The Next Level of Long-Term Agent Autonomy — Proactively Acquiring Knowledge and Abilities

Blue Sky Ideas Track

Hermine J. Grosinger
Örebro University
Örebro, Sweden
hermine.grosinger@oru.se

ABSTRACT

For an artificial agent operating long-term under real-world conditions it is not enough to be able to act on orders given by the human. Even being able to act proactively (anticipatory, self-initiated) does not suffice. The reason roots in the unrealistic assumption that the proactive agent from the start and always knows everything it needs to know and has all the abilities it requires. We argue that the agent has to be able to *proactively learn* new knowledge and abilities according to how the dynamic environment evolves. We identify challenges and directions towards proactive learning. Our focus is on formal methods which lend themselves to doing the necessary reasoning but also give suggestions how these might be integrated with machine learning. The ideas envisioned in this paper can advance (M)AS (Multi-Agent Systems) research and have the potential to enhance collaborations of hybrid human-AI systems.

KEYWORDS

Proactivity; Autonomy; Human-AI Systems; Long-term Real-world Agents; Epistemic Reasoning; Ability; Learning

ACM Reference Format:
Hermine J. Grosinger. 2025. The Next Level of Long-Term Agent Autonomy — Proactively Acquiring Knowledge and Abilities: Blue Sky Ideas Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 6 pages.

1 INTRODUCTION

Research on agent-based systems rests on a solid body of findings [49]. This paper is on agents which are *proactive* in *long-term, real-world* settings. Although being addressed [11, 23, 26, 44], this is an unsolved problem. Narrowing down our focus, we want to show the need for agents on an even higher level of autonomy and call for agents that not only *proactively act* but *proactively learn* new knowledge and new abilities, that is, the agents can self-initiated decide which new knowledge or ability to learn and when. First we define the term *proactivity* as we use it here (similar to Weiss [49]'s *proativeness*):

Proactivity is a comprehensive (human-like), self-initiated and anticipatory behavior as opposed to a simple reactive behavior.

Example 1. *A domestic robot ProRo self-initiated brings Bob his medicine so he can take it as ProRo foresees that otherwise Bob would not have taken it and fallen ill.*

This is difficult because ProRo needs to use joint cognitive capabilities to produce proactive behavior: ProRo needs to understand the current context (Bob does not take his medicine); it needs to anticipate the future (Bob will be ill because he is not taking his medicine); it needs to mentally simulate acting alternatives and compute their effects (Bob can take his medicine after ProRo brought it); it needs to understand what is preferable (Bob not being ill is preferable); it might also need to use further capabilities such as epistemic reasoning (Bob does not know he needs to take medicine); etc. AI researchers [5, 22, 26, 32] have adopted the definition of human proactivity established in organizational psychology, as being *anticipatory* and *self-initiated* action to impact people and their environment [21]. It has been found that proactive AI systems are preferred [5], more easily accepted [38], and trusted [31] by humans. Several recent works address this question. Some of them use epistemic reasoning to initiate action based on the human's false belief, based on modal logic and an epistemic planner [36, 44]; or based on HTN planning for belief alignment in a human-robot collaborative task [17, 18, 43]. While epistemic reasoning can be useful for proactive acting, it is indispensable for *proactive learning*, which we here introduce as a new concept [1] and an extension of proactive acting . We propose the following definition:

Proactive Learning is a comprehensive (human-like), self-initiated and anticipatory behavior of deciding when to learn what new knowledge or new ability as opposed to a simple reactive learning behavior.

Hence, in contrast to classical Machine Learning (ML), where machines learn an objective given by the human, we advocate for an agent that itself decides its own learning objectives.

Example 2. *Ann's robot proactively* acquires *the* new ability *to pick up the lunch from a delivery service and bring it to Ann. But before it needs to* learn a new piece of knowledge*, the door code, to be able to re-enter the house afterwards. Bob's robot proactively* learns the new ability *of opening drawers, so it can* acquire the knowledge *where the medicine for Bob is.*

[1] In contrast, in machine learning "proactive learning" (an extension of "active learning") is about obtaining class labels for unlabeled data or ranking preferences [14].

