Built to Learn: An Intrinsically Motivated Approach to Incremental Learning in Robotics

Letícia M. Berto[1], **Ricardo R. Gudwin**[2], **Esther L. Colombini**[1]

[1]Institute of Computing – University of Campinas
Campinas, Brazil

[2]School of Electrical and Computer Engineering – University of Campinas
Campinas, Brazil

{leticia.berto,esther}@ic.unicamp.br, gudwin@unicamp.br

Abstract. *This research explores how robots can autonomously adapt to complex environments using a human-inspired cognitive development framework. By integrating intrinsic motivation and decision-making, robots employed reinforcement learning to balance internal states. An extended model introduced the concept of pleasure, enabling context-aware prioritization of needs and preferences. A child-inspired cognitive architecture guided learning, interaction, and survival, giving rise to distinct robot personalities that shaped human engagement. Building on this, an advanced model with Theory of Mind allowed two robots to collaborate through mutual understanding of needs. We validated our models on both simulated and real robots with varied architectures.*
Keywords: *Cognitive architecture, human-robot interaction, robotics, intrinsic motivation, theory of mind, reinforcement learning.*

Student level: Ph.D. **Date of conclusion:** 17/03/2025. Thesis available at https://hdl.handle.net/20.500.12733/30160.

1. Introduction

From scoring goals and scrubbing floors to assisting in surgeries, defusing bombs, and even exploring Mars—robots are no longer confined to science fiction; they've become part of everyday life. Many of these machines fall under the category of service robots, designed to perform useful tasks for humans beyond traditional industrial automation [G. S. VIRK and GELIN 2008]. As their capabilities expand, service robots are becoming especially valuable in aging societies, where they can help fill critical labor gaps in areas like healthcare and caregiving.

Unlike industrial robots, service robots operate in dynamic, unpredictable environments and must interact with people and other agents. This requires them to learn and adapt rather than follow pre-programmed routines. Humans develop these abilities through interaction with the world from an early age—learning from both success and failure [Piaget 1952]. Inspired by this, fields like developmental and cognitive robotics aim to create robots with human-like intelligence by drawing from disciplines such as psychology, neuroscience, and philosophy [Vernon et al. 2016, Lungarella et al. 2003, Cangelosi and Schlesinger 2015].

A critical aspect of human learning is social interaction. As Vygotsky [Vygotsky 1978] and others have shown, children develop cognitively and emotionally through engagement with others. They exchange knowledge, build values, and develop behavior through relationships with peers and adults. Emotion and motivation also play key roles in learning and decision-making, and are essential components in building socially intelligent robots.

Guided by these principles and inspired by Turing's vision of intelligence [Turing 1950], our research explores how robots can incrementally develop cognitive and social abilities. Through a series of experiments driven by intrinsic motivation, we structured our approach into four developmental phases (see Figure 1) that enable robots to gradually acquire increasingly complex behaviors.

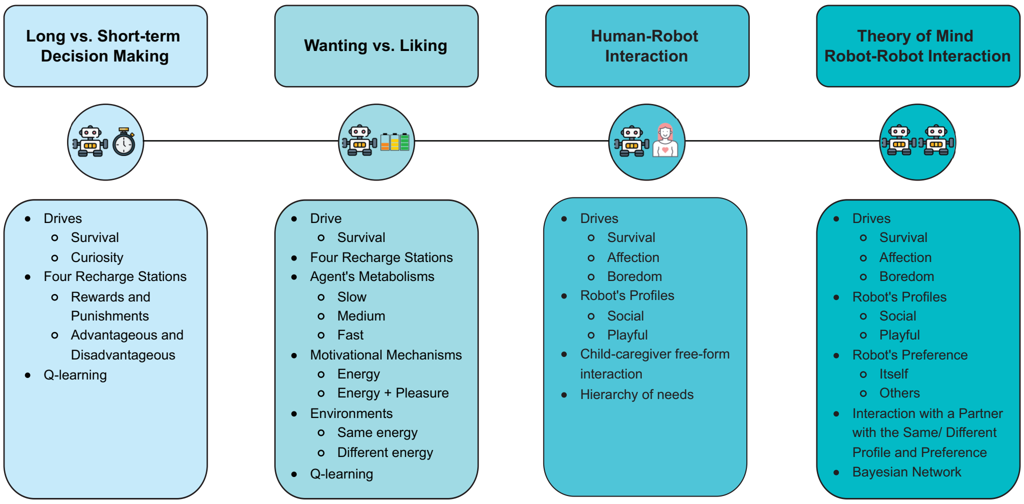

Figura 1. Aspects explored in this study, organized into four main research topics. The level of complexity progressively increases from left to right.

