A Motivation-Driven Incremental Learning Framework for Robotics

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

[1]Institute of Computing – University of Campinas (UNICAMP)
Campinas, Brazil

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

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

Abstract. *This thesis presents a computational framework for intrinsically motivated autonomous agents, formalizing intrinsic motivation (IM) as a multi-objective reinforcement learning problem integrating drive regulation, hedonic valuation, hierarchical need prioritization, and Theory of Mind, enabling adaptive, long-term decision-making and socially aware interaction. Validated on simulated and physical robotic platforms, the framework demonstrates improved behavioral stability, policy reshaping via hedonic modulation, and enhanced cooperation among agents with compatible motivational profiles and one altruistic agent. This work extends beyond robotics, offering a computational model of IM that bridges cognitive science and autonomous decision-making.*

1. Introduction

Over the past decades, robotics has advanced rapidly, with social robotics emerging as a transformative field. The vision of robots collaborating seamlessly with humans is increasingly realistic, yet real-world deployment remains challenging: robots often require constant supervision or fail when environments change. To move beyond scripted behavior, robots must anticipate tasks, adapt to dynamic contexts, and operate autonomously. Human development provides a natural inspiration for achieving this autonomy, particularly in motivation, learning, and decision-making. Motivation enables humans to prioritize competing needs, act proactively, and adapt to evolving circumstances. Translating this principle to robotics allows agents to self-regulate their behavior, reducing reliance on explicit instructions while enhancing adaptability in complex environments. Moreover, human behavior is socially embedded: decisions are shaped not only by personal needs but also by the needs of others and environmental context. Incorporating these social and contextual dynamics is critical for creating robots capable of meaningful collaboration.

Building on these principles, we designed an incremental cognitive architecture that supports intrinsic motivation, learning, and adaptive decision-making. Robots gradually acquire cognitive and social capabilities through a structured, four-phase developmental process (Figure 1). The system integrates multiple decision-making strategies-including rule-based control, reinforcement learning, and Bayesian reasoning- allowing autonomous adaptation in both simulated and physical environments.

---

*This study was financed, in part, by the São Paulo Research Foundation (FAPESP), Brasil. Process Number #2021/07050-0, and has been partly supported by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

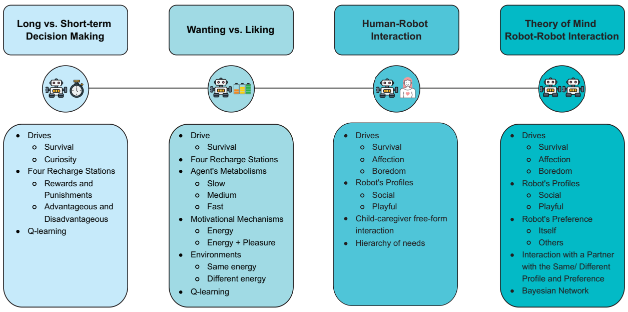

Figura 1. Research topics. The level of complexity increases from left to right.

1.1. Objective

This work[1] addresses nine research questions, examined in depth in the thesis and summarized in this paper. First, it investigates whether a cognitive agent can learn to balance multiple competing needs based on Hull's drive reduction theory (RQ1), and whether long-term decision-making strategies grounded in this theory lead to improved learning outcomes compared to short-term strategies (RQ2). It further examines whether an artificial agent can balance need satisfaction ("*wanting*") with hedonic evaluation ("*liking*") (RQ3). The research also explores how environmental factors shape an agent's learning strategies (RQ4) and how different artificial agents adapt their learning processes under identical environmental conditions (RQ5). In social contexts, it analyzes how interaction dynamics vary as a function of different robot profiles (RQ6), whether people can recognize distinct robot needs and how they evaluate these profiles (RQ7), and whether an agent can enhance its recognition of others' motivations by leveraging its own needs to bootstrap a Theory of Mind (RQ8). Finally, it investigates how cognitive architecture, agent profiles, and individual preferences jointly influence decision-making during social interactions between autonomous agents (RQ9).

