai-exclusive-data.

[2] Arjun R Akula, Keze Wang, Changsong Liu, Sari Saba-Sadiya, Hongjing Lu, Sinisa Todorovic, Joyce Chai, and Song-Chun Zhu. 2022. CX-ToM: Counterfactual explanations with theory-of-mind for enhancing human trust in image recognition models. *Iscience* 25, 1 (2022).

[3] Simon Baron-cohen. 1999. Evolution of a Theory of Mind? In *The Descent of Mind: Psychological Perspectives on Hominid Evolution*. Oxford University Press, 1–31.

[4] Simon Baron-Cohen. 2000. Theory of mind and autism: A review. *International review of research in mental retardation* 23 (2000), 169–184.

[5] Simon Baron-Cohen, Alan M Leslie, and Uta Frith. 1985. Does the autistic child have a "theory of mind"? *Cognition* 21, 1 (1985), 37–46.

[6] Cindy Beaudoin, Élizabel Leblanc, Charlotte Gagner, and Miriam H Beauchamp. 2020. Systematic review and inventory of theory of mind measures for young children. *Frontiers in psychology* 10 (2020), 2905.

[7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.

[8] Virginia Braun and Victoria Clarke. 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and psychotherapy research* 21, 1 (2021), 37–47.

[9] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.

[10] Cynthia Breazeal. 2004. *Designing sociable robots*. MIT press.

[11] Rodney A Brooks. 1990. Elephants don't play chess. *Robotics and autonomous systems* 6, 1-2 (1990), 3–15.

[12] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023).

[13] Hendrik Buschmeier and Stefan Kopp. 2018. Communicative listener feedback in human-agent interaction: Artificial speakers need to be attentive and adaptive. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*. 1213–1221.

[14] Pew Research Center. 2023. *Public Awareness of Artificial Intelligence in Everyday Activities*. Technical Report. Pew Research Center. https://www.pewresearch.org/wp-content/uploads/sites/20/2023/02/PS_2023.02.15_AI-awareness_REPORT.pdf

[15] Luca Cerniglia, Letizia Bartolomeo, Micaela Capobianco, Sara Lucia M Lo Russo, Fabiana Festucci, Renata Tambelli, Walter Adriani, and Silvia Cimino. 2019. Intersections and divergences between empathizing and mentalizing: Development, recent advancements by neuroimaging and the future of animal modeling. *Frontiers in behavioral neuroscience* 13 (2019), 212.

[16] George Chalhoub, Martin J Kraemer, Norbert Nthala, and Ivan Flechais. 2021. "It did not give me an option to decline": A Longitudinal Analysis of the User Experience of Security and Privacy in Smart Home Products. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–16.

[17] Mai Lee Chang, Samantha Reig, Alicia Lee, Anna Huang, Hugo Simão, Nara Han, Neeta M Khanuja, Abdullah Ubed Mohammad Ali, Rebekah Martinez, John Zimmerman, et al. 2025. Unremarkable to Remarkable AI Agent: Exploring Boundaries of Agent Intervention for Adults With and Without Cognitive Impairment. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–26.

[18] Marco Cristani, Ramachandra Raghavendra, Alessio Del Bue, and Vittorio Murino. 2013. Human behavior analysis in video surveillance: A social signal processing perspective. *Neurocomputing* 100 (2013), 86–97.

[19] Tiago Cunha, Carlos Soares, and André CPLF de Carvalho. 2018. Metalearning and Recommender Systems: A literature review and empirical study on the algorithm selection problem for Collaborative Filtering. *Information Sciences* 423 (2018), 128–144.

[20] Scott Davidoff, Min Kyung Lee, Charles Yiu, John Zimmerman, and Anind K Dey. 2006. Principles of smart home control. In *International conference on ubiquitous computing*. Springer, 19–34.

[21] Martin Davies and Tony Stone. 1995. Folk psychology: The theory of mind debate. (1995).

[22] Hanne De Jaegher and Ezequiel Di Paolo. 2007. Participatory sense-making: An enactive approach to social cognition. *Phenomenology and the cognitive sciences* 6, 4 (2007), 485–507.

[23] Manoj Deshpande and Brian Magerko. 2024. Embracing embodied social cognition in AI: moving away from computational theory of mind. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–7.

[24] Sandra Devin and Rachid Alami. 2016. An implemented theory of mind to improve human-robot shared plans execution. *ACM/IEEE International Conference on Human-Robot Interaction* 2016-April (2016), 319–326. doi:10.1109/HRI.2016.7451768

[25] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 278–288.