This work makes substantial and multifaceted contributions to the field of cognitive and developmental robotics, culminating in **twelve publications** in journals and conferences (with two currently under review or in progress). Our research began with a deep investigation into human cognitive development, particularly drawing from Piaget's theory of child development. Inspired by his stages of cognitive growth, we designed a series of incremental, developmentally-aligned experiments for robots [Berto et al. 2024b]. These experiments enabled us to assess the cognitive evolution of a "baby robot" interacting autonomously with its environment, leveraging cognitive architectures without external intervention [Berto et al. 2025b, Rossi et al. 2022]. Recognizing the critical role of motivation in autonomous behavior and decision-making, we developed and validated a motivational system for robotic agents. Our initial studies examined the trade-off between short-term and long-term decisions when balancing competing internal needs [Berto et al. 2021]. We then enriched this framework by incorporating pleasure as a motivational factor, exploring its effects across diverse agents and environmental contexts [Berto et al. 2024a].

Given the inherently social nature of human development, we extended our architecture to support human-robot interaction (HRI) scenarios [Berto et al. 2025c, Berto et al. 2024c]. Furthermore, we explored Theory of Mind as a tool to enhance robot-robot interactions, enabling more nuanced and peer-aware behavior in social settings [Berto et al. 2025a], and we plan to extend this capability to enhance social understanding and responsiveness in HRI contexts as well.

To support adaptive learning in motivated decision-making, we investigated various methods for optimizing internal drives through reinforcement and dual-system approaches [Rossi et al. 2025]. While prior research often emphasizes persistent, long-term needs as motivational drivers, we identified a critical gap in addressing temporary, fluctuating needs. Addressing this, we conducted experiments with a simulated humanoid agent, comparing decision-making outcomes using drive-based versus impulse-based motivational systems within a cognitive architecture framework [Rossi et al. 2024].

As all of our research is grounded in cognitive architectures, we also contributed tools to facilitate their development and interpretation. Notably, we introduced a visualization tool to aid in the design and analysis of intelligent agents [Áureo Marques et al. 2022]. Furthermore, we initiated a novel interdisciplinary connection between cognitive robotics and NeuroIS, proposing cognitive architectures—rooted in neuroscience—as a framework to model intelligent agents as natural cognitive systems [Baima et al. 2024].

As the main contributions of this work, we emphasize: (i) A formal framework integrating key psychological principles with robotics; (ii) A progressive series of experiments investigating the impact of social interaction and the integration of increasingly complex cognitive functions into the architecture; (iii) A motivational system incorporating multiple needs and balancing long-term vs. short-term decision-making; (iv) A motivation-based learning model driven by *wanting* (needs) and *liking* (subjective evaluation, or pleasure); (v) A curiosity- and affect-driven cognitive architecture for human-robot interaction; (vi) A Theory of Mind-driven motivational framework for social interaction; and (vii) Datasets collected in the experiments [Berto and Colombini 2025c, Berto and Colombini 2025a, Berto and Colombini 2025b].

Finally, the methodologies and systems developed through this research have been implemented and validated across a wide range of robotic platforms, including the simulated Pioneer P3DX, iCub, Pepper, Nao, and the simulated humanoid Marta [Begazo 2020]—demonstrating the flexibility and robustness of our approach across both simulated and physical environments.

2. Motivated Behavior — Incremental Approach

Hull's Drive Reduction Theory [Hull 1943], rooted in the concept of homeostasis [Cannon 1939], proposes that motivation arises from the body's need to maintain internal balance. A drive represents a state of arousal caused by unmet biological or physiological needs, with its intensity determining the strength of the resulting behavior. While the drive energizes behavior, it is learning—through reinforcement—that directs it. Behaviors that reduce the drive are reinforced, leading the individual to learn which actions are effective in restoring balance.

Building on Hull's Drive Reduction Theory, this research developed a motivational system with progressively increasing complexity to explore motivated behavior in autonomous robots. The incremental approach is depicted in Figure 2. The first two phases focus on individual learning, inspired by Piaget's theory. Phase 1 follows Hull's drive reduction theory, where the agent learns to balance multiple drives to maximize its well-being. In Phase 2, the cognitive architecture is adjusted to create agents with distinct profiles, exploring how even small modifications affect perception, learning, and decision-making. We also introduce preferences as a key motivator, examining how different environments and profiles influence these processes.