1.2. Main Scientific Contributions to Computer Science

This work resulted in **12 peer-reviewed journal and conference publications**, with one additional manuscript currently under review. The research was also **awarded as the Best PhD Thesis at the 2025 National Contest of PhD and MSc Theses on Robotics (CTDR 2025)**. This thesis advances the state of the art in Cognitive Robotics and Autonomous Systems through a unified framework that integrates psychological principles, motivational modeling, and cognitive architectures for autonomous agents. At its core, the work develops a formal computational model of intrinsic motivation as a multi-objective reinforcement learning problem, where internal drives are represented as dynamic state variables shaping both reward and policy learning. Unlike prior rule-based or threshold-driven approaches, this model enables agents to autonomously balance competing needs while preserving long-term behavioral stability, providing a scalable formulation of homeostatic regulation within Markov Decision Process frameworks.

Building on this foundation, the thesis introduces a novel distinction between *wanting* (drive-based necessity) and *liking* (hedonic valuation), showing how subjective pleasure influences exploration, learning, and policy selection. The work further implements

---

[1] This is a short summary of Letícia Berto's Ph.D. thesis. Full text is available at https://hdl.handle.net/20.500.12733/30160.

a hierarchical, context-sensitive mechanism inspired by Maslow's hierarchy of needs, allowing dynamic prioritization of drives based on internal state and environmental conditions during social interaction. Together, these contributions advance computational models of adaptive, multi-criteria decision-making and validate the impact of hedonic and hierarchical modulation on agent behavior. A Theory-of-Mind extension integrates probabilistic belief representation with motivational reasoning, demonstrating that cooperative behavior in social settings emerges not from belief inference alone, but from the interaction of ToM with compatible motivational profiles and altruistic preferences. This provides a computational perspective on social reasoning, highlighting the interdependence of belief and motivational architectures in autonomous agents.

The thesis also presents a modular, incremental cognitive architecture that unifies perception, memory, motivation, reasoning, learning, and social inference. This architecture supports curiosity- and affect-driven learning, enabling autonomous, socially adaptive behavior across human–robot interaction scenarios. It has been validated across multiple robotic platforms—including iCub, Nao, Pepper, Pioneer P3DX, and the simulated Marta—demonstrating portability, scalability, and robustness. Experiments include 18 controlled multi-agent studies, 15 motivational experiments in varied environmental contexts, and human–robot interaction studies involving 36 participants, generating datasets that provide a structured methodology for evaluating motivational and cognitive architectures in both individual and social contexts.

By integrating intrinsic motivation, hedonic valuation, hierarchical prioritization, social reasoning, and curiosity-driven learning, this work contributes a framework for cognitively and socially competent autonomous systems. Its validation across simulated and physical platforms, in collaboration with the Italian Institute of Technology and the University of Manchester, demonstrates the robustness and applicability of the approach, advancing autonomous cognitive robotics at both national and international levels.

2. Related Work

To develop autonomous and adaptive agents, researchers have increasingly focused on integrating motivational systems, homeostatic regulation, and hedonic influences into decision-making processes. Early works used drives to regulate behavior via predefined rules or homeostatic thresholds [Konidaris and Barto 2006, Salichs and Malfaz 2011, Breazeal et al. 1998], while others extended this to social contexts [Cao et al. 2017, Breazeal 2004, Cañamero 1997]. Hedonic values and pleasure have been shown to influence adaptive behavior and decision-making, though they remain underexplored in robotics [Cos et al. 2013, Lewis and Canamero 2016]. In human-robot interaction (HRI), social robots require adaptive cognitive mechanisms to recognize and respond to human cues, predict intentions, and adjust behavior dynamically [Goodrich et al. 2008, Breazeal 2003, Scassellati 2002]. ToM has emerged as a key component, enabling robots to infer mental states and motivations of humans or other agents, thereby improving trust, collaboration, and interaction quality [Romeo et al. 2022, Yu et al. 2024, Hellou et al. 2023]. Probabilistic ToM and trust models have been successfully applied in developmental scenarios, allowing robots to anticipate false beliefs, track preferences, and adapt actions in multi-agent contexts [Vinanzi et al. 2019, Hellou et al. 2024].

Despite these advances, no prior work integrates: (i) multi-drive regulation for short- versus long-term decision-making, (ii) hedonic valuation versus needs across agents and environments using reinforcement learning, (iii) hierarchical prioritization and social understanding in free-form interactions with diverse robot profiles, and (iv) ToM reasoning within a computational architecture validated on different physical robots.