[26] Jodi Forlizzi. 2018. Moving beyond user-centered design. *interactions* 25, 5 (2018), 22–23.

[27] I-Ning Fu, Kuan-Lin Chen, Meng-Ru Liu, Dai-Rong Jiang, Ching-Lin Hsieh, and Shih-Chieh Lee. 2023. A systematic review of measures of theory of mind for children. *Developmental Review* 67 (2023), 101061.

[28] Yosuke Fukuchi, Masahiko Osawa, Hiroshi Yamakawa, Tatsuji Takahashi, and Michita Imai. 2022. Conveying intention by motions with awareness of information asymmetry. *Frontiers in Robotics and AI* 9 (2022), 783863.

[29] Shaun Gallagher. 2001. The practice of mind. Theory, simulation or primary interaction? *Journal of consciousness studies* 8, 5-6 (2001), 83–108.

[30] Shaun Gallagher. 2004. Understanding interpersonal problems in autism: Interaction theory as an alternative to theory of mind. *Philosophy, Psychiatry, & Psychology* 11, 3 (2004), 199–217.

[31] Shaun Gallagher. 2008. Understanding others: embodied social cognition. In *Handbook of cognitive science*. Elsevier, 437–452.

[32] Shaun Gallagher and Somogy Varga. 2015. Social cognition and psychopathology: a critical overview. *World Psychiatry* 14, 1 (2015), 5–14.

[33] Vittorio Gallese and Alvin Goldman. 1998. Mirror neurons and the simulation theory of mind-reading. *Trends in cognitive sciences* 2, 12 (1998), 493–501.

[34] Gallup, Inc. 2025. Americans Use AI in Everyday Products Without Realizing It. Online. *Gallup News* (2025). https://news.gallup.com/poll/654905/americans-everyday-products-without-realizing.aspx

[35] Richard GT Gipps. 2004. Autism and intersubjectivity: Beyond cognitivism and the theory of mind. *Philosophy, Psychiatry, & Psychology* 11, 3 (2004), 195–198.

[36] Alison Gopnik and Henry M Wellman. 1992. Why the child's theory of mind really is a theory. (1992).

[37] Maaike Harbers, Karel Van Den Bosch, and John Jules Meyer. 2009. Modeling agents with a theory of mind. *Proceedings - 2009 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2009* 2 (2009), 217–224. doi:10.1109/WI-IAT.2009.153

[38] Laura M. Hiatt, Anthony M. Harrison, and J. Gregory Trafton. 2011. Accommodating human variability in human-robot teams through theory of mind. *IJCAI International Joint Conference on Artificial Intelligence* (2011), 2066–2071. doi:10.5591/978-1-57735-516-8/IJCAI11-345

[39] Karen Holtzblatt and Hugh Beyer. 1997. *Contextual design: defining customer-centered systems*. Elsevier.

[40] Jennifer Hu, Felix Sosa, and Tomer Ullman. 2025. Re-evaluating Theory of Mind evaluation in large language models. *Philosophical Transactions B* 380, 1932 (2025), 20230499.

[41] Chien-Ming Huang and Andrea Lockerd Thomaz. 2010. Joint Attention in Human-Robot Interaction.. In *AAAI fall symposium: dialog with robots*.

[42] Razan Jaber, Sabrina Zhong, Sanna Kuoppamäki, Aida Hosseini, Iona Gessinger, Duncan P Brumby, Benjamin R Cowan, and Donald Mcmillan. 2024. Cooking with agents: Designing context-aware voice interaction. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–13.

[43] Sebastian Kahl and Stefan Kopp. 2023. Intertwining the social and the cognitive loops: socially enactive cognition for human-compatible interactive systems. *Philosophical Transactions of the Royal Society B* 378, 1875 (2023), 20210474.

[44] Michal Kosinski. 2023. Theory of mind may have spontaneously emerged in large language models. *arXiv preprint arXiv:2302.02083* 4 (2023), 169.

[45] Minae Kwon, Erdem Biyik, Aditi Talati, Karan Bhasin, Dylan P Losey, and Dorsa Sadigh. 2020. When humans aren't optimal: Robots that collaborate with risk-aware humans. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 43–52.

[46] Min Kyung Lee, Jodi Forlizzi, Sara Kiesler, Paul Rybski, John Antanitis, and Sarun Savetsila. 2012. Personalization in HRI: A longitudinal field experiment. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. 319–326.