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Note that Example 1 highlights many cognitive abilities (context awareness, anticipation, mental simulation, reasoning on preferences), which are useful or indeed necessary for agents to be proactive. However, here we will focus on discussing those cognitive abilities that are most crucial for proactive learning, highlighted in Example 2 (epistemic reasoning, reasoning on abilities, learning). While unsolved, long-term proactive acting is being addressed today, whereas poactive learning is a completely unexplored field. It is required to develop new theories and new computational methods which enable an agent to reason on what knowledge or ability it lacks, why it should acquire the new knowledge or new ability, who knows the knowledge it needs, etc. Work on such highly autonomous agents can advance the field of (M)AS, more precisely, (M)AS with proactive agents acting long-term in dynamic environments. It may also facilitate collaboration in human-AI teams. We state the following challenges that are necessary to be addressed:

I. Develop a theory and computational methods to proactively decide when to acquire what new knowledge.

II. Develop a theory and computational methods to proactively decide when to acquire what new ability.

III. Integrate results of **I**. and **II**. into one unified theory and combined computational methods to proactively decide when to acquire what new knowledge and when to acquire what new ability.

2 PROACTIVE LEARNING

Have we not already "solved" proactive agent behavior? The short answer: it depends on how you define proactivity (and how you define "solved"). One might say that the early work on Belief Desire Intention architectures for modeling rational agents [40] (which in turn is based on the considerations about intention and intentional action in Bratman [10]) provides means to model proactive behavior. One might consider Wooldridge and Parsons [51]'s work presenting BDI-agents that are able to choose whether to deliberate (reconsider) or act and higher-level control functions that can choose which such strategies to apply, as proactive behavior. However, we have a more general, extrinsic view. Instead of considering an agent's intrinsic drives, our understanding of proactivity is based on the recent works of Grosinger [22] and Lorini [33] who speak of behavior that is *autonomously initiating action taking into account future state development as well as anticipating potential consequences of the agent's actions on (the mind) of other agents and the environment.* Proactive behavior, following the definition above, requires (at least) reasoning about the current context, making predictions about future state development, doing mental simulations of applying actions, doing preference and epistemic reasoning. This is the understanding of proactivity we propose to extend.

Why proactive learning? Again the short answer first: because we live in a dynamic world (that is, the real world). A proactive agent operating long-term in such a world makes autonomous acting decisions by reasoning on what she currently knows and what she currently is able to do, among other things. However, it is an unrealistic assumption that an agent knows all she needs to know independent of how the world will evolve. It is unrealistic to assume that she observes and knows all changes in the world. Concerning the agent's abilities, we cannot realistically assume that the agent has knowledge of all actions' effects and knows which strategy she should form for some goal. In long-term dynamic settings, an agent will have to adapt. It might be necessary to gain new knowledge or update existing knowledge in order to make informed (acting) decisions. It might be necessary to not only reason about what the agent can already do, but also about what she can learn to do and its potential effects, in order to behave intelligently. It might be necessary that the agent first acquires new knowledge before she can learn a new ability, or vice versa, it might be that the agent first needs to learn a new ability in order to retrieve a new knowledge. This and further intertwined learning of new knowledge and abilities is addressed by proactive learning.

Why logic? In the following subsections we analyze how to approach the problem of proactively learning new knowledge and new abilities. We do this giving several directions, with different combined methodologies, but focus on logic based aspects. We argue that it is favorable to model the cognitive abilities needed with expressive formal languages (see also Lorini [33]). The reasoning necessary to create proactive learning behavior can conveniently be based on these formal models. We target learning new knowledge that is not only data but has meaning and can be attributed to an agent. The agent should be able to reason on some other agent's knowledge and make inferences from that; the agent needs to be able to reason on her own knowledge and must be able to infer what knowledge she might be missing and which new knowledge she needs to learn. Similarly, the proactive learning of new abilities requires it to be possible for the agent to understand what abilities she already has and which she needs to acquire, when, and why. Formal methods including formal logic are well suited to do this deliberation on what to learn and when, which then may use machine learning to, for example, learn a certain motor skill, in a neuro-symbolic solution.

