My Body, My Perceptions: A Shift from Computationalism to Embodied Cognition in BDI-agent-based Embedded Systems

Nilson Lazarin
Federal Fluminense University (UFF)
Niterói, Brazil
Federal Center for Technological Education Celso Suckow da Fonseca (Cefet/RJ) - Rio de Janeiro, Brazil
nilson.lazarin@cefet-rj.br

Carlos Pantoja
Federal Center for Technological Education Celso Suckow da Fonseca (Cefet/RJ) - Rio de Janeiro, Brazil
Linköping University (LiU)
Linköping, Sweden
carlos.pantoja@cefet-rj.br

Jose Viterbo
Federal Fluminense University (UFF)
Niterói, Brazil
viterbo@uff.br

ABSTRACT

Autonomous agents perceive and act in an environment and can also be part of a Multi-agent System (MAS), cooperating or competing with one another to achieve the system goals. When embodied in physical devices, this entity is given a body composed of sensors and actuators to manifest itself in the real-world environment without remote control or external processing. Currently, the approach used to build embedded agent-based systems is based on the construction of conventional MAS, rooted in the Computational Metaphor, a theory in cognitive science that views mental processes as computational processes, with the brain functioning as the center of cognition, disregarding the physical body of the agent. Inspired by Embodied Cognition — another theory in cognitive science that redefines the computational metaphor, emphasizing the importance of the physical body in cognitive abilities — this paper discusses the need to expand the representation of the agent body when embedded in hardware platforms. In addition, we propose expanding the Belief-Desire-Intention (BDI) agents' perceptions by considering interoception (internal body sensations), exteroception (external environmental stimuli), and proprioception (awareness of the body's position and movement). Finally, to evaluate this proposal, a scenario using the Multi-agent Oriented Programming (MAOP) paradigm is proposed and assessed.

KEYWORDS

BDI model; Multi-Agent Oriented Programming; Embedded MAS

ACM Reference Format:
Nilson Lazarin, Carlos Pantoja, and Jose Viterbo. 2026. My Body, My Perceptions: A Shift from Computationalism to Embodied Cognition in BDI-agent-based Embedded Systems. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 10 pages. https://doi.org/10.65109/QIVX3835

1 INTRODUCTION

The Belief-Desire-Intention (BDI) model [14] is one of the most widely used models for constructing cognitive agents and architectures in autonomous systems [25]. In the past, they were applied in purely software-based [63, 87] or augmented reality applications [27, 29] due to the technological limitations of the 1990s. In the early 2000s, they could act indirectly in the physical world using the remote control of robots or cyber-physical systems [24, 57]. Since the 2010s, with the advent of Single-Board Computers (SBC) (e.g., Raspberry Pi), a whole Multi-Agent System (MAS) can be embedded in a hardware platform, and its agents were finally situated in the real-world environment [47, 50, 61, 67].

The approach used to develop a BDI-agent-based Embedded System is strongly based on the construction of conventional MAS, utilizing the *computational metaphor* — a theory from traditional cognitive science where mental processes in the brain are the core of cognition [81]. However, the *status quo* does not consider the *physical reality* of the agent's body situated in the real-world environment, as can be verified by the absence of a body representation in the dimensions provided by the Multi-Agent Oriented Programming (MAOP) paradigm [9]: the agent dimension, where individual autonomous entities exist; the environmental dimension, a computational representation of artifacts from the environment where agents are acting; and the organization dimension, where agents are organized.

From a dualistic philosophical point of view, a body can be understood as the prison of the entity or even the physical part at its service [60]. In contrast, from an anatomical perspective, a body can be understood as a set of different apparatuses (e.g., respiratory, digestive, urinary) composed of distinct elements (e.g., lungs, liver, kidneys) that constitute the entity [88]. In other words, the body, from the point of view of the entity that governs it, is not an artifact of the environment but rather the entity itself.

This demonstrates a *Theoretical* and *Practical-knowledge* gap – when an absence of a model exists, and the practice differs from what the theory proposes [19]. Once there is no agent's body representation in the MAOP paradigm, the designer needs to represent the agent's body as an environmental artifact or use implicit representations. Moreover, this representation makes it challenging to develop agent-based embedded systems in areas such as robotics, where physical interaction with the environment is required. In addition, the absence of an explicit distinction between bodily and environmental percepts may lead to ambiguities in the agent's deliberation process, particularly in embedded domains where internal and external signals often share the same structure. A more expressive representation of these distinct sources of information may help agents reason more clearly about their own physical state, reducing potential confusion and contributing to more reliable

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/QIVX3835

decision-making. Thus, this research relies on Embodied Cognition [81], a theory from cognitive science, to better represent the agents' bodies, guided by two questions: 1) *How can a cognitive agent perceive its own body?* 2) *Is it possible to extend the MAOP to represent a notion of the agent's body?*

This paper proposes a novel model to resignify the body representation in the agents' minds, shifting from traditional *computationalism* to *embodied cognition*. Following the Design Science Research (DSR) methodology [26], we present the novel $MAOP^{+b}$ model, which extends the MAOP to incorporate a new *construct* named *mechatropsychosocial*, adding a body representation for the agent. Now, the agent can perceive its body from the perspective of embodied cognition by using three constructs: *interoceptions* (perceptions within the body), *exteroceptions* (perceptions outside the body), and *proprioceptions* (kinesthesia, which enables the perception of the body's movement and location). Lastly, the model and *constructs* are instantiated in a JaCaMo [10] extension named $JACAMO^{+b}$.

