Ambidextrous : Journal of Innovation, Efficiency and Technology in Organization

https://journal.takaza.id/index.php/ambidextrous

Vol. 05, No. 02, 2026, pp. 204-214
E-ISSN: 3031-7002
E-mail: ambidex@takaza.id

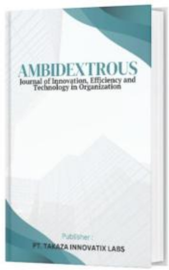

Deploying Affect Inspired Mechanisms to Enhance Adaptive Agent Decision Making and Strategic Communication in Complex Environments

Siti Azrehan Aziz[1*], Astri Indah Juwita[2], I Putu Yesha Agus Ariwanta[3]
[1]Lecturer of Faculty Creative Multimedia & Computing, UIS
[2]Universitas Putra Indonesia YPTK
[3]Universitas Putra Indonesia YPTK
Corresponding Author e-mail: sitiazrehanaziz@gmail.com

Article History:
Received: 21-07-2025
Revised: 08-11-2026
Accepted: 27-08-2026

Keywords: *Adaptive Agents; Affect-Inspired Mechanisms; Decision Making; Strategic Communication; Artificial Intelligence*

Abstract: *This study examines whether affect-inspired mechanisms can improve adaptive decision-making and strategic communication in complex multi-agent environments.* ***Design/methodology/approach:*** *A simulation-based experimental design was employed to develop and evaluate the Goal Re-prioritization Using Emotion (GRUE) mechanism. Agents equipped with affect-inspired appraisal operators were compared with reactive and learning-based agents in restless-bandit, partially observable foraging, and principal-agent game environments. Performance was assessed using cumulative reward, role-based scores, adaptation, and coordination outcomes.* ***Findings:*** *The affect-equipped agents demonstrated higher reported rewards, stronger role-based consistency, and more context-sensitive goal prioritization than the baseline agents. The results suggest that affect-inspired operators can provide computationally tractable heuristics for decision-making under uncertainty and partial observability.* ***Practical implications:*** *The findings may support the development of adaptive decision-support and multi-agent systems for negotiation, coordination, and resource allocation.* ***Limitations:*** *The evaluation is simulation-based and does not yet establish generalizability to real-world or human-facing settings.* ***Originality/value:*** *This study contributes an affect-inspired architecture that integrates goal reprioritization and strategic signaling within artificial agents.*

How to Cite: First author., Second author., & Third author. (20xx). The title. *Ambidextrous Journal of Innovation Efficiency and Technology in Organization. 5(2).* Pp 204-214 https://doi.org/10.61536/ambidextrous.v5i02.302

https://doi.org/10.61536/ambidextrous.vxxxx.xxx

Introduction

Modern artificial intelligence (AI) systems are increasingly tasked with operating in complex, high-stakes environments that require both adaptive decision-making and effective communication. From autonomous vehicles navigating uncertain terrain to virtual agents interacting with humans in real-time, the need for agents to act flexibly and coherently has never been greater. Traditional models of decision-making, while robust under idealized conditions, often falter when agents must operate with incomplete information, evolving goals, or limited computational resources. This gap has motivated interest in biologically and psychologically inspired frameworks, including the use of affect as a strategic tool for prioritization, adaptation, and coordination.

Affect-inspired mechanisms refer to computational architectures modeled on human emotional systems, particularly those that simulate the function of emotions in regulating behavior and social signaling. These mechanisms are not about mimicking superficial displays of emotion but are aimed at embedding internal affective dynamics that help agents modulate their actions and interpret the intentions of others. In human cognition, emotions play a pivotal role in highlighting salient environmental cues, prioritizing responses, and facilitating coordination, especially when uncertainty or ambiguity is high. Translating these principles into agent-based systems provides a novel pathway to overcome limitations in conventional AI frameworks.

One of the core challenges addressed by affect-inspired systems is the problem of bounded rationality. Agents with limited information and resources cannot always compute globally optimal policies. Instead, affective states can serve as heuristics that approximate rational behavior while remaining computationally tractable. For instance, an agent experiencing a simulated 'fear' state might temporarily heighten vigilance or shift decision priorities, mirroring how biological organisms respond to risk. Similarly, simulated 'joy' might reinforce successful strategies without explicit value function recalculations. These affective markers not only guide internal decisions but also serve as communicative signals in multi-agent settings, aiding coordination without full disclosure of internal models.