3 CHALLENGES AND DIRECTIONS FOR I. – III.

3.1 Proactive Knowledge Learning (I.)

To proactively decide to learn a new piece of knowledge, the agent needs to understand *introspection*, This property of knowledge is expressed in the *Introspection Axioms*: $\models K_i\phi \Rightarrow K_iK_i\phi$, says that, if agent i knows ϕ then, agent i knows that she knows ϕ (positive introspection); and $\models \neg K_i\phi \Rightarrow K_i\neg K_i\phi$, says that, if agent i does not know ϕ, then agent i knows that she does not know ϕ (negative introspection) [16]. These axioms enable the agent to understand (i) that she is lacking knowledge; and (ii) which knowledge she is lacking. The capability allowing to make such cognitive inferences is called *Epistemic Reasoning*. It enables the agent to reason about her own and other agents' (including humans') knowledge and beliefs. It is an important capability for collaborating agents, hence, also in hybrid human-AI agent teams. Reasoning in general about own and other agents' mental states is referred to as *Theory of Mind (ToM)* in psychology [7] and in AI [46]. Work in AI exists investigating proactive decisions if, what and when to inform the human about their false beliefs [44]. Beyond that, we here propose to enable the agent to proactively decide if, what and when to learn a new knowledge ϕ from observation or from another agent who knows ϕ. There are different ways to model knowledge. A vast body of work in Epistemic Logic (EL) [2] builds on the possible worlds semantics which is based mathematically on Kripke models [16, 27]. These are used for modeling that the agent believes what she considers possible in all accessible worlds. Alternatives are belief sets (knowledge bases) and Bayesian models (see, for example, Gärdenfors and Makinson [20]). Epistemic states in belief sets are modeled as sets of sentences based on a language $\mathcal{L}$ that the agent believes to be true. Bayesian models (see also Baker et al. [3]), model an agent's epistemic state by a probability function over possible worlds representing degrees of belief. We need to be able to model *change* of knowledge states to achieve a proactively learning agent. In belief sets this is done by the operations *expansion* (accept new information ϕ), *contraction* (reject existing information ϕ), and *revision* (accept ϕ if not existing, reject ϕ if existing). For Bayesian models, update of belief states is done by Bayesian inference. Furthermore, the dynamic version of EL, *Dynamic Epistemic Logic (DEL)* [47] can be used. It allows to update both the ontic and the epistemic aspects of the state (*epistemic model* $\mathcal{M}$) with an action (*event model* $\mathcal{E}$) through a *product update* to result in a new epistemic model $\mathcal{M}'$, denoted $\mathcal{M} \otimes \mathcal{E} = \mathcal{M}'$.

Epistemic reasoning can be combined with reasoning on causality to find out what new knowledge to learn. Consider the following variant of Example 1.

Example 3. *ProRo knows the following causal chain concerning Bob: Not knows medicine $\Longrightarrow$ Not takes medicine $\Longrightarrow$ Sick. Based on this intertwined causal and epistemic reasoning, ProRo proactively decides to learn the new knowledge of the location of the medicine and bring it to Bob so he can take it.*

The combination of causal and epistemic reasoning is under-investigated, but may inform proactive learning. Barbero et al. [6] stress the importance of causal reasoning research in AI. Their approach is based on a standard causal model [24] with exogenous (causally independent) and endogenous (causally dependent) variables, as well as a structural function which describes the relation between variables. They make this model epistemic by allowing several (instead of only one) possible valuations of variables. An alternative approach is proposed by Ding et al. [13], which is a fusion of causal models and Kripke models, thus using possible world semantics, which they call *Causal Kripke Models*. Pearl and Mackenzie [39] propose *(Structural) Causal Models, SCMs*, which can combine data and causal knowledge. Madumal et al. [34] propose action influence models, which are SCMs that are learned by reinforcement learning.

After discussing epistemic and causal reasoning for proactive learning, we have come to the question of how to do learning itself. Baltag et al. [4] propose *Dynamic Logic for Learning Theory (DLLT)*. The agent can use their introduced modality $[o]\phi$ ('after the evidence o is observed, ϕ will hold') and the learning operator $L(\vec{o})$ to map every evidence (sequence of observations $\vec{o}$) to a conjecture (the agent's strongest belief after observing $\vec{o}$). Works like Charrier et al. [12] can serve as a basis for learning from announcements, denoted $!\phi$. One may consider to employ Large Language Models (LLMs) for doing ToM tasks. Kosinski [30] present a study of GPT4[3] which successfully passes 95% of widely used false-belief tests[4].However, our position is that we need other methods (possibly combined with data-based methods such as the above) to approach the kind of proactive learning we suggest in this paper. We envision agents to be able to reason about other agents' beliefs, introspect their own knowledge and understand which knowledge they are lacking, track their own and other agents' beliefs changing over time and use all this information for their proactive learning decisions — a kind of reasoning for which formal methods are suitable.