3. Motivated Behavior - Incremental Approach

Drawing on Hull's Drive Reduction Theory [Hull 1943], we developed a motivational system of increasing complexity, as illustrated in Figure 2a. The first two phases focus on individual learning and homeostasis. In Phase 1, the agent learns to mitigate internal "tensions" by balancing multiple drives to maximize its overall well-being. In Phase 2, we adjusted the architecture to explore the emergence of distinct robot profiles. By introducing preferences as a primary motivator, we analyzed how minor cognitive variations influence the agent's perception, learning, and decision-making across varied environments. The subsequent phases integrate social dimensions. In Phase 3, incorporating Maslow's Hierarchy of Needs [Maslow 1981], the agent dynamically prioritizes drives based on temporal and environmental contexts. Social awareness begins, as agents learn to infer others' motivations by observing emotional expressions during interaction. Finally, Phase 4 introduces ToM modules, allowing agents to represent the internal states, goals, and motivations of other entities. This capability refines decision-making by enabling the robot to distinguish between collaborative opportunities and scenarios where cooperation is unfeasible, thereby navigating complex social and environmental dynamics. Each phase used a shared cognitive architecture (Figure 2b), adapted to different robots and experimental goals, with ToM modules introduced only in the final phase.

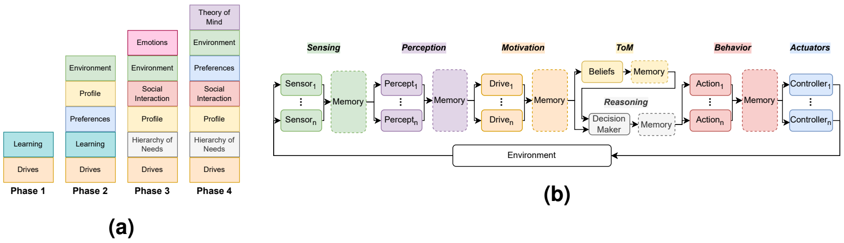

Figura 2. (a) Incremental additions to the cognitive architecture; each phase corresponds to a main research topic (see Fig. 1). (b) Generic framework.

4. Balancing Needs: Long-Term vs. Short-Term Decision-Making

Decision-making is central to intelligent behavior and involves balancing short-term reactivity with long-term planning. In robotics, this balance is critical: purely reactive strategies lead to inefficient or unstable behavior, while purely long-term planning reduces adaptability. Inspired by the Iowa Gambling Task (IGT) [Bechara et al. 1994], we designed a simplified robotic experiment to study motivational decision-making under uncertainty [Berto et al. 2021]. Under the defined learning policy, the agent pursues two primary objectives: satisfying curiosity through environmental exploration and maintaining survival by visiting power stations.

Using the motivational model, the agent learned to regulate multiple drives through short- and long-term strategies. Short-term decision-making led to impulsive, curiosity-driven behavior (Figure 3a), whereas long-term decision-making (Figure 3b) produced a more balanced regulation of both drives, favoring sustained survival. While both drives were assigned equal priority, the results highlight the need for dynamic prioritization-such as a hierarchy of needs-to better support adaptive behavior. Overall, the findings align with IGT outcomes, showing that long-term decision-making enables more stable and effective drive regulation than short-term strategies.

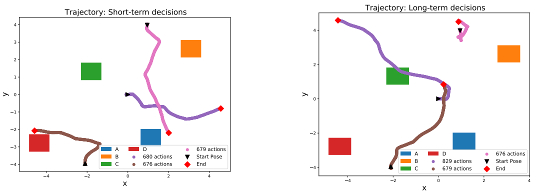

(a) Short-term **(b) Long-term**

Figura 3. Robot trajectory and power stations (A–B harmful, C–D beneficial).

5. *Wanting* vs. *Liking*

In real-world systems, multiple *needs* increase decision-making complexity, making learned behavior harder to interpret. To investigate this, we conducted 15 experiments with agents using two motivational models (*wanting*; *wanting* + *liking*) in two 20×20 grid-world environments containing four recharge stations differing in energy and hedonic value. Agents had three metabolic rates-slow, regular, and fast-to evaluate how internal characteristics influence behavior in identical environments [Berto et al. 2024].