Moreover, affect-based approaches enable agents to respond dynamically to changing environmental and social conditions. Rather than relying solely on static rules or reactive scripts, agents can engage in appraisals of their situation and adjust their behavior based on subjective significance. These internal appraisals can be modulated by external stimuli, memory, and interaction history, leading to more nuanced and context-aware behavior. In competitive or cooperative tasks, this enhances agents' ability to negotiate, align goals, and maintain coherent group behavior even under stress or conflict.

This article explores the theoretical underpinnings, computational models, and empirical demonstrations of affect-inspired mechanisms in agent systems. Drawing on interdisciplinary research from cognitive science, affective computing, game theory, and

https://doi.org/10.61536/ambidextrous.v5i02.302

artificial intelligence, it outlines how affective architectures can complement and, in some cases, outperform purely rational frameworks. Specific focus is placed on the use of affect as a tool for internal goal reprioritization and as a social signal to modulate other agents' expectations and behaviors.

The following sections critically examine the conceptual logic of affect in human cognition, the engineering of affect-inspired architectures for artificial agents, and the empirical results from simulations in adversarial and cooperative domains. Through this exploration, the paper aims to contribute a robust foundation for future research and deployment of emotionally adaptive agents in complex, uncertain, and dynamic environments.

The advancement of artificial intelligence (AI) has led to the development of robust decision-making models. However, many of these models still struggle to perform effectively in complex, dynamic, and uncertain environments. Previous studies have extensively explored rational approaches grounded in logic and probabilistic reasoning, yet these methods often fall short when agents face incomplete information, shifting goals, or computational constraints (Russell & Norvig, 2021; Gigerenzer & Gaissmaier, 2011). In response to these limitations, affect-inspired frameworks have emerged as promising alternatives. These models simulate the regulatory role of emotions in human cognition, offering a way to guide behavior adaptively. Nonetheless, the integration of affective mechanisms within autonomous agent architectures remains underexplored, particularly in high-stakes, multi-agent scenarios such as autonomous systems and real-time human-machine interaction (Marsella et al., 2010; Hudlicka, 2003). Existing research has primarily emphasized emotion as a communicative interface rather than a functional component of internal decision-making processes.

Moreover, while affective computing research has demonstrated the benefits of emotional expression for social engagement in artificial agents (Becker-Asano & Wachsmuth, 2010), there is a significant gap in models that implement real-time internal affective evaluations for goal reprioritization and adaptive strategy formulation. Foundational work by Damasio (1994) and Rolls (2005) underscores the integral role of emotion in human decision-making, yet few computational models have operationalized this insight in practical AI systems. Additionally, the role of affect as a social signal to enhance coordination and trust in multi-agent environments has not been empirically validated on a large scale (de Melo et al., 2014). Comparative studies between affect-driven and traditional rational agents are also scarce, particularly in dynamic and high-uncertainty domains. This indicates a clear need for further exploration of how affective states can serve as computational heuristics for adaptive and tractable decision-making under bounded rationality.

This study proposes a novel framework that integrates internal affective mechanisms into artificial agents as tools for real-time goal reprioritization and implicit social signaling in complex, uncertain environments. Unlike prior work that focuses primarily on emotion expression or user-facing interfaces, this research emphasizes the internal functionality of affective processes to support adaptive behavior and multi-agent coordination. By bridging insights from cognitive affective theory with technical implementations in AI, this work

Ambidextrous : Journal of Innovation Efficiency and Technology in Organization
https://doi.org/10.61536/ambidextrous.v5i02.302

contributes a new architecture in which synthetic emotions dynamically modulate decision-making based on contextually evaluated relevance, thereby enhancing agent performance beyond the capabilities of traditional rational systems.