3.2 Proactive Ability Learning (II.)

By *ability* we mean abstract skill. By learning abilities we mean acquiring in a systematic way new high-level behaviors, that the agent can reason on in an abstract way. High-level behaviors can be comprised of low-level actions, which might be learned by classical reinforcement learning in order to, for example, gain the motor skills to play table tennis [37]. Having an ability or strategy can be understood as knowing how to achieve a goal through applying actions. The early work on *Propositional Dynamic Logic (PDL)* [19] can model that $[a]p$, which says that "when program a terminates, assertion p holds". To reason on which ability the agent is lacking (does not *know* how to reach the goal), the agent also needs to do epistemic reasoning. Alternating-time Temporal Epistemic Logic (ATEL) [45] (based on ATL [1]) can model epistemic notions and ability notions together. Using the modality $\langle\langle\rangle\rangle$, which means "bringing about", ATEL can, for example, express $\mathrm{K}_a\langle\langle a\rangle\rangle\Diamond\phi$, which means agent a knows that she can ensure/achieve that eventually ϕ holds. Constructive Strategic Logic (CSL) [28] goes beyond AT(E)L. It differentiates between knowing *that* (operator K) and knowing *how* (operator $\mathbb{K}$), meaning, knowing which strategy is available to bring about ϕ. For example, $K_a\mathbb{K}_b\langle\langle b\rangle\rangle\mathsf{win}$, says that agent a knows *that* agent b knows *how* (can identify a strategy) to win. Wang [48] introduces a logic able to express "knowing how to achieve the goal ϕ, given ψ", which can also be understood as "having the *ability* to achieve goal ϕ, given ψ", formally, $\mathcal{K}h(\psi, \phi)$. Areces et al. [2] use the knowing-how operator to express that the agent knows how to achieve ϕ given ψ if there is a *plan* that when executed in any state where ψ holds will end in states satisfying ϕ. Hammond et al. [25] present a game-theoretic approach for reasoning about causal relationships and estimating effects of strategic behavior. Motamed et al. [35] propose a probabilistic temporal logic which can express $\Diamond^n_{\bowtie r}\Phi$, "the agent can act in the next n steps such that Φ will hold with probability $\bowtie r$", where $\bowtie \in \{<, =, >\}$.

It is conceivable that formal logic approaches for reasoning on abilities as the ones above can be combined in a neuro-symbolic approach to enable proactive learning. First the formal, high-level reasoner proactively infers *(i.) if* to learn a new ability, *(ii.) what* new ability should be learned and *(iii.) when.* After a high level decision is made to learn a new ability, machine learning such as reinforcement learning can be employed to learn *(iv.) how* to do the new ability. Qualitative planning is another potential way to go. Based on logic it allows the agent to formulate and reason explicitly

[2]Epistemic logic is concerned with reasoning about knowledge. Doxastic logic is concerned with reasoning about belief. As is common, we here use the term epistemic logic to denote both, reasoning about knowledge and belief.

[3]GPT4 - Generative Pre-trained Transformer 4, https://openai.com/product/gpt-4
[4]False belief tasks test an agent's ability to understand another agent's beliefs, whether they are false and how this other agent will react on the basis of these beliefs [50].

on knowledge about action effects. The agent can learn by streams of observations of action effects from actions executed by herself or by other agents, as well as by history [8, 9] or announcement [12]. Epistemic planning (EP) finds plans considering both ontic and epistemic states and ontic and epistemic actions. EP might be used to find generic epistemic policies of whose actions some might have to be learned by the agent to achieve her goal. Again, in a neuro-symbolic approach, qualitative learning can be used to reason on which actions and their effects the agent knows and which she needs to know for conducting an epistemic policy; machine learning might be used to learn how to execute actions in the policy.

3.3 Unified Theory of Proactive Learning (III.)

Learning new knowledge (**I.**) and learning new abilities (**II.**) can be combined in a unified theory.