Then, we conducted a Descriptive Evaluation with an Illustrative Scenario. The scenario comprises a BDI agent-based system embedded in a robotic vacuum cleaner, accountable for autonomously cleaning a campus. The evaluation compares the agent's capability to perceive its own body with related works selected in a mapping review, considering full compliance with the MAOP paradigm, the agent's body representation, and the agent's bodily perceptions. As a result, our model demonstrated that it is capable of representing the descriptions of different agents' bodies, whereas the others do not do so explicitly or at all. The body is represented in the agent's mind based on specific beliefs, discerning them from perceptions of the environment. In this way, expressiveness is added to the knowledge held by the agent, facilitating decision-making that also considers bodily characteristics. Furthermore, the proposed scenario demonstrates that state-of-the-art works fail when the agent must make a decision involving conflicting environmental and bodily beliefs.

In summary, the contributions of this paper are: *(i)* new *constructs* to a better understanding of the embodied cognition applied to embedded BDI-based multi-agent system: the definition of the *mechatropsychosocial* entity, concerned with the mechatronical, psychological, and social aspects of autonomous entities; besides that, three sources of perceptions in the cognitive process: *interoceptions*, *exteroceptions*, and *proprioceptions*; *(ii)* a *model*, named $MAOP^{+b}$, that extends the MAOP, including the complementary notion of body, allowing an agent to perceive its body considering those new sources; *(iii)* an *instantiation*, named $JACAMO^{+b}$ that complies with the notions of agents, bodies, workspaces, organizations, to program a MAS using the agent, environmental, and organizational dimensions.

Finally, this paper is organized as follows: in Section 2, we present some background about MAOP; in Section 3, we present some related works considering the body of an agent or a MAS; in Section 4, we present our model of body and percepts for systems with a body; in Section 5 we present a evaluation using our proposed model and related works, in Section 5.1 we provide instructions to guarantee the reproducibility of this research; and finally in Section 6 the conclusions are presented.

2 BACKGROUND

Some approaches have explored how an agent's body or sensorimotor capabilities could influence cognition. Urban and Schmidt [91, 92] proposed the PECS model to integrate physical, emotional, cognitive, and social factors, but its "physical" dimension captured only internal physiological states rather than a concrete body interacting with the world. Vernon [94] argued that cognition emerges from autonomous sensorimotor coupling, which later inspired operational implementations such as the iCub humanoid robot [56]; however, neither the conceptual principles of Enaction nor the architecture embodied in iCub offered a formal or generalizable model of an agent's body within a cognitive architecture. Additionally, Beetz et al. [7] developed the CRAM framework for cognition-enabled manipulation; however, despite its tight coupling with low-level sensorimotor data, it did not provide a formal or declarative representation of the agent's body as part of a cognitive architecture.

Later, Boissier et al. [9] proposed the MAOP paradigm, the first successful integration of Agent-Oriented Programming (AOP), Organization-Oriented Programming (OOP), and Environmental-Oriented Programming (EOP), thereby unifying cognitive agents, modular artifacts, and organizational structures within a high-level programming model where agents are the primary abstractions of the system, organized along three orthogonal dimensions that together define the structure and dynamics of MAS.

The *organizational dimension* is the level at which agents are structured by organizational specifications, typically grounded in the Moise model [39], which provides a normative and structural framework that supports coordination and dynamic reorganization of the system at runtime [9, 41]. The *agent dimension*, in turn, is the level at which individual autonomous entities exist, and it is typically materialized by adopting a BDI-based AOP framework, such as Jason [9, 11, 12]. The relationship between the agent and organization dimensions can be understood as analogous to Durkheim's notion of *social cohesion* – where individuals internalize shared norms and rules into their consciousness, which permeate the group's way of being and acting, leading to the emergence of collective behavior [58]. An example of this analogy is presented by Hafiene et al. [38], who use constitutive and regulative norms to govern interactions among BDI agents on the Reddit social network, thereby ensuring trust and facilitating structured collaboration.