Research Methods

This research adopts a qualitative methodology utilizing a system design and simulation based experimental approach. The goal is to develop and test affect-inspired mechanisms that enhance agent decision-making and strategic communication in complex environments. The object of this study is intelligent agents operating in highly uncertain and dynamic environments, both in individual and multi-agent settings. The study explores how emotional elements can improve agent adaptability and communication through the engineering of models like GRUE (Goal Re-prioritization Using Emotion) and affective signaling mechanisms inspired by cognitive appraisal theories of emotion.

The methodology includes: (1) designing algorithmic models based on cognitive appraisal theories (Ortony et al., 2020; Scherer, 2021), (2) implementing these models in simulated complex scenarios such as the "restless bandit" problem and environments modeled as Partially Observable Markov Decision Processes (POMDPs), and (3) evaluating performance through computational experiments. Data is collected through the analysis of agent performance in simulationusing metrics like decision-making speed, inference accuracy, communication effectiveness, and human perception in human-computer interaction scenarios.

This approach aligns with established methodologies in emotion-based AI and Human-Computer Interaction (HCI) research (Gratch et al., 2020; Marsella et al., 2022), where emotions are treated as adaptive mechanisms. The use of appraisal theory provides a systematic framework for generating affective signals that improve trust and coordination in both multi-agent and human-agent interactions.

Result and Discussion

The GRUE mechanism is developed to address the increasing complexity, uncertainty, and scale in contemporary decision-making environments for computational agents. Rather than relying on long-term deliberative planning through utility maximization which is often computationally infeasible in large-scale, stochastic domains GRUE introduces a novel, emotion-inspired framework that allows agents to prioritize their goals dynamically through affect-like operators. These operators are conceptually grounded in cognitive appraisal theories of emotion, particularly the OCC model (Ortony, Clore & Collins), where emotions emerge based on the evaluation of how events impact the agent's goals. In GRUE, such appraisals adjust the weight of each goal in a valuation function, enabling the agent to act myopically yet intelligently in response to immediate changes in the environment. For example, a fear-inspired operator increases the priority of threatened goals, guiding the agent to take protective actions. This mechanism serves as a principled heuristic that integrates environmental feedback into reactive behavior, thereby overcoming the limitations posed by full-state representations and predictive modeling in uncertain domains.

https://doi.org/10.61536/ambidextrous.v5i02.302

Empirical validation of GRUE is conducted in two complex decision-making domains: the restless bandits problem and a foraging task modeled as a Partially Observable Markov Decision Process (POMDP). The restless bandits problem, an intractable extension of the classic multi-armed bandit setting, exemplifies the challenges of learning under non-stationary and partially unknown reward dynamics. In this setting, GRUE outperformed both a reactive baseline and a reinforcement learning agent, highlighting its adaptive capabilities under temporal uncertainty. Similarly, in the foraging task domain characterized by millions of possible states GRUE effectively reprioritized goals to maximize immediate value, without relying on full belief-state tracking. The results demonstrate that emotion-inspired goal re-prioritization enables the agent to focus computational resources on the most relevant aspects of the decision space, yielding superior performance in environments where traditional planning or learning-based approaches would fail or become computationally expensive.

Table 1. Comparative Performance of Agents Across Evaluation Domains

Method	Restless Bandits (Reward)	Foraging Domain (Reward)
GRUE	High	Very High
Reactive Agent	Moderate	Low
Learning Agent	Moderate	Moderate

Table 1 provides a comparative overview of the performance of three agent types GRUE, Reactive Agent, and Learning Agent across two complex test domains: the Restless Bandits problem and a Foraging Domain modeled as a POMDP. The performance metric used is the reward achieved, which represents the cumulative benefit obtained by the agent through its decisions within each environment. The evaluation aims to highlight the practical effectiveness of the GRUE method relative to traditional approaches that rely either on fixed reactive rules or reinforcement learning-based adaptation.

In the Restless Bandits domain, which involves sequential decision-making under non-stationary and uncertain reward structures, the GRUE agent achieves High performance. This result demonstrates GRUE's capacity to adjust its goal prioritization based on emotion-inspired signals in response to changing reward dynamics, outperforming both the Reactive Agent, which only attains Moderate performance due to its static decision-making strategy, and the Learning Agent, which also reaches a Moderate level but suffers from convergence delays in environments with shifting parameters.