LEARN NEW KNOWLEDGE $\longrightarrow$ LEARN NEW ABILITIES. The agent proactively decides to acquire a new ability σ. To learn σ, the agent not only needs to learn action models' effects, but she also needs to learn new knowledge ϕ. Hence, the agent learns knowledge ϕ that is necessary for acquiring ability σ. She learns ϕ, for example, from observation streams $\vec{o}$ and from announcements $!\phi$. When factual and action knowledge are retrieved, the agent can finally use the previously computed epistemic policy which represents ability σ (all actions are now known to the agent). Example: PRORO learns the ability to bring the lunch box to Ann. For this, PRORO first needs to acquire the knowledge of what the door code is.

LEARN NEW KNOWLEDGE $\longleftarrow$ LEARN NEW ABILITIES. The agent proactively decides to acquire a new piece of knowledge ϕ. It can be retrieved by an announcement $!\phi$ or by observation stream $\vec{o}$ that shows ϕ. But in order to acquire ϕ, the agent needs to learn the ability σ. EP can find an epistemic policy for the epistemic goal ϕ; ability reasoning and qualitative learning may find which of the action models of the policy the agent does not know and may need to acquire, for example, from observation streams $\vec{o}$ or from announcements $!\alpha$ telling the effects of action models. Example: PRORO knows that Bob's medicine is in either drawer A, B or C. PRORO needs to learn the ability to open drawers to gain the new knowledge in which of the drawers Bob's medicine is.

Chains of intermixed learning of new knowledge and learning of new abilities are conceivable (learn new knowledge to learn new abilities so that yet another new knowledge can be learned, etc.).

Note on *scalability*: The methods suggested here are expressive but costly. One proposal to address this is by Shvo et al. [44] who use a KD45 logic and the epistemic planner RP-MEP [36] for proactive robot assistance, reasoning about human belief. RP-MEP can reason over complex action theories in large state spaces by limiting belief nesting depth and permitted formulas, which increases efficiency. Works like these are promising proposals for scalability in the field.

3.4 Ethical Challenges in Proactive Learning

An agent that proactively can decide what new knowledge or ability to learn has very high autonomy. This can be a significant benefit as it can facilitate hybrid human-agent teaming in long-term real world settings and makes collaboration with artificial agents more natural for humans [29]. However, it is not obvious how the benefit of such agents should be measured. How should we compare the behavior of agent 1 that chooses to learn ability σ_1 and knowledge ϕ_1 in a certain situation with the behavior of another agent 2 which instead decides to learn ability σ_2 and knowledge ϕ_2. Investigating this question and developing benchmarks for proactively acting and proactively learning agents is a topic of future research.

High agent autonomy is not only beneficial. It might pose ethical challenges such as *Value Alignment (VA)* [41]. VA requires autonomous agents to align their goals and behaviors with human values[5]. When values are not aligned, the agent "single-minded" optimizes getting maximal rewards, but is not doing what humans *actually* want it to do ("wire-heading"). For example, PRORO with misaligned values hides the medicine from Bob so he cannot misplace it. While achieving the positive effect that Bob never misses to take his medicine (PRORO always brings it to him), the behavior is hurting Bob's human dignity. Sanneman and Shah [42] suggest *Transparent VA*, using human feedback about the value learning process to verify alignment or identify and amend gaps in agent models. Russell [41] proposes endowing agents with uncertainty about humans' preferences. Proactively learning agents might take into account this uncertainty, so they are "open" to collect additional information reassuring values are aligned with the human.

Trustworthiness: The EU High-level Expert Group on Artificial Intelligence drafts ethics guidelines for trustworthy AI [15], stating that human autonomy is not to be undermined and humans should have the ability to over-ride decisions made by an AI system. *Transparency* is particularly important for highly autonomous systems as suggested here. Knowledge-based methods are conducive to transparency and trustworthiness through human-readable white-box models. Certain measures should be taken for trustworthiness and transparency in neuro-symbolic approaches, where, for instance, knowledge-based methods compute which new ability should be learned; reinforcement learning is used for learning to do the ability.