[47] Q Vera Liao, Matthew Davis, Werner Geyer, Michael Muller, and N Sadat Shami. 2016. What can you do? Studying social-agent orientation and agent proactive interactions with an agent for employees. In *Proceedings of the 2016 acm conference on designing interactive systems*. 264–275.

[48] Q Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. 2023. Designerly understanding: Information needs for model transparency to support design ideation for AI-powered user experience. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21.

[49] Jieun Lim, Youngji Koh, Auk Kim, and Uichin Lee. 2024. Exploring context-aware mental health self-tracking using multimodal smart speakers in home environments. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.

[50] Shih-Yun Lo, Elaine Schaertl Short, and Andrea L Thomaz. 2020. Planning with partner uncertainty modeling for efficient information revealing in teamwork.

In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. 319–327.
[51] Michal Luria, Judeth Oden Choi, Rachel Gita Karp, John Zimmerman, and Jodi Forlizzi. 2020. Robotic futures: Learning about personally-owned agents through performance. In *Proceedings of the 2020 ACM designing interactive systems conference*. 165–177.
[52] Michal Luria, Rebecca Zheng, Bennett Huffman, Shuangni Huang, John Zimmerman, and Jodi Forlizzi. 2020. Social boundaries for personal agents in the interpersonal space of the home. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–12.
[53] Ziqiao Ma, Jacob Sansom, Run Peng, and Joyce Chai. 2023. Towards A Holistic Landscape of Situated Theory of Mind in Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 1011–1031.
[54] Yuanyuan Mao, Shuang Liu, Qin Ni, Xin Lin, and Liang He. 2024. A review on machine theory of mind. *IEEE Transactions on Computational Social Systems* (2024).
[55] Marco Matarese, Francesco Rea, and Alessandra Sciutti. 2022. Perception is only real when shared: A mathematical model for collaborative shared perception in human-robot interaction. *Frontiers in Robotics and AI* 9 (2022), 733954.
[56] Damian EM Milton. 2012. On the ontological status of autism: The 'double empathy problem'. *Disability & society* 27, 6 (2012), 883–887.
[57] Rachel LC Mitchell and Louise H Phillips. 2015. The overlapping relationship between emotion perception and theory of mind. *Neuropsychologia* 70 (2015), 1–10.
[58] Steven Moore, Q Vera Liao, and Hariharan Subramonyam. 2023. fAIlureNotes: Supporting Designers in Understanding the Limits of AI Models for Computer Vision Tasks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
[59] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 72–78.
[60] Jeesun Oh, Wooseok Kim, Sungbae Kim, Hyeonjeong Im, and Sangsu Lee. 2024. Better to ask than assume: Proactive voice assistants' communication strategies that respect user agency in a smart home environment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–17.
[61] Pegasystems Inc. 2024. What Consumers Really Think About AI: A Global Study. https://www.pega.com/ai-survey.
[62] Michael Plastow. 2012. 'Theory of mind'II: difficulties and critiques. *Australasian Psychiatry* 20, 4 (2012), 291–294.
[63] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences* 1, 4 (1978), 515–526.
[64] David V. Pynadath and Stacy C. Marsella. 2005. PsychSim: Modeling theory of mind with decision-theoretic agents. *IJCAI International Joint Conference on Artificial Intelligence* (2005), 1181–1186.
[65] Omid Rafieian and Hema Yoganarasimhan. 2023. AI and personalization. *Artificial intelligence in marketing* (2023), 77–102.
[66] Hannes Rakoczy. 2022. Foundations of theory of mind and its development in early childhood. *Nature Reviews Psychology* 1, 4 (2022), 223–235.
[67] Samantha Reig, Michal Luria, Elsa Forberger, Isabel Won, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2021. Social robots in service contexts: Exploring the rewards and risks of personalization and re-embodiment. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*. 1390–1402.
[68] Samantha Reig, Michal Luria, Janet Z Wang, Danielle Oltman, Elizabeth Jeanne Carter, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2020. Not some random agent: Multi-person interaction with a personalizing service robot. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 289–297.
[69] Simone G Shamay-Tsoory and Judith Aharon-Peretz. 2007. Dissociable prefrontal networks for cognitive and affective theory of mind: a lesion study. *Neuropsychologia* 45, 13 (2007), 3054–3067.
[70] Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2023. Clever hans or neural theory of mind? stress testing social reasoning in large language models. *arXiv preprint arXiv:2305.14763* (2023).