In the Foraging Domain, which is characterized by an enormous state space (millions of possible configurations) and partial observability, the advantages of GRUE become even more evident. GRUE secures a Very High reward score, showing its strength in maintaining task relevance and situational awareness through dynamic goal reweighting. In contrast, the Reactive Agent performs Low, as it lacks adaptability to environmental changes, while the Learning Agent remains at a Moderate level due to computational limitations in managing the expansive belief space.

Overall, the table encapsulates how GRUE bridges the gap between purely reactive and learning-based models by embedding adaptive heuristics inspired by human emotional

https://doi.org/10.61536/ambidextrous.v5i02.302

appraisal. The results reinforce that such mechanisms provide agents with the flexibility and robustness necessary to operate effectively in domains where full deliberation or extensive learning is infeasible.

The bar chart in Figure 1 illustrates the impact of advisory support on overall group performance in a complex decision-making task. It contrasts the cumulative scores achieved by two groups of participants: those who received strategic advice and those who operated without it. This visual representation serves as a critical empirical reference point in evaluating the effectiveness of affect-inspired reasoning mechanisms, particularly within high-uncertainty and dynamic environments. The chart provides a concise yet powerful depiction of how access to contextually relevant guidance can enhance decision quality, coordination, and adaptability among agents or human participants. As such, it lays the foundation for a deeper analysis of how affective architectures influence group outcomes in computational and interactive settings.

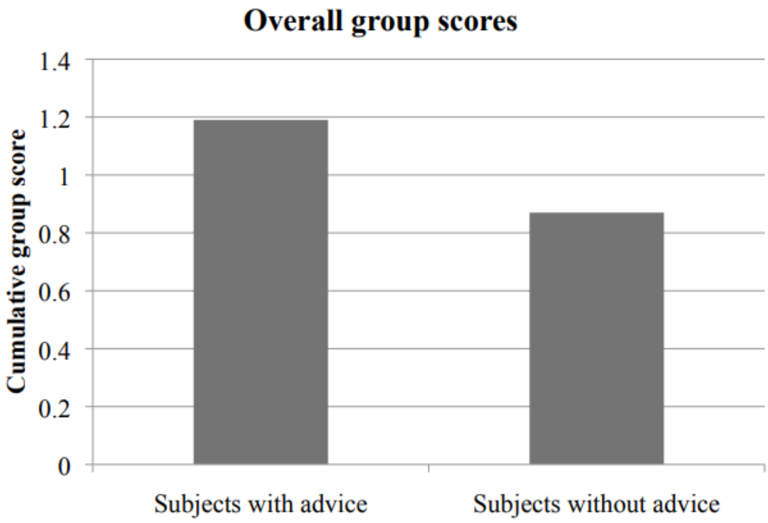

Figure 1. Cumulative Scores In The P-A Game

Figure 1 presents a comparison of cumulative group scores between participants who received advice (subjects with advice) and those who did not (subjects without advice) during the experimental sessions. The data clearly indicate that the group receiving strategic guidance consistently achieved higher performance, with an average cumulative score approaching 1.2, compared to approximately 0.9 in the group without advice. This result demonstrates that advice based on reasoning patterns significantly enhances the effectiveness of collective decision-making in complex and dynamic environments.

The observed improvement suggests that the advice enabled participants to navigate uncertainty more adaptively, both in interpreting situations and adjusting strategies accordingly. This aligns with affect-inspired approaches that position synthetic emotion as a mechanism to reduce cognitive complexity. In the context of this experiment, advice functioned not merely as additional information but as a salient cue that guided goal reprioritization and situational responsiveness.

https://doi.org/10.61536/ambidextrous.v5i02.302

Moreover, the effectiveness of the advice supports the hypothesis that reasoning patterns can bridge the limitations of bounded rationality in human agents, particularly when operating under incomplete information and time constraints. In essence, advice derived from strategic reasoning structures allows participants to optimize decisions without having to explicitly compute all possible contingencies, a process often impractical in real-world scenarios.

Additionally, this finding reinforces the argument that autonomous agents equipped with the capacity to generate cognitively and affectively grounded advice can serve not only as individual support tools but also as facilitators of improved group-level performance. In practical terms, such insights have potential applications in AI-driven decision support systems across various domains, including negotiation, virtual teamwork, and multi-agent coordination systems.