4 CONCLUSION

In the field of long-term real world agent autonomy, a large part of the research has addressed the question how to act, given a goal by the human. Another part of the field investigates what the agent should do and when proactively, meaning, self-initiated and anticipatory. We argue for the need of yet higher autonomy, which is beyond *proactively acting* agents. We call for *proactively learning* agents. We propose a new research direction to investigate how the AI agent proactively can decide what new knowledge and what new ability to learn and when. This comprises enabling the agent to reason about her own and others' knowledge, to reason about which ability she lacks and which ability she needs, to reason about which knowledge to learn for learning a new ability, and which ability to learn for learning a new knowledge. We point out challenges of this new research area and suggest directions to approach them, focusing on knowledge/logic-based methods. The significance of this new research direction roots in that humans expect proactivity from their collaborators, and it can therefore contribute to facilitate collaboration in hybrid human-AI teams.

ACKNOWLEDGMENTS

This research is funded by the Swedish Research Council, 2021-05542.

[5]https://futureoflife.org/ai/align-artificial-intelligence-with-human-values/

REFERENCES

[1] Rajeev Alur, Thomas A Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *ACM* 49, 5 (2002), 672–713. https://doi.org/10.1145/585265.585270

[2] Carlos Areces, Raul Fervari, Andrés R. Saravia, and Fernando R. Velázquez-Quesada. 2023. Uncertainty-Based Knowing How Logic. arXiv:2304.01022 [cs.LO]

[3] Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. 2011. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proc Annual meeting of the cognitive science society*, Vol. 33. eScholarship University of California, California, USA, 2469–2474.

[4] Alexandru Baltag, Nina Gierasimczuk, Aybüke Özgün, Ana L Vargas Sandoval, and Sonja Smets. 2019. A dynamic logic for learning theory. *Journal of Logical and Algebraic Methods in Programming* 109 (2019), 100485. https://doi.org/10.1016/j.jlamp.2019.100485

[5] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh PN Rao, and Minoru Asada. 2017. Efficient human-robot collaboration: when should a robot take initiative? *International Journal of Robotics Research* 36, 5-7 (2017), 563–579.

[6] Fausto Barbero, Katrin Schulz, Sonja Smets, Fernando R. Velázquez-Quesada, and Kaibo Xie. 2020. Thinking About Causation: A Causal Language with Epistemic Operators. In *Dynamic Logic. New Trends and Applications*, Manuel A. Martins and Igor Sedlár (Eds.). Springer International Publishing, Cham, 17–32.

[7] 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. https://doi.org/10.1016/0010-0277(85)90022-8

[8] Thomas Bolander and Nina Gierasimczuk. 2017. Learning to act: qualitative learning of deterministic action models. *Journal of Logic and Computation* 28, 2 (2017), 337–365. https://doi.org/10.1093/logcom/exx036 arXiv:https://academic.oup.com/logcom/article-pdf/28/2/337/24261856/exx036.pdf

[9] Thomas Bolander, Nina Gierasimczuk, and Andrés Occhipinti Liberman. 2021. Learning to Act and Observe in Partially Observable Domains. https://doi.org/10.48550/ARXIV.2109.06076

[10] Michael E Bratman. 1987. *Intention, plans, and practical reason*. Harvard University Press, Cambridge, MA. http://books.google.de/books?id=I0nuAAAAMAAJ

[11] Sera Buyukgoz, Jasmin Grosinger, Mohamed Chetouani, and Alessandro Saffiotti. 2022. Two ways to make your robot proactive: Reasoning about human intentions or reasoning about possible futures. *Frontiers in Robotics and AI* 9 (2022). https://doi.org/10.3389/frobt.2022.929267

[12] Tristan Charrier, Andreas Herzig, Emiliano Lorini, Faustine Maffre, and François Schwarzentruber. 2016. Building epistemic logic from observations and public announcements. In *15th International Conference on Principles of Knowledge Representation and Reasoning (KR 2016), co-located with DL 2016 and NMR 2016*. AAAI Press (Association for the Advancement of Artificial Intelligence), Le Cap, South Africa, 268–277. https://hal.science/hal-04109521

[13] Yiwen Ding, Krishna Manoorkar, Apostolos Tzimoulis, Ruoding Wang, and Xiaolong Wang. 2023. Causal Kripke Models. *Electronic Proceedings in Theoretical Computer Science EPTC, Theoretical Aspects of Rationality and Knowledge TARK* 379 (July 2023), 185–200. https://doi.org/10.4204/eptcs.379.16

[14] Pinar Donmez and Jaime G Carbonell. 2008. Proactive learning: cost-sensitive active learning with multiple imperfect oracles. In *Proceedings of the 17th ACM conference on Information and knowledge management*. 619–628.