In the next two phases, we introduce social factors, drawing from Vygotsky's theory. Phase 3 applies Maslow's hierarchy of needs [Maslow 1981, Maslow 1943] to prioritize drives, altering motivation over time and in varying contexts. The agent begins to explore the motivations of others by observing emotional expressions in social interactions. Finally, in Phase 4, the agent gains Theory of Mind capabilities, understanding that others have different goals and motivations, which refines its decision-making. This phase also addresses situations where collaboration is not possible, enhancing the robot's ability to navigate social and environmental dynamics.

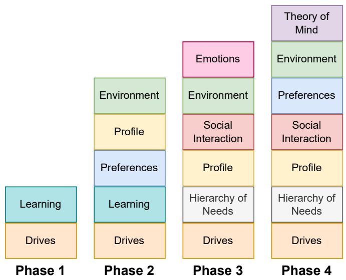

Figura 2. Incremental modules added to the cognitive architecture across research phases, with increasing complexity from left to right. Each phase aligns with a main research topic, as shown in Figure 1.

This incremental approach is essential for understanding motivated behavior as it enables the autonomous agent to adapt to more complex scenarios. By integrating both internal drives and social influences, the agent's decision-making evolves with context, offering a more sophisticated model of how autonomous systems balance personal and social needs. This approach lays the groundwork for robots to interact with humans in more emotionally aware and intelligent ways.

For each research phase, we developed a cognitive architecture based on a generic framework (Figure 3), customizing each component according to the specific robot, experimental setup, and objectives. Notably, the Theory of Mind modules were implemented only in the final phase.

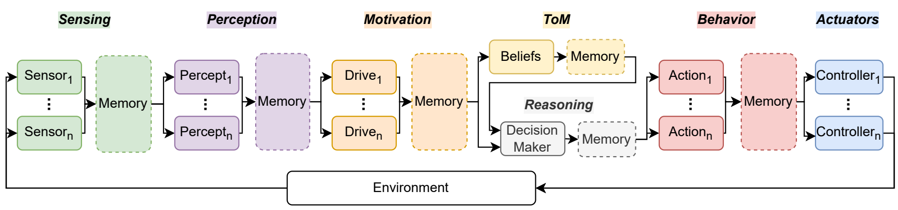

Figura 3. Generic framework used in all experiments.

3. Balancing Drives: Long-Term vs. Short-Term Decision-Making

Our goal was to investigate the motivational mechanism with multiple *drives*, and considering short and long-term decisions in this experiment. In short-term decisions, the agent learns how to reduce each *drive* reasoning only in the immediate effect, which we call "motivated behavior". In contrast, in the long-term decisions, the agent understands the consequences of their actions over time when trying to reduce the *drives* during the learning process, which we call "emergent motivated behavior" [Berto et al. 2021].

According to the defined learning policy, the agent is driven by two primary objectives: satisfying its curiosity through environmental exploration and fulfilling its survival needs by visiting power stations. Among these stations, A and B are ultimately detrimental over time, whereas C and D offer long-term benefits. Figures 4a and 4b present the outcomes of the testing phase after the agent has learned the optimal policy. The results show that under a short-term decision strategy, the agent prioritizes curiosity. This is expected, as exploring the environment yields immediate rewards without any short-term penalties—though it reduces the agent's lifespan due to neglecting power stations. In contrast, the long-term decision strategy leads to a more balanced satisfaction of both drives: the agent still values exploration but also recognizes the critical importance of visiting power stations for sustained survival.

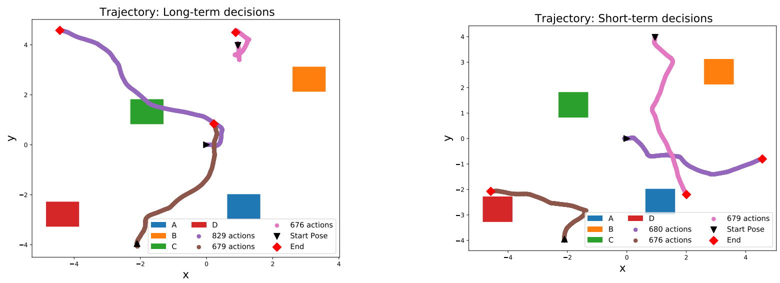

(a) Long-term decision **(b) Short-term decision**

Figura 4. Robot's trajectory. The squares are the power stations. Stations A and B are disadvantageous in the long term, while C and D are beneficial.

4. *Wanting vs. Liking*

We propose a motivation-based computational model to study how agents balance internal needs and pleasure when learning and making decisions [Berto et al. 2024a]. Grounded in the wanting vs. liking framework, we conducted 15 experiments with agents of varying metabolic rates in two distinct environments. Simulations took place in a 20×20 grid world with four fixed recharge stations offering different energy values.