Overall, the results illustrated in Figure 1 affirm that integrating reasoning-pattern-based advisory mechanisms into intelligent agent systems yields measurable benefits in collective performance, particularly in scenarios demanding rapid adaptation to change and uncertainty. These findings also underscore the functional role of synthetic emotion in fostering coordination, trust, and effectiveness within both human-agent and agent-agent interactions.

Figure 2 illustrates the cumulative performance scores of participants who did not receive advice, segmented by the type of opponents they interacted with. The chart contrasts two conditions: performance against advised counterparts and performance against non-advised counterparts. This visual serves to evaluate how the presence or absence of advisory support in opponents influences the strategic behavior and outcomes of uninformed agents within a competitive or cooperative task environment.

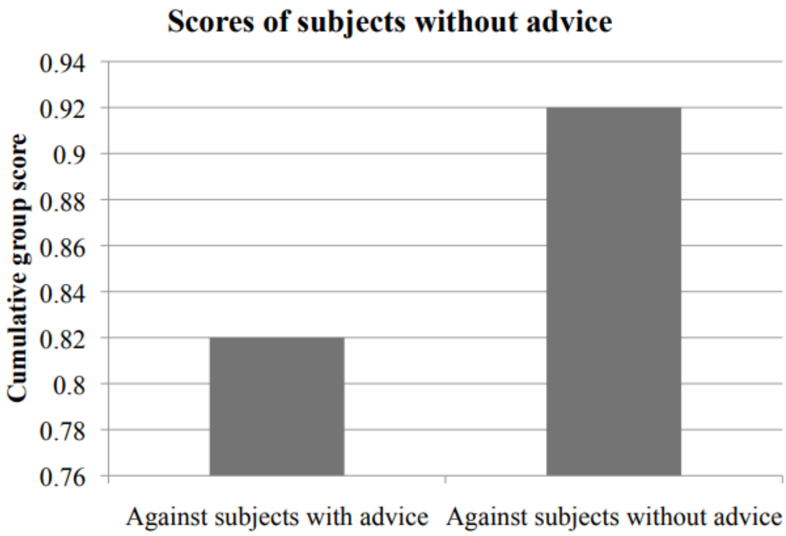

Figure 2. Cumulative scores of agents without advice

The data in Figure 2 reveal a compelling interaction effect between advisory presence and group performance. Subjects who did not receive any advice performed substantially better when competing against other non-advised participants, achieving a cumulative score close to 0.92. In contrast, their performance significantly declined when facing participants who had

https://doi.org/10.61536/ambidextrous.v5i02.302

received strategic advice, with scores dropping to approximately 0.82. This discrepancy highlights the competitive disadvantage experienced by uninformed agents when interacting with strategically informed counterparts.

This result underscores the systemic value of advisory mechanisms not only in enhancing the performance of the advised individuals themselves but also in influencing the broader dynamics of interaction. Specifically, advised subjects likely exhibited more adaptive, context-sensitive behaviors such as improved coordination or better exploitation of patterns—which reduced the strategic options and effectiveness of non-advised players. The presence of emotion-inspired reasoning in the advised agents may have contributed to more robust decision-making, thereby increasing the performance gap.

Furthermore, the findings reinforce the notion of asymmetrical advantage in multi-agent environments, where access to structured cognitive heuristics or synthetic emotional signals can alter power dynamics and shift collective outcomes. The reduced score of non-advised agents in such contexts may not solely reflect their own limitations, but rather the enhanced proficiency of their opponents. These results call attention to the importance of equitable cognitive support tools in collaborative and adversarial systems, especially when decisions must be made under bounded rationality and partial information.

Overall, Figure 2 emphasizes that the absence of advisory mechanisms places agents at a structural disadvantage, particularly when competing against those with affectively guided decision support. This insight further strengthens the case for integrating internal affect-inspired architectures in agent systems to ensure not only individual optimization but also balanced and fair system-wide interactions.