[15] EU. 2019. High-Level Expert Group on Artificial Intelligence: Draft Ethics Guidelines for Trustworthy AI. https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai.

[16] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. 2004. *Reasoning About Knowledge*. MIT Press, London, England. https://books.google.se/books?id=xHmlRamoszMC

[17] Anthony Favier, Shashank Shekhar, and Rachid Alami. 2022. Robust Planning for Human-Robot Joint Tasks with Explicit Reasoning on Human Mental State. https://hal.laas.fr/hal-03909319

[18] Anthony Favier, Shashank Shekhar, and Rachid Alami. 2023. Anticipating False Beliefs and Planning Pertinent Reactions in Human-Aware Task Planning with Models of Theory of Mind. https://hal.science/hal-04163435

[19] Michael J Fischer and Richard E Ladner. 1979. Propositional dynamic logic of regular programs. *J. Comput. System Sci.* 18, 2 (1979), 194–211. https://doi.org/10.1016/0022-0000(79)90046-1

[20] P. Gärdenfors and D. Makinson. 2008. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. College Publications. https://books.google.se/books?id=zO17PAAACAAJ

[21] Adam M Grant and Susan J Ashford. 2008. The dynamics of proactivity at work. *Research in Organizational Behavior* 28 (2008), 3–34.

[22] Jasmin Grosinger. 2022. On proactive human-AI systems. In *AIC Workshop on AI and Cognition*. AIC Workshop, Örebro, Sweden.

[23] Jasmin Grosinger, Federico Pecora, and Alessandro Saffiotti. 2019. Robots that maintain equilibrium: Proactivity by reasoning about user intentions and preferences. *Pattern Recognition Letters* 118 (2019), 85–93. https://doi.org/10.1016/j.patrec.2018.05.014 Cooperative and Social Robots: Understanding Human Activities and Intentions.

[24] J.Y. Halpern. 2016. *Actual Causality*. MIT Press. https://books.google.se/books?id=9cvxDAAAQBAJ

[25] Lewis Hammond, James Fox, Tom Everitt, Ryan Carey, Alessandro Abate, and Michael Wooldridge. 2023. Reasoning about causality in games. *Artificial Intelligence* 320 (2023), 103919. https://doi.org/10.1016/j.artint.2023.103919

[26] Helen Harman and Pieter Simoens. 2020. Action graphs for proactive robot assistance in smart environments. *Journal of Ambient Intelligence and Smart Environments* 12, 2 (2020), 1–21.

[27] Kaarlo Jaakko Juhani Hintikka. 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, NY, USA.

[28] Wojciech Jamroga and Thomas Ågotnes. 2007. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics* 17, 4 (2007), 423–475.

[29] Ujjwal KC and Jacques Chodorowski. 2019. A case study of adding proactivity in indoor social robots using belief–desire–intention BDI model. *Biomimetics* 4, 4 (2019), 74.

[30] Michal Kosinski. 2023. Theory of Mind May Have Spontaneously Emerged in Large Language Models. arXiv:2302.02083 [cs.CL]

[31] Matthias Kraus, Marvin Schiller, Gregor Behnke, Pascal Bercher, Michael Dorna, Michael Dambier, Birte Glimm, Susanne Biundo, and Wolfgang Minker. 2020. "Was That Successful?" On Integrating Proactive Meta-Dialogue in a DIY-Assistant Using Multimodal Cues. In *Proceedings of the 2020 International Conference on Multimodal Interaction* (Virtual Event, Netherlands) *(ICMI '20)*. Association for Computing Machinery, New York, NY, USA, 585–594. https://doi.org/10.1145/3382507.3418818

[32] Tingting Liu, Erli Lyu, Jiaole Wang, and Max Q.-H. Meng. 2022. Unified Intention Inference and Learning for Human–Robot Cooperative Assembly. *IEEE Transactions on Automation Science and Engineering* 19, 3 (2022), 2256–2266. https://doi.org/10.1109/TASE.2021.3077255

[33] Emiliano Lorini. 2024. Designing Artificial Reasoners for Communication. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 2690–2695.