Figure 5 presents results from the first 12 experiments. Across all cases, agents adapted their behavior based on internal needs, increasing steps taken and converging to a zero-reward optimal policy. Notably, slow metabolism agents showed temporary performance drops, likely due to parameter settings or unfamiliar situations, but recovered over time. In EXP01–EXP12, each agent experienced four scenarios under identical conditions. In the simplest setup (EXP01, 05, 09), where all recharge stations offered equal energy, agents remained close to stations as metabolic demands increased. Slower metabolism made it easier to manage drive and maximize survival. In environments with varied station values (EXP02, 06, 10), agents learned to maintain homeostasis, yielding similar rewards despite the added complexity.

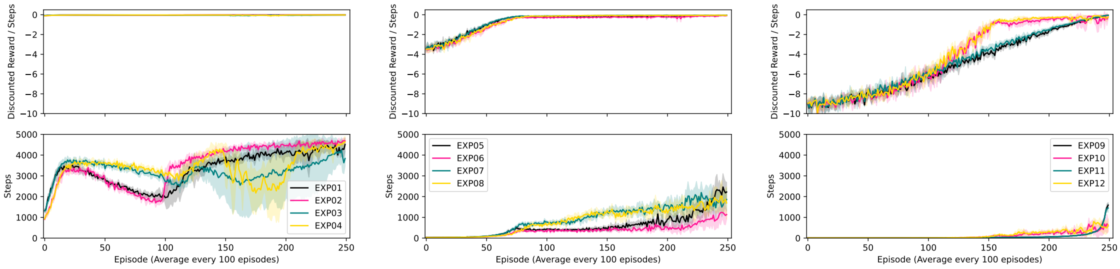

(a) Slow Metabolism **(b) Regular Metabolism** **(c) Fast Metabolism**

Figura 5. Average discounted reward per average actions (top) and number of actions (bottom) per episode for each metabolism in each experiment.

The learning time proved critical: faster metabolism agents faced more pressure to learn quickly, especially when stations were far or insufficiently rewarding. Despite the challenge, they successfully adapted. Environmental configuration significantly influenced learning efficiency. Agents performed better when station values aligned with their needs—low for slow, medium for regular, and high for fast metabolism. Initial battery level and position also shaped outcomes.

Motivation influenced behavior in agent- and environment-specific ways. Slow metabolism agents balanced need and pleasure, preferring high-energy or medium-energy stations as appropriate. Regular agents were more affected by hedonic value, often choosing stations with higher pleasure. Fast metabolism agents prioritized survival, aligning with Maslow's hierarchy, as even the most rewarding station failed to fully meet their needs. Future work could explore environments with richer hedonic options for these agents. In summary, both environment and motivational modeling strongly influence agent behavior and learning outcomes.

5. Curiosity and Affect-Driven Cognitive Architecture for HRI

This work explores how autonomous agents—humans and cognitively motivated robots with varying internal drives and value systems—can understand each other's needs through free-form interaction. Grounded in human-robot interaction, we experimentally examine how altering the relative importance of internal drives, within a fixed cognitive architecture, influences robot behavior and shapes interaction dynamics with human partners [Berto et al. 2025c].

Specifically, we investigate how different robot profiles—designed with varying levels of dependency and expectations—affect human perception and response. The goal is to assess whether humans recognize these profiles and adapt their behavior to help the robot reach a more pleasant state. The experiment was designed as a caregiver-child interaction scenario, allowing the iCub robot to act autonomously while supporting rich, natural exchanges with the human partner. The setup featured the robot positioned across from the participant, with a table full of toys placed between them (Figure 6). To capture both subjective and objective aspects of the interaction, we collected data from three sources: participant questionnaires, robot-logged behavioral data, and video recordings from an external camera.

A total of 36 participants (17 female, 19 male; mean age = 26.1, SD = 4.99), from five countries (Italy, Germany, France, Lebanon, Egypt), interacted with two robot

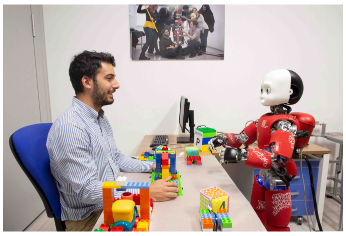

Figura 6. A participant and the robot iCub during the interaction.

profiles—Playful and Social—in a counterbalanced order. Participants were not informed about the differences between profiles. Two groups were formed: PS (Playful first) and SP (Social first). All participants provided informed consent and received € 12 compensation. The study was approved by the Regional Ethical Committee (Comitato Etico Regione Liguria, Application IIT_wHiSPER) and conducted in accordance with the Declaration of Helsinki.