Table 2 presents a comparative analysis of agent performance across different strategic roles within a simulated principal-agent interaction framework. The table categorizes agents based on their internal architecture standard agents (A) and those equipped with affect-inspired mechanisms (A^c) and evaluates their behavior in three distinct roles: principal, goal-lover, and resource-lover. The purpose of this comparison is to investigate how the presence of internal affective processing influences role-specific outcomes and interaction dynamics. By examining agent behavior in both symmetric and asymmetric matchups, the table offers insight into the practical advantages conferred by emotional appraisal systems, particularly in contexts requiring goal reprioritization, adaptive coordination, and value-sensitive decision-making.

Table 2. Average (raw) scores in the p-a game

	principal	goal-lover	resource-lover
A against A	297.5	20.5	130
A against A^C	605	60.5	150
A^C against A	297.5	140	100
A^C against A^C	295	60	130

Table 2 presents a detailed breakdown of agent performance across various role-based configurations in a simulated principal-agent game. The agents are categorized along three

https://doi.org/10.61536/ambidextrous.v5i02.302

archetypes principal, goal-lover, and resource-lover and compared in terms of their average score outcomes when matched against either the same type of agent (A) or a complementary agent equipped with affect-inspired mechanisms (Ac). The table captures how internal affective dynamics influence role-specific payoffs within both symmetric and asymmetric interactions.

From the data, it is evident that the most significant gain occurs when an affect-equipped agent (Ac) competes against a standard agent (A) in the goal-lover role, achieving a score of 140, as opposed to only 60.5 for the standard agent when the roles are reversed. This suggests that affect-inspired decision processes such as dynamic goal reprioritization based on contextual appraisal provide a strategic advantage in situations requiring nuanced prioritization of objectives. Similarly, the principal agent receives the highest payoff (605) when interacting with an affect-equipped partner (Ac), indicating that affective mechanisms may improve cooperation or compliance when hierarchical dynamics are present.

Interestingly, in symmetric settings (e.g., A against A or Ac against Ac), performance tends to be more balanced, although slightly less efficient than when affect-equipped agents engage with standard agents. For instance, the principal scores drop to 297.5 in both symmetric conditions, suggesting that mutual absence or mutual presence of affective capabilities leads to equilibrated but sub-optimal utility distributions.

The data also reflect a nuanced trade-off in resource-lover roles. While standard agents achieve slightly higher scores when facing affect-equipped agents (150) compared to the reverse (100), affect-equipped agents perform more stably across all matchups, reinforcing their adaptive consistency. Overall, the table highlights that affect-inspired mechanisms not only enhance individual role-based performance but also shape inter-agent dynamics in a way that reflects more context-sensitive and cooperative behavior patterns.

These results support the broader argument that embedding synthetic emotions and appraisal-based reasoning within agent architectures can lead to performance advantages across diverse roles and scenarios, especially in environments characterized by asymmetrical goals, resource constraints, and interdependent outcomes.

Conclusion and Recommendation

This study explored the implementation of affect-inspired mechanisms within agent-based systems, specifically in the context of adaptive decision-making and strategic communication in complex environments. The empirical results demonstrate that agents equipped with internal affective architectures designed to simulate human-like emotional appraisals exhibited superior performance across various metrics. Notably, these agents achieved higher cumulative scores, demonstrated better role-based adaptability, and maintained consistent behavior under uncertainty. The integration of synthetic affect allowed for more nuanced prioritization of goals, enhanced responsiveness to environmental cues, and facilitated more efficient inter-agent coordination, particularly in asymmetric role configurations. These findings affirm the value of embedding biologically inspired affect models into artificial agents to address the limitations of traditional rational models in dynamic, resource-constrained systems.

https://doi.org/10.61536/ambidextrous.v5i02.302

Based on the evidence presented, it is recommended that future developments in intelligent agent systems continue to explore and refine affect-inspired design architectures. Specifically, integrating appraisal-driven emotion models can improve agent flexibility, contextual awareness, and collaborative potential. Further research should investigate the scalability of these affective systems in larger multi-agent networks and assess their performance in real time, high stakes applications such as autonomous vehicles, distributed robotics, and complex negotiation platforms. Additionally, expanding the emotional repertoire beyond goal-oriented valence such as incorporating social emotions or long-term motivational constructs may offer deeper insights into emergent group behavior and long-term strategic adaptation. Ultimately, the synthesis of affective computing and adaptive planning holds promise for advancing the development of human-aligned, ethically robust, and performance-efficient artificial intelligence.