[34] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2020. Explainable Reinforcement Learning through a Causal Lens. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 03 (Apr. 2020), 2493–2500. https://doi.org/10.1609/aaai.v34i03.5631

[35] Nima Motamed, Natasha Alechina, Mehdi Dastani, Dragan Doder, and Brian Logan. 2023. Probabilistic Temporal Logic for Reasoning about Bounded Policies. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, Edith Elkind (Ed.). IJCAI, Macao, China, 3296–3303. https://doi.org/10.24963/ijcai.2023/367 IJCAI 2023 : 32nd International Joint Conference on Artificial Intelligence ; Conference date: 19-06-2023 Through 25-06-2023.

[36] Christian Muise, Vaishak Belle, Paolo Felli, Sheila McIlraith, Tim Miller, Adrian R Pearce, and Liz Sonenberg. 2022. Efficient multi-agent epistemic planning: Teaching planners about nested belief. *Artificial Intelligence* 302 (2022), 103605. https://doi.org/10.1016/j.artint.2021.103605

[37] Katharina Mülling, Jens Kober, Oliver Kroemer, and Jan Peters. 2013. Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research* 32, 3 (2013), 263–279. https://doi.org/10.1177/0278364912472380 arXiv:https://doi.org/10.1177/0278364912472380

[38] Amit Kumar Pandey. 2018. Socially Intelligent Robots, the Next Generation of Consumer Robots and the Challenges. In *ICT Innovations 2016*, Georgi Stojanov and Andrea Kulakov (Eds.). Springer International Publishing, Cham, 41–46.

[39] Judea Pearl and Dana Mackenzie. 2018. *The book of why*. Basic Books, NY, USA.

[40] Anand S Rao and Michael P Georgeff. 1991. Modeling rational agents within a BDI-architecture. In *Proceedings of Principles of Knowledge Representation and Reasoning (KR)*. Second International Conference on Principles of Know ledge Representation and Reasoning, KR91, Cambridge, Massachusetts, USA, 473–484.

[41] Stuart Russell. 2019. *Human compatible: Artificial intelligence and the problem of control*. Penguin, USA.

[42] Lindsay Sanneman and Julie Shah. 2023. Transparent Value Alignment. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 557–560. https://doi.org/10.1145/3568294.3580147

[43] Shashank Shekhar, Anthony Favier, and Rachid Alami. 2024. An Epistemic Human-Aware Task Planner which Anticipates Human Beliefs and Decisions. arXiv:2409.18545 [cs.RO] https://arxiv.org/abs/2409.18545

[44] Maayan Shvo, Ruthrash Hari, Ziggy O'Reilly, Sophia Abolore, Sze-Yuh Nina Wang, and Sheila A. McIlraith. 2022. Proactive Robotic Assistance via Theory of Mind. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 9148–9155. https://doi.org/10.1109/IROS47612.2022.9981627

[45] Wiebe van der Hoek and Michael Wooldridge. 2003. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia logica* 75 (2003), 125–157.

[46] Hans van Ditmarsch and Willem Labuschagne. 2007. My beliefs about your beliefs: a case study in theory of mind and epistemic logic. *Synthese* 155 (2007), 191–209.

[47] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. 2007. *Dynamic Epistemic Logic*. Vol. 337. Springer Science & Business Media, Dordrecht, The Netherlands.
[48] Yanjing Wang. 2018. A logic of goal-directed knowing how. *Synthese* 195, 10 (2018), 4419–4439.
[49] Gerhard Weiss (Ed.). 2013. *Multiagent Systems, second edition* (2 ed.). MIT Press, Cambridge, Massachusetts, Chapter 15, 695–757.
[50] Heinz Wimmer and Josef Perner. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13, 1 (1983), 103–128. https://doi.org/10.1016/0010-0277(83)90004-5
[51] Michael J. Wooldridge and Simon Parsons. 1998. Intention Reconsideration Reconsidered.. In *ATAL (Lecture Notes in Computer Science, Vol. 1555)*, Jörg P. Müller, Munindar P. Singh, and Anand S. Rao (Eds.). Springer, Paris, France, 63–79. http://dblp.uni-trier.de/db/conf/atal/atal1998.html#WooldridgeP98