For each metabolism, the first scenario had all recharge stations providing equal energy, followed by a second scenario with distinct energy values. The third and fourth scenarios repeated these configurations with the addition of a distinct hedonic (pleasure) value at each station. Across all scenarios, agents adapted their behavior based on internal needs, gradually converging to zero-reward optimal policies. Slow metabolism agents remained close to stations when energy demands increased and balanced need with pleasure, while faster metabolism agents prioritized survival over hedonic value. Environmental configuration significantly influenced learning efficiency, with agents performing best when station values aligned with their metabolic requirements. Figure 4 presents the results for the first 12 experiments, showing how metabolic rate, environmental configuration, and motivational model affect exploration and station visits. Statistical analysis of station visits (ANOVA followed by Tukey's test, $\alpha = 95\%$) confirms significant differences across environments and motivational models, demonstrating that both internal and external factors shape learned decision-making policies.

6. Curiosity and Affect-Driven Cognitive Architecture for HRI

We hypothesize that social interaction-and the ability to understand others-can be grounded in motivational drives that shape decision-making processes. Drawing on Maslow's hierarchy of needs, we consider social affiliation a need typically prioritized below physiological requirements such as survival, highlighting a central challenge in robotics: dynamically balancing competing needs within a unified framework.

Building on this perspective, this work investigates how autonomous agents-humans and cognitively motivated robots with distinct internal drives and value systems-can infer each other's needs through free-form interaction. Within an HRI framework, we

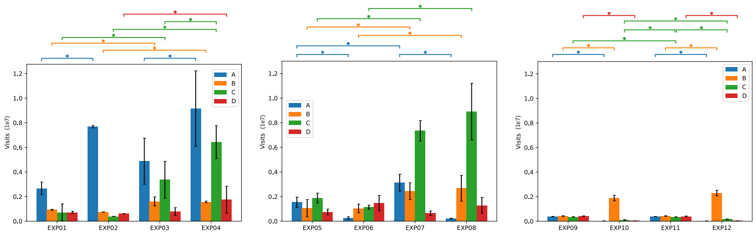

(a) Slow **(b) Regular** **(c) Fast**

Figura 4. Mean ± SD visits per station across metabolisms; * indicates $p < 0.05$.

examine how modulating the relative importance of internal drives within a fixed cognitive architecture influences robot behavior and shapes interaction dynamics with human partners [Berto et al. 2025b]. We focus on how different robot profiles, characterized by varying levels of dependency and expectations, affect human perception and behavior. Specifically, we assess whether humans recognize these profiles and adapt their actions to help the robot reach a more pleasant internal state. The experiment adopts a caregiver–child interaction paradigm, allowing the iCub robot to behave autonomously while engaging in natural interactions with participants. Facing the participant, the robot stood on the other side of a table filled with toys (Figure 5a). Data were collected through questionnaires, robot-logged behavioral measures, and external video recordings.

Thirty-six participants (17 female, 19 male; mean age 26.1, SD = 4.99) interacted with two robot profiles-Playful and Social- without being informed of the profile differences. Participants were assigned to two groups (PS and SP), 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. Figure 5b illustrates that, as anticipated, the proportion of time allocated to the *Recharge* state is notably low for both robot profiles. Conversely, the duration of time spent in the *Idle* state exhibits considerable variability. Upon closer examination of the data, it was observed that lower values corresponded to instances where participants provided insufficient stimuli to the robot, resulting in prolonged periods in either the *Interact* or *Play* state. Conversely, higher values indicated scenarios in which participants provided ample stimuli to the robot, satisfying its drives for an extended period. The distribution of the *Play* and *Interact* states is concentrated below and above the identity line, respectively, indicating that participants tended to engage the robot for longer periods in states that aligned with its inherent profile.