We analyzed the average behavior of the PS and SP groups (Figure 7) and found similar patterns across both. As expected, the Playful robot spent more time in the *Play* state, while the Social robot spent more time either *playing* or *idle*. This was likely due to participants perceiving the Social robot as more needy, leading to increased touch and reducing the need for it to request interaction.

The Social robot also spent more time in the *Idle* state, consistent with enhanced comfort from frequent physical and visual engagement. It also showed greater variability in interactive and idle behaviors compared to the Playful profile. Both robots spent minimal time in the *Recharge* state, a design choice to maintain participant engagement.

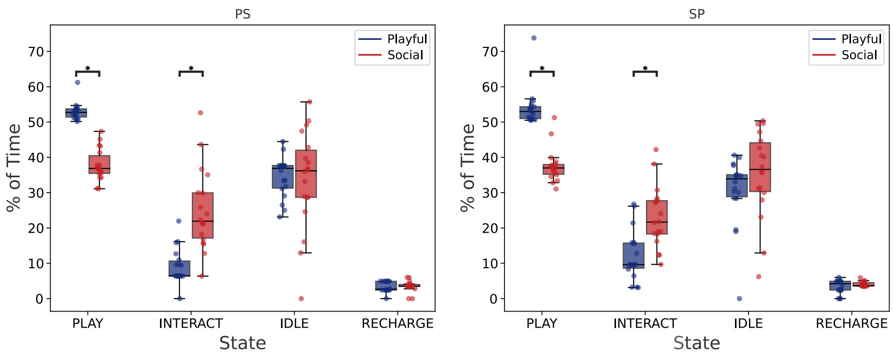

Figura 7. Average percentage of interaction time spent in the primary states, depicted across sessions.

Furthermore, we examined participants' behavior to determine whether they intuitively understood the robot's needs across different states, focusing on touch stimuli (Figure 8). Overall, participants frequently touched the robot in all states—including during *Recharge*—and adjusted their touch behavior according to the robot's profile. Touch had a stronger influence on the Social robot's state dynamics, leading to higher variability in its responses across conditions.

In conclusion, we explored how people interact with two robot profiles without knowing their differences—mirroring first-time encounters with strangers. Our results

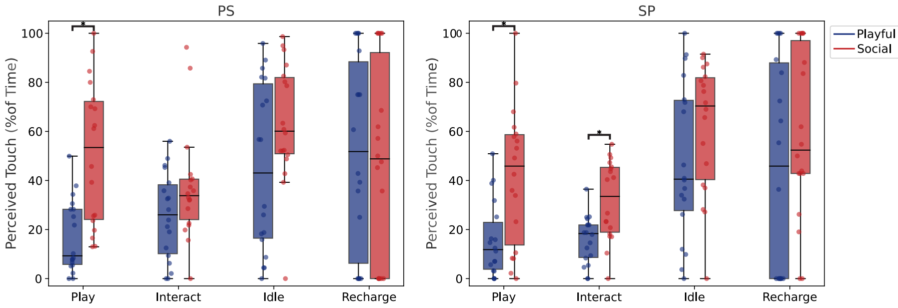

Figura 8. Average percentage of time the robot perceived the touch stimuli in each main state, depicted across sessions.

show that simply reordering motivational priorities can lead to distinct robot behaviors and perceptions. While the robot had its own preferences, interaction dynamics were shaped largely by human engagement. Despite variability across participants, the robot's behavior aligned with its profile and was generally understood—especially in the playful condition. These findings highlight the role of a robot's "personality" in shaping human perception and interaction.

6. A Theory of Mind-Driven Motivational Framework for Social Interaction

In this study, we address challenges in socially adaptive robotics by implementing a motivational system guided by internal drives and a simplified Theory of Mind (ToM). This allows robots to recognize their own needs while considering those of a partner, enabling them to adapt to unfamiliar individuals without prior knowledge [Berto et al. 2025a].

We model the robot as an older child—socially aware but still developing perspective-taking skills. It starts with initial assumptions about the partner's motivations and refines them through interaction. To investigate this, we designed two agents with the same foundational cognitive architecture but differing developmental profiles. MAR (Motivated Autonomous Robot) serves as a baseline, while MARTOM (Motivated Autonomous Robot with Theory of Mind) includes ToM capabilities, allowing it to infer and respond to others' needs. MARTOM can operate in two modes: Self-Priority, focusing on its own needs, and Other-Priority, assisting its partner unless its own needs become critical.