[71] Hua Shen, Tiffany Knearem, Reshmi Ghosh, Kenan Alkiek, Kundan Krishna, Yachuan Liu, Ziqiao Ma, Savvas Petridis, Yi-Hao Peng, Li Qiwei, et al. 2024. Towards bidirectional human-ai alignment: A systematic review for clarifications, framework, and future directions. *arXiv preprint arXiv:2406.09264* 2406 (2024), 1–56.
[72] Hari Subramonyam, Divy Thakkar, Andrew Ku, Juergen Dieber, and Anoop K Sinha. 2025. Prototyping with prompts: Emerging approaches and challenges in generative ai design for collaborative software teams. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
[73] Tomer Ullman. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv preprint arXiv:2302.08399* (2023).
[74] Sarah Theres Völkel, Daniel Buschek, Malin Eiband, Benjamin R Cowan, and Heinrich Hussmann. 2021. Eliciting and analysing users' envisioned dialogues with perfect voice assistants. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–15.
[75] Alan Richard Wagner. 2013. Developing Robots that Recognize when they are being Trusted. In *2013 AAAI Spring Symposium Series*.
[76] Qiaosi Wang and Ashok K Goel. 2022. Mutual theory of mind for human-AI communication. *arXiv preprint arXiv:2210.03842* (2022).
[77] Qiaosi Wang, Michael Madaio, Shaun Kane, Shivani Kapania, Michael Terry, and Lauren Wilcox. 2023. Designing responsible ai: Adaptations of ux practice to meet responsible ai challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
[78] Qiaosi Wang, Koustuv Saha, Eric Gregori, David Joyner, and Ashok Goel. 2021. Towards mutual theory of mind in human-ai interaction: How language reflects what students perceive about a virtual teaching assistant. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
[79] Qiaosi Wang, Xuhui Zhou, Maarten Sap, Jodi Forlizzi, and Hong Shen. 2025. Rethinking theory of mind benchmarks for llms: Towards a user-centered perspective. *arXiv preprint arXiv:2504.10839* (2025).
[80] Laura M Watrin-Avino, Franziska J Forbes, Martin C Buchwald, Katja Dittrich, Christoph U Correll, Felix Bermpohl, and Katja Bödeker. 2023. Affect Recognition, Theory of Mind, and Empathy in Preschool Children with Externalizing Behavior Problems—A Group Comparison and Developmental Psychological Consideration. *Children* 10, 9 (2023), 1455.
[81] Henry M Wellman. 2018. Theory of mind: The state of the art. *European Journal of Developmental Psychology* 15, 6 (2018), 728–755.
[82] Qian Yang, Nikola Banovic, and John Zimmerman. 2018. Mapping machine learning advances from hci research to reveal starting places for design innovation. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–11.
[83] Qian Yang, Justin Cranshaw, Saleema Amershi, Shamsi T Iqbal, and Jaime Teevan. 2019. Sketching nlp: A case study of exploring the right things to design with language intelligence. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
[84] Qian Yang, Alex Scuito, John Zimmerman, Jodi Forlizzi, and Aaron Steinfeld. 2018. Investigating how experienced UX designers effectively work with machine learning. In *Proceedings of the 2018 designing interactive systems conference*. 585–596.
[85] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
[86] Yuan Yao, Li Huang, Yi He, Zhijun Ma, Xuhai Xu, and Haipeng Mi. 2023. Reviewing and reflecting on smart home research from the human-centered perspective. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–21.
[87] Nur Yildirim. [n. d.]. *Discovering the Right Things to Design with Artificial Intelligence*. Ph. D. Dissertation. Carnegie Mellon University.
[88] Nur Yildirim, Alex Kass, Teresa Tung, Connor Upton, Donnacha Costello, Robert Giusti, Sinem Lacin, Sara Lovic, James M O'Neill, Rudi O'Reilly Meehan, et al. 2022. How experienced designers of enterprise applications engage AI as a design material. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–13.
[89] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supritha Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, et al. 2023. Creating design resources to scaffold the ideation of AI concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2326–2346.
[90] Nur Yildirim, Susanna Zlotnikov, Deniz Sayar, Jeremy M Kahn, Leigh A Bukowski, Sher Shah Amin, Kathryn A Riman, Billie S Davis, John S Minturn, Andrew J King, et al. 2024. Sketching ai concepts with capabilities and examples: ai innovation in the intensive care unit. In *Proceedings of the 2024 CHI conference on human factors in computing systems*. 1–18.
[91] Lance Ying, Katherine M Collins, Lionel Wong, Ilia Sucholutsky, Ryan Liu, Adrian Weller, Tianmin Shu, Thomas L Griffiths, and Joshua B Tenenbaum. 2025. On benchmarking human-like intelligence in machines. *arXiv preprint arXiv:2502.20502* (2025).
[92] Sojeong Yun and Youn-kyung Lim. 2025. What If Smart Homes Could See Our Homes?: Exploring DIY Smart Home Building Experiences with VLM-Based Camera Sensors. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.