The *environmental dimension*, lastly, provides a computational representation of the environment where agents and resources are situated. Following the Agents and Artifacts (A&A) [77] model, this dimension is organized into *workspaces* and populated with *artifacts*. *Workspaces* are logical places, containers that store a dynamic collection of artifacts and agents, used to define the situational topology, establishing a scope for interactions, event observability, and the activities carried out by agents [8, 10]. *Artifacts*, in turn, are useful for modularizing each environment, they: are basic "computational bricks" that allow the representation of any elements that can be perceived or used by the agents [8]; support both individual and collective activities of agents [74, 75]; and allows interaction not only with internal computational entities but also with real-world components, such as sensors, devices, or web services [10].

Boissier et al. [9] describe agents as situated in the environment. In *Situatedness*, agents are not isolated reasoning entities:

their beliefs are dynamically updated by observable properties of artifacts, and their actions are executed as operations, thereby coupling decision-making with ongoing changes in the environment [9, 10, 74]. However, *Situatedness* differs from *Embodiment*. For example, Ricci et al. [76] distinguish the environment into *exogenous*, conceived as external physical or computational contexts, and *endogenous*, which are explicitly designed as part of the MAS itself. While *Situatedness* defines the informational and interactive coupling between agent and environment, *Embodiment* considers the agent's physical body, morphology, and sensorimotor apparatus as constitutive elements of intelligence, as advocated by the literature on Embodied AI [17, 42, 99]. So, the MAOP relationship between the agent and environment dimensions, when considering actions in the physical world (exogenous environment), can be metaphorically compared to phenomena described by Kardec, in which discarnate spirits from another dimension manipulate physical artifacts [2]. An example is the Chameleon Agent Architecture [28], in which a BDI-agent manipulates a robot as an avatar [53] — a corporeal representation that, by historical analogy, recalls the original meaning of the term "*avatar*" as the incarnation of a deity [6].

Thus, while MAOP "situates" agents in an environment of computational artifacts, it does not consider the agent's body, which can contribute to shaping agents' cognition and behavior. From the perspective of the agent dimension, all perceptions are derived from the environmental dimension and stored as generic perceptions; there is no difference in the physical body perceptions as internal senses, or its awareness of position and movement in space. As a result, identifying where the agent's body begins or ends when developing an embedded BDI-agent-based system is not trivial.

From an engineering perspective, modeling the agent's body as an artifact creates a structural limitation. Artifacts are shareable by design: any agent that focuses on them may perceive their properties and act upon them. However, this behavior is inadequate for apparatuses that constitute the body, which are inherently private and under the exclusive control of a single agent. Therefore, multiple agents running on the same embedded platform may conflict over hardware control, exposing a Practical Gap in MAOP's ability to distinguish between environmental resources and an agent's own physical composition.

This absence of body representation on the MAOP contrasts with the literature on embodied intelligence, which emphasizes that cognition cannot be dissociated from the agent's physical body. Chrisley [20] says that Embodiment involves more than mere physical realization, requiring a body whose morphology and sensorimotor constraints directly shape cognition. Similarly, Cangelosi et al. [17] define embodied intelligence as the strict coupling between brain, body, and environment, mediated by perceptual and motor systems, with principles such as morphological computation and sensorimotor coordination. More recently, Hughes et al. [42] reaffirm that the body must be placed at the center of discussions on intelligence, since embodied systems provide robustness, adaptability, and the capacity to handle unstructured real-world interactions in ways that disembodied AI cannot. Therefore, embodied AI highlights the constitutive role of the physical body in shaping cognition, behavior, and the very boundaries of intelligence [99].

3 BDI-AGENT'S BODY IN LITERATURE

To identify related works focusing on the use of BDI agents in the context of embedded systems with MAOP; we conducted a mapping using the Snowballing [97] method. It is important to note that the goal of this mapping was not to identify frameworks that already include a body model, but to identify all BDI-based solutions capable of managing physical devices and fulfilling the MAOP. Since Boissier et al. [9] is the starting point for the MAOP paradigm, this publication was adopted as the seed for the initial set, and one complete round (*forward* and *backward*) of snowballing was conducted.

In the *Forward Snowballing* step, we looked for citations of the *seed paper* in the SCOPUS and Web of Science (WoS) databases, which together totaled 536 records as of August 2025 (318 in SCOPUS and 218 in WoS). To refine the results and exclude publications unrelated to embedded systems, a filter was applied to restrict those whose titles or abstracts contained the following expressions: "robot", "physical", "device", "embed*", "embod*", or "body". After this, removing duplicates, 63 unique candidate papers remained. Right away, all abstracts were reviewed, from which 28 publications were identified as relevant to the subject. However, of these, only 24 were downloaded, because four were not open access or could not be obtained from the authors. Finally, these available publications were submitted to an initial reading, during which papers were selected that present an *approach*, *framework*, *architecture*, *platform*, *model*, or *middleware* for agent programming. Thus, 12