Experiments (shown in Table 1) were conducted in a controlled university lab using two face-to-face robots, Pepper and NAO, with a table of toys between them (Figure 9). NAO ran the MAR architecture, while Pepper operated as MARTOM. Two main interaction scenarios were tested: (1) MAR interacting with MARTOM (Figure 9a), to examine differences in decision-making under a shared motivational system; and (2) both robots running MARTOM (Figure 9b), to explore behavior arising from varying profiles and priorities. Each session lasted 10 minutes, during which both robots acted autonomously.

We organized our analysis around three main investigation goals, structuring the results accordingly. In all experiments, Agent A's configuration remained fixed, while Agent B's varied. This setup allowed us to assess the effects of each modification by observing changes in Agent A's responses. The outcomes are summarized in Table 2.

Tabela 1. Experiment configuration, including the ID, cognitive architecture, profile, and preference used by each agent.

EXP ID	Agent A			Agent B		
	CogArch	Profile	Preference	CogArch	Profile	Preference
1	MARTOM	Playful	Self	MAR	Playful	-
2	MARTOM	Playful	Other	MAR	Playful	-
3	MARTOM	Social	Self	MAR	Social	-
4	MARTOM	Social	Other	MAR	Social	-
5	MARTOM	Social	Self	MAR	Playful	-
6	MARTOM	Social	Other	MAR	Playful	-
7	MARTOM	Playful	Self	MAR	Social	-
8	MARTOM	Playful	Other	MAR	Social	-
9	MARTOM	Playful	Self	MARTOM	Playful	Self
10	MARTOM	Playful	Other	MARTOM	Playful	Other
11	MARTOM	Social	Self	MARTOM	Social	Self
12	MARTOM	Social	Other	MARTOM	Social	Other
13	MARTOM	Playful	Self	MARTOM	Playful	Other
14	MARTOM	Social	Self	MARTOM	Social	Other
15	MARTOM	Playful	Self	MARTOM	Social	Self
16	MARTOM	Playful	Other	MARTOM	Social	Other
17	MARTOM	Playful	Self	MARTOM	Social	Other
18	MARTOM	Playful	Other	MARTOM	Social	Self

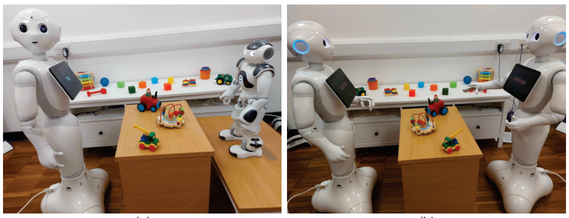

(a) (b)

Figura 9. Experimental setup.

We investigated how changes in cognitive architecture, profiles, and preferences affect agent interactions, conducting 18 experiments with these factors introduced incrementally.

Our results show that simply adding Theory of Mind (ToM) does not guarantee cooperative behavior. When one agent uses the MAR architecture and the other is a MARTOM agent with a self-serving preference—or when both agents are self-prioritizing MARTOMs—both prioritize their own needs, with little regard for the other (Experiments 1, 3, 5, 7, 9, 11, 15). In these cases, ToM increases awareness of the partner's needs but does not influence behavior unless coupled with a cooperative preference.

Similar outcomes occurred when a MAR agent interacted with an altruistic MARTOM that had a different profile. The MAR agent, by design, does not assist, while the MARTOM remains preoccupied with its own urgent needs. Here, the profile had a stronger impact than the preference (Experiments 6 and 8).

This pattern persisted in experiments where both agents shared the same cognitive architecture but differed in profile or preference (Experiments 15–18). These results suggest that ToM alone is insufficient; profiles and preferences play a critical role in shaping decision-making. Further work is needed to explore richer decision models and alternative drive configurations.

Despite these challenges, promising behaviors emerged. When agents shared the same profile, two effective patterns were observed: (I) altruistic MARTOMs paired with MAR agents reduced MAR's stress (Experiments 2 and 4), though the benefit was not reciprocal; (II) interactions between MARTOM agents—where at least one prioritized the other's needs—led to strong cooperation and mutual satisfaction (Experiments 10, 12,

Tabela 2. Summary of results considering each parameter variation.