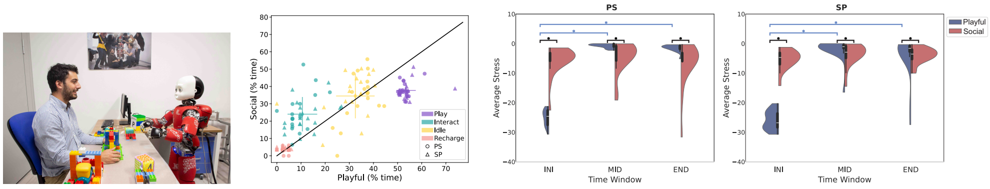

(a) **(b)** **(c)**

Figura 5. (a) Interaction setup, (b) time allocation across states (crosses indicate state averages), (c) robot stress levels over the interaction.

To examine stress dynamics during interaction (measured by being below homeostasis), we divided each session into three phases: beginning (INI), middle (MID), and end (END), as illustrated in Figure 5c. Stress levels were highest at the beginning of the interaction, reflecting the initial imbalance of the robot's internal drives and participants' exploratory behavior. As the interaction progressed, participants adapted to the robot's responses, leading to a reduction and subsequent stabilization of stress during the middle and end phases. We conducted two levels of statistical analysis: (i) within-profile comparisons across time windows and (ii) between-profile comparisons at each phase. A significant decrease in stress over time was observed only for the Playful robot, indicating that participants progressively learned how to regulate its internal state. Comparisons between profiles revealed significant differences across all phases: the Playful robot exhibited higher stress at the beginning, whereas the Social robot showed higher stress than the Playful during the middle and end of the interaction.

Interaction order further influenced stress regulation. Participants who interacted first with the Playful robot (PS group) generally induced lower stress overall, likely due to higher engagement and richer stimulation early on. In contrast, participants who first encountered the Social robot (SP group) showed higher stress levels in the subsequent Playful interaction, suggesting that prior experience led them to reduce stimulation. Specifically for the playful robot, the higher stress levels shown by the SP group indicate that participants understood the robot's profile better by not providing unnecessary input stimuli, which led to faster increases in stress.

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

To support socially adaptive robot behavior, agents must balance their own needs with those of others and adjust decisions accordingly. In this work, we investigate this challenge using a motivational system combined with a Theory of Mind mechanism, allowing robots to consider both internal drives and inferred partner needs during interaction [Berto et al. 2025a]. We model the robot as an older child with emerging perspective-taking abilities, starting from initial assumptions about a partner's motivations and refining them through interaction. Two agents sharing the same cognitive architecture were designed: MAR (Motivated Autonomous Robot), serving as a baseline, and MARTOM, which incorporates ToM. MARTOM operates in two modes: Self-Priority, focusing on its own needs, and Other-Priority, assisting the partner unless its own needs become critical.


(a) (b)

Figura 6. Experimental setup.

Experiments (Table 1) were conducted in a controlled laboratory setting using two face-to-face robots, Pepper and NAO, positioned across a table of toys (Figure 6). NAO ran MAR, while Pepper ran MARTOM. Two interaction scenarios were tested: MAR interacting with MARTOM, and both robots running MARTOM with different priority settings. Each session lasted 10 minutes, during which both robots acted fully 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. Experiments including IDs, architecture, profile, and preference.

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

We conducted 18 experiments to analyze how cognitive architecture, agent profiles, and preferences influence interaction. Results show that ToM alone does not ensure cooperation: when agents prioritize their own needs-either as MAR agents or self-focused MARTOM agents-interactions remain largely non-cooperative, despite awareness of the partner's state. In these cases, profiles often had a stronger effect than preferences on behavior. Cooperative behavior emerged only when ToM was combined with compatible profiles and cooperative preferences. Altruistic MARTOM agents reduced partner stress when interacting with MAR agents, though benefits were not mutual. Strongest cooperation occurred when at least one MARTOM agent prioritized the other's needs, leading to mutual satisfaction. These findings indicate that ToM must be coupled with appropriate motivational profiles to shape decision-making effectively. While the current setup assumes homogeneous agents and short-term reasoning, the architecture shows strong potential for scalable, adaptive social interaction without prior knowledge of partners, supporting its applicability to real-world scenarios.