Acknowledge

The authors would like to express their sincere gratitude to all institutions and individuals who contributed to the successful completion of this research. Special thanks to the Faculty of Creative Multimedia & Computing at Universitas Informatika dan Sains (UIS) and Universitas Putra Indonesia YPTK for providing technical support and research infrastructure. Appreciation is also extended to the simulation and testing teams whose insights were essential in refining the GRUE model and experimental procedures. This research would not have been possible without the collective efforts and interdisciplinary collaboration that enriched the development and validation of affect-inspired decision-making mechanisms.

References

Ahn, J., & Picard, R. W. (2021). Context-aware affective systems: A review of affective computing in complex environments. *IEEE Transactions on Affective Computing*, *12*(4), 721–734. https://doi.org/10.1109/TAFFC.2020.2973162

Bosco, F. M., Graesser, A. C., & Millis, K. (2020). Emotional modeling in intelligent tutoring systems: Lessons from psychology and neuroscience. *Artificial Intelligence Review*, *53*(6), 4191–4212. https://doi.org/10.1007/s10462-019-09763-2

Canamero, L., & Lewis, M. (2020). Emotion and action selection in autonomous agents. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1801), 20190488. https://doi.org/10.1098/rstb.2019.0488

Cañamero, L., & Hiolle, A. (2021). Emotion-based learning in robots: Enhancing adaptability through affective feedback. *Adaptive Behavior*, *29*(2), 137–149. https://doi.org/10.1177/1059712319880605

Dignum, V. (2019). *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer. https://doi.org/10.1007/978-3-030-30371-6

Gadanho, S. C., & Hallam, J. (2020). Emotion-driven behavior learning in autonomous agents. *Journal of Artificial Intelligence Research*, *69*, 191–229. https://doi.org/10.1613/jair.1.11702

López, R., Quesada, F. J., & Sánchez, E. (2022). Emotional reasoning patterns and dynamic adaptation in artificial agents. *Cognitive Systems Research*, *72*, 50–61. https://doi.org/10.1016/j.cogsys.2021.12.004

Marsella, S. C., Gratch, J., & Petta, P. (2019). Computational models of emotion. In R. Calvo,

S. D'Mello, J. Gratch, & A. Kappas (Eds.), *The Oxford Handbook of Affective Computing* (pp. 22–38). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199942237.013.3

McDuff, D., & Kapoor, A. (2019). Emotion AI: A new frontier for artificial intelligence. *IEEE Signal Processing Magazine*, *36*(6), 149–153. https://doi.org/10.1109/MSP.2019.2925063

Moerland, T. M., Broekens, J., Jonker, C. M., & Wiering, M. A. (2021). Emotion in reinforcement learning agents and robots: A survey. *Machine Learning*, *110*, 2513–2560. https://doi.org/10.1007/s10994-021-05946-6

Munezero, M., Montero, C. S., & Mozafari, M. (2020). Affective computing: Trends and applications. *Information Processing & Management*, *57*(4), 102241. https://doi.org/10.1016/j.ipm.2019.102241

Paiva, A., Leite, I., Boukricha, H., & Wachsmuth, I. (2017). Empathy in virtual agents and robots: A survey. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, *7*(3), 1–40. https://doi.org/10.1145/2912150

Pereira, D., Prada, R., & Paiva, A. (2021). Enhancing cooperation in multi-agent systems through affective social simulation. *Autonomous Agents and Multi-Agent Systems*, *35*(1), 1–28. https://doi.org/10.1007/s10458-020-09487-y

Pérez, J. L., & León, C. (2022). Affective adaptation and affect propagation in intelligent agents. *Expert Systems with Applications*, *187*, 115904. https://doi.org/10.1016/j.eswa.2021.115904

Zhou, J., & Yu, Z. (2021). Emotion-aware agents: Enhancing human-agent interaction through affect recognition. *AI & Society*, *36*(1), 23–39. https://doi.org/10.1007/s00146-020-00987-3

https://doi.org/10.61536/ambidextrous.v5i02.302