Variation	EXP ID	Agent A					Agent B				
		Actions To Itself	Actions To Other	Priority To Help Other	Success ToM Affection	Success ToM Exploration	Actions To Itself	Actions To Other	Priority To Help Other	Success ToM Affection	Success ToM Exploration
Cognitive Architecture	1	100%	0%	0%	40%	100%	100%	0%	-	-	-
	9	100%	0%	0%	43.48%	100%	100%	0%	0%	43.48%	100%
	3	100%	0%	0%	100%	13.33%	100%	0%	-	-	-
	11	100%	0%	0%	100%	13.7%	100%	0%	0%	100%	13.7%
	2	60%	40%	40%	93.33%	73.33%	40%	0%	-	-	-
	10	26.09%	73.91%	17.39%	95.65%	86.96%	4.35%	95.65%	39.13%	91.3%	86.96%
	4	65.91%	34.09%	34.09%	72.73%	88.64%	36.36%	0%	-	-	-
	12	24%	76%	16%	88%	96%	4%	96%	36%	92%	92%
Profile	9	100%	0%	0%	43.48%	100%	100%	0%	0%	43.48%	100%
	15	100%	0%	0%	97.3%	24.32%	100%	0%	0%	37.84%	100%
	10	26.09%	73.91%	17.39%	95.65%	86.96%	4.35%	95.65%	39.13%	91.3%	86.96%
	16	100%	0%	0%	97.06%	17.65%	100%	0%	0%	32.35%	100%
Preferences	9	100%	0%	0%	43.48%	100%	100%	0%	0%	43.48%	100%
	13	34.48%	24.14%	0%	86.21%	58.62%	27.59%	41.38%	34.48%	93.1%	86.21%
	11	100%	0%	0%	100%	13.7%	100%	0%	0%	100%	13.7%
	14	29.41%	20.59%	0%	61.76%	88.24%	23.53%	35.29%	29.41%	85.29%	94.12%
More than one parameter	5	100%	0%	0%	47.73%	84.09%	100%	0%	-	-	-
	6	95%	5%	5%	42.5%	95%	100%	0%	-	-	-
	7	100%	0%	0%	100%	22.22%	100%	0%	-	-	-
	8	100%	0%	0%	100%	20%	100%	0%	-	-	-
	17	100%	0%	0%	97.22%	22.22%	100%	0%	0%	36.11%	100%
	18	100%	0%	0%	97.22%	22.22%	100%	0%	0%	36.11%	100%

13, 14).

While we used agents with identical needs and drive dynamics for simplicity, future studies could explore heterogeneous agents with differing motivations and rates of change. Additionally, our current system limits ToM to short-term support; extending this to include longer-term modeling may yield more adaptive, collaborative behavior.

Despite some limitations, our findings highlight the value of Theory of Mind in social interactions. It enhances adaptability and enables more collaborative behavior. The proposed architecture effectively supports interaction with unfamiliar agents, allowing the robot to infer and respond to the partner's motivations without prior knowledge. While the experimental setup imposed constraints on drives and actions, the architecture demonstrates strong potential for real-world deployment.

7. Conclusion

This work investigated the behavior of intrinsically motivated autonomous robots in both social and non-social settings. Using a minimal cognitive architecture inspired by human development, we focused on internal drives as the foundation for autonomous decision-making. Grounded in psychological theories—particularly Hull's drive reduction theory—we developed a flexible motivational model that allowed robots to act purposefully and adaptively, without relying on scripted behaviors. Our results demonstrate the value of intrinsic motivation, adaptation, and interaction as core components of cognitively autonomous systems.

Building on this foundation, several directions emerge for future research. Integrating emotional states into decision-making could help robots respond more sensitively to context, enhancing social engagement and adaptability. Likewise, adding long-term memory would allow robots to personalize interactions by recalling past experiences, promoting continuity and deeper human-robot relationships.

As robots begin to operate in team-based environments, understanding multi-agent dynamics becomes essential. Future work should explore how agents with different motivations and cognitive capabilities coordinate, negotiate, and collaborate in complex, real-world scenarios.

Finally, trust remains a critical factor in effective human-robot interaction. Studying how trust is built, maintained, or repaired—from both human and robotic perspectives—will support the development of agents that are not only autonomous but also socially reliable and cooperative.

Referências

Baima, R. L., Berto, L. M., and Roth, T. (2024). Expanding the scope—cognitive robotics meets neurois. In Davis, F. D., Riedl, R., Brocke, J. v., Léger, P.-M., Randolph, A. B., and Müller-Putz, G. R., editors, *Information Systems and Neuroscience*, pages 195–203, Cham. Springer Nature Switzerland.

Begazo, M. F. T. (2020). A learning-based model-free controller for decoupled humanoid robot walking. *Instituto de Ciência e Tecnologia de Sorocaba - Universidade Estadual Paulista "Júlio de Mesquita Filho"(UNESP) (Master thesis)*.

Berto, L., Costa, P., Simões, A., Gudwin, R., and Colombini, E. (2024a). A motivational-based learning model for mobile robots. *Cognitive Systems Research*, 88:101278.

Berto, L., Hellou, M., Sciutti, A., Gudwin, R., Colombini, E., and Cangelosi, A. (2025a). A theory of mind-driven motivational framework for social interaction with autonomous cognitive robots. In *2025 IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. Accepted.