8. Conclusion

This work investigates intrinsically motivated autonomous robots operating in both social and non-social environments through a minimal, developmentally inspired cognitive architecture. By grounding decision-making in internal drives rather than predefined scenarios or scripted interactions, the proposed motivational model enables adaptive, goal-directed behavior. The results underscore the central role of intrinsic motivation and interaction in achieving cognitive autonomy.

This research advances cognitive robotics by formalizing and validating an intrinsically motivated, self-regulating framework inspired by Hull's drive reduction theory. The integration of multi-drive motivational systems supports adaptive learning and sustainable long-term autonomy. By extending motivational processes to social cognition and Theory of Mind, the framework enhances socially intelligent robotics and human–robot interaction. Its progressive experimental methodology and open research resources promote reproducibility and interdisciplinary collaboration, offering both foundational theoretical contributions and scalable solutions for cognitively autonomous, socially adaptive

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%

robotic systems. Future work will focus on incorporating emotional modulation and long-term memory mechanisms to enable context-sensitive and personalized interactions. Expanding the framework to multi-agent systems with heterogeneous motivational profiles represents another key direction. Additionally, investigating mechanisms of trust formation and regulation will be critical for developing socially reliable, cooperative robots capable of operating effectively in real-world environments.

Referências

Bechara, A., Damasio, A. R., Damasio, H., and Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1-3):7–15.

Berto, L., Costa, P., Simões, A., Gudwin, R., and Colombini, E. (2024). 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 motivational framework for social interaction with autonomous cognitive robots. In *2025 34th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 429–436.

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

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.

Breazeal, C. (2003). Toward sociable robots. *Robotics and autonomous systems*, 42(3):167–175.

Breazeal, C. (2004). *Designing sociable robots*. MIT press, United States of America.

Breazeal, C. et al. (1998). A motivational system for regulating human-robot interaction. In *Aaai/iaai*, pages 54–61.

Cañamero, D. (1997). A hormonal model of emotions for behavior control. *VUB AI-Lab Memo*, 2006:1–10.

Cao, H.-L., Gómez Esteban, P., Albert, D. B., Simut, R., Van de Perre, G., Lefeber, D., and Vanderborght, B. (2017). A collaborative homeostatic-based behavior controller for social robots in human–robot interaction experiments. *International Journal of Social Robotics*, 9(5):675–690.

Cos, I., Cañamero, L., Hayes, G. M., and Gillies, A. (2013). Hedonic value: Enhancing adaptation for motivated agents. *Adaptive Behavior*, 21(6):465–483.

Goodrich, M. A., Schultz, A. C., et al. (2008). Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction*, 1(3):203–275.

Hellou, M., Vinanzi, S., and Cangelosi, A. (2023). Bayesian theory of mind for false belief understanding in human-robot interaction. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1893–1900. IEEE.

Hellou, M., Vinanzi, S., and Cangelosi, A. (2024). Where is my favourite toy? inferring the mental states of users in false belief understanding. In *2024 IEEE International Conference on Development and Learning (ICDL)*, pages 1–8. IEEE.

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

Konidaris, G. and Barto, A. (2006). An adaptive robot motivational system. In *International Conference on Simulation of Adaptive Behavior*, pages 346–356. Springer.

Lewis, M. and Canamero, L. (2016). Hedonic quality or reward? a study of basic pleasure in homeostasis and decision making of a motivated autonomous robot. *Adaptive Behavior*, 24(5):267–291.

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

Romeo, M., McKenna, P. E., Robb, D. A., Rajendran, G., Nesset, B., Cangelosi, A., and Hastie, H. (2022). Exploring theory of mind for human-robot collaboration. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 461–468. IEEE.

Salichs, M. A. and Malfaz, M. (2011). A new approach to modeling emotions and their use on a decision-making system for artificial agents. *IEEE Transactions on affective computing*, 3(1):56–68.

Scassellati, B. (2002). Theory of mind for a humanoid robot. *Autonomous Robots*, 12:13–24.

Vinanzi, S., Patacchiola, M., Chella, A., and Cangelosi, A. (2019). Would a robot trust you? developmental robotics model of trust and theory of mind. *Philosophical Transactions of the Royal Society B*, 374(1771):20180032.

Yu, C., Serhan, B., and Cangelosi, A. (2024). Top-tom: Trust-aware robot policy with theory of mind. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7888–7894. IEEE.