A Reflexive Thematic Analysis

Table A1: Mapping between design recommendations, themes, and representative example codes from our reflexive thematic analysis of the session transcripts. Only example codes are included here due to space limitations. While the table displays the themes within separate recommendation groups, this structure is simplified for presentation; in our analysis the themes frequently overlapped and contributed to multiple design recommendations.

Design Recommendation	Themes	Example Codes
Designing ToM-enabled AI that is situated in the social context	ToM-enabled AI collecting data to infer mental states	ToM-enabled AI data collection equipment can introduce unnaturalness
		ToM-enabled AI can collect data through network of existing devices
		ToM-enabled AI needs to collect data from multiple modalities and sources to infer mental states
		Types of data that can be used to make mental state inferences (e.g., "nonverbal cues and subtle gestures," "need to access massive personal data," "physiological data," "compare and contrast historical data and current state," "user's reactions to the AI can indicate their mental states," "contextual and environmental data.")
	ToM-enabled AI acting on mental state inferences to align with social norms	ToM AI should be considerate of everyone's mental states in group settings
		ToM AI should be mindful of what information should it disclose to whom
	Data challenges of designing & developing ToM-enabled AI	Technical difficulty in designing an architecture to support massive streams of data
		Data are expensive and difficult to obtain, especially during the initial stage of a project
Designing ToM-enabled AI to support dynamic mental states	ToM-enabled AI acting on mental state inferences to support longitudinal interactions	ToM-enabled AI can correct/revise/update its interpretation of the user based on user feedback over time, rather than fixed assumptions
		AI should anticipate and respond to sudden changes in mental states
		Timing of ToM-enabled AI responses matters
	Concerns of user losing control over AI's collection and sharing of information	AI sharing information to others
		AI acting on mental state information without user control
	Concerns of user consent and acceptance	Users might not be comfortable/accepting of the passive and invasive forms of data collection (e.g., camera always collecting and analyzing)
		Challenge and importance of getting user consent on data access
		People might not want their mental states to be known to the AI or other people
	Tradeoffs between cost and benefit of designing and developing ToM-enabled AI features	How to balance user privacy and the amount of data ToM-enabled AI requires
		Concerns of using ToM-enabled AI might be bigger than the use cases
		Risk of AI misinterpreting users outweighs the convenience of ToM-enabled AI
Designing ToM-enabled AI to attune for subjective individual mental states	Technical challenge in ensuring the accuracy of ToM inferences	People have very nuanced mental states and behaviors, difficult to generalize
		Physiological data can be up to interpretation when inferring mental states, data can represent multiple valid mental states
		People's behaviors may not always truly represent their mental states.
	ToM-enabled AI doesn't fit neatly into the current tech design and development paradigm	Practitioners are used to design for explicit user problems and commands instead of implicit and nuanced mental states
		Practitioners are skeptical on whether ToM-enabled AI products/features can drive company profit
		Practitioners work within constraints of data availability
	ToM-enabled AI making inferences about human mental states	Techniques to engineer ToM-enabled AI seems technically plausible currently
		ToM-enabled AI should infer mental states through probability distribution, which is different than the current classifications and deterministic models
		We need both generalized solution that can be scaled but also personalized solution to enable ToM-enabled AI
	Opportunities of ToM-enabled AI Applications	To design adaptive AI responses tailored to user's transient mental states, improve UX of AI systems
		ToM-enabled AI can better ensure responses are accurate and adaptive to what the user needs, instead of random guessing
		ToM AI can proactively respond to user's transient mental states like agentic AI

B Affinity Diagramming

We used affinity diagramming to cluster the sticky notes participants generated in activity 1 and 2, and the storyboards. The tables below show the example themes and subthemes, and sample data. Note that we also connected sticky notes across the activities for better analytical context— for example, we connected the AI techniques generated in activity 2 to infer specific mental states in activity 1 during our analysis. We are not able to represent those analytical nuances in these tables due to space constraints.