Berto, L., Rossi, L., Costa, P., Simões, A., Gudwin, R., and Colombini, E. (2025b). Where are my toys going? exploring incremental learning in object tracking and permanence. *IEEE Transactions on Cognitive and Developmental Systems*. Paper in progress.

Berto, L., Rossi, L., Rohmer, E., Costa, P., Gudwin, R., Simões, A., and Colombini, E. (2024b). Piagetian experiments to devrobotics. *Cognitive Systems Research*, 83:101170.

Berto, L., Tanevska, A., Cirne, A., Costa, P., Simões, A., Gudwin, R., Rea, F., Colombini, E., and Sciutti, A. (2024c). Curiosity and affect-driven cognitive architecture for hri. In *2024 IEEE International Conference on Development and Learning (ICDL)*. Late-Breaking Results.

Berto, L., Tanevska, A., Cirne, A., Costa, P., Simões, A., Gudwin, R., Rea, F., Colombini, E., and Sciutti, A. (2025c). Curiosity and affect-driven cognitive architecture for hri. *IEEE Transactions on Affective Computing*, pages 1–18.

Berto, L. M. and Colombini, E. L. (2025a). Human-robot interaction: motivated autonomous robots. `https://doi.org/10.25824/redu/ZWJPCW`, Repositório de Dados de Pesquisa da Unicamp, V1.

Berto, L. M. and Colombini, E. L. (2025b). Theory of mind: motivational framework for social interaction with robots. `https://doi.org/10.25824/redu/MVWJGS`, Repositório de Dados de Pesquisa da Unicamp, V1.

Berto, L. M. and Colombini, E. L. (2025c). Wanting vs. liking: motivational framework for mobile robots. `https://doi.org/10.25824/redu/BWRUOL`, Repositório de Dados de Pesquisa da Unicamp, V1.

Berto, L. M., Costa, P. D. P., Simoes, A. S., Gudwin, R. R., and Colombini, E. L. (2021). An iowa gambling task-based experiment applied to robots: A study on long-term decision making. In *2021 IEEE International Conference on Development and Learning (ICDL)*, pages 1–6.

Cangelosi, A. and Schlesinger, M. (2015). *Developmental Robotics: From Babies to Robots*. The MIT Press.

Cannon, W. B. (1939). The wisdom of the body.

G. S. VIRK, S. M. and GELIN, R. (2008). Iso standards for service robots. In Marques, L., de Almeida, A., Tokhi, M. O., and Virk, G. S., editors, *Advances in Mobile Robotics*, pages 133–138. World Scientific.

Hull, C. L. (1943). Principles of behavior: An introduction to behavior theory.

Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.

Maslow, A. H. (1943). A theory of human motivation. *Psychological review*, 50(4):370.

Maslow, A. H. (1981). *Motivation and personality*. Prabhat Prakashan, New Delhi.

Piaget, J. (1952). *The origins of intelligence in children.* The origins of intelligence in children. W W Norton & Co, New York, NY, US.

Rossi, L., Berto, L., Costa, P., Gudwin, R., Colombini, E., and Simões, A. (2025). Dual or unified: Optimizing drive-based reinforcement learning for cognitive autonomous robots. *Cognitive Systems Research*. First round review.

Rossi, L., Berto, L., Rohmer, E., Costa, P., Gudwin, R., Colombini, E., and Simões, A. (2022). Aprendizado procedimental e sensório-motor em robôs cognitivos. In *Anais Estendidos do XIV Simpósio Brasileiro de Robótica e XIX Simpósio Latino-Americano de Robótica*, pages 73–84, Porto Alegre, RS, Brasil. SBC.

Rossi, L. L., Berto, L., Costa, P. D. P., Gudwin, R., Colombini, E., and Simões, A. (2024). Drives and impulses: Shaping motivation and procedural learning for humanoid robots. In *2024 IEEE International Conference on Development and Learning (ICDL)*, pages 1–8.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236):433–460.

Vernon, D., von Hofsten, C., and Fadiga, L. (2016). Desiderata for developmental cognitive architectures. *Biologically Inspired Cognitive Architectures*, 18:116–127.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*, volume 86. Harvard university press.

Áureo Marques, Coletta, L., Silva, A., Paraense, A., Berto, L., Costa, P., Colombini, E., Simões, A., and Gudwin, R. (2022). Visualization tools for monitoring and debugging a cognitive architecture using cst. *Procedia Computer Science*, 213:528–535. 2022 Annual International Conference on Brain-Inspired Cognitive Architectures for Artificial Intelligence: The 13th Annual Meeting of the BICA Society.