B.1 Activity 1. Brainstorming Mental States

Example Theme	Example Subtheme	Sample Data (sticky notes)
Emotional mental states	Emotional mental states due to situation	[S13] Alice feels upset because she lost her favorite china
		[S7] he's confused about how to cook the dish.
		[S4] Passengers in the car can be frustrated bc they need to get to work
	Hidden and implicit emotions	[S6] Lily might have past experience with Miami that she hates it. Trauma associated with Miami. AI suggestion is not vibing with her
		[S8] she is not comfortable speaking up about her worries because jack and amy are really excited.
Intentions	About situation	[S4] Pedestrian is waiting for his car (automatic vehicle), just not that one.
		[S3] Peter is testing the AI assistant
	For the future	[S11]A friend of the pedestrian is going to drive by and hand them something– maybe they left the sweater at their friend's house.
		[S13] she wants to warn her son to be careful with nice things at home
Knowledge	Lack of knowledge about self/situation	[S8] she doesnt know anything about the location - she is skeptical about where they are going
		[S7] If its their first date Sarah wont know if Jamie know how to cook or not
	Lack of knowledge about AI action/solution	[S6] Jack and Amy don't know much about AI so they just accept AI's solution.
		[S9] husband doesnt know what the future of the gift would look like (high probability of it being broken by the son)

B.2 Activity 2. AI Techniques Brainstorming

Example Theme	Example Subtheme	Sample Data (sticky notes)
What data can AI collect	Visual cues	[S2] AI can monitor lily's facial expression, not maintaining eye contact, eyes dodging
		[S13] There could be a camera that could detect if something similar has happened before
	Voice and speech cues	[S9] AI can identify Alice's tone and speech, voice data
		[S10] AI can detect unusual loud noises (key clutter, drop noise) compared to normal days
	Physiological/biometric information	[S7] Heart rate and other sign of nervousness
		[S12] AI could also monitor her temperature to gauge her mental state
	Personal data	[S7] AI can track whether Jaime made google searches on Sarah to figure out what she likes.
		[S2] AI has access to your personal albums, Ai recognize videos and photo of that place
How can AI make inferences	Multiple sources	[S11]Ride share apps have designated pick-up stops. Check and see if where the pedestrian is standing is near the pick-up spot. Location of the pedestrian (AV can collect data from other apps)
		[S9] Every object at home is connected in some ways and AI can identify if there is frequent displacement from their existing place - IoT
	AI inferring social patterns	[S2] AI could track lily's history when engaging with these events, she's always shy or doesn't like to speak up.
		[S8] AI can know the relationship dynamics between the group
	Track behavioral patterns longitudinally	[S3] AI can recognize how frequently Peter talks to the AI assistant
		[S1] Previous recordings would be useful. Interesting and difficult part. Instead of judging current state, the AI should infer and gain previous learnings about the person's personality.
	AI self monitoring	[S5]AI just turns itself off for the night. Lina starts walking up to the AI in an angry manner...
		[S4] AI should provide rationales of its abnormal AI behavior.

B.3 Activity 3. Storyboarding

Themes	Example Subthemes	Sample Data (Researcher Notes/Observation)
AI forms	Home decor	[S5] Photo frame on the wall
		[S7, S10, S12] Alexa / Nest /tablet
	Robotic AI assistant	[S13, S1, S4]
AI data collection	Continuous data collection	[S13] collect historical data
		[S2] collect digital footprint - browsing history content
	Multimodal data collection	[S8] calendar + personal communications + personality + physiological data
		[S4] facial expression + posture + gaze direction -> from others phones / services
		[S5] Camera + sensors - images/ visual information, detect things on the flooe
	Connected devices	[S4]data type - from other AI agents
		[S9]data type - Sensors for interconnected devices
AI inferring mental states	Comparing historical data against current	[S13] comparing historical behaviour with current behaviour to get mental states
		[S1]check delta as compared to previous behaviours
	Most probable cause of mental state	[S8] highest most likely issue - probability - output as JSON file
		[S4]most probable action that AV thinks is applicable
AI solutions	Confirming before acting	[S3] confirm inclination and action
		[S11]confirming - ask pedestrian for clarification
		[S1] ask what it can do to cheer up
	Offering suggestions	[S13] Suggestive - giving son advice and emotional supprt
		[S2] suggestive - proposes new date after looking at calendars of everyone in group
		[S10] suggestive - offers lina suggestions for her favourite food and tv show
	Subtle background actions	[S5] corrects action - stops coffee machine and energetic music
		[S12]AI stops interactions with lina
		[S5] goes to sleep to avoid intrusion and offer space
	Invasive actions	[S1] play fav music, tell a joke, dance
		[S11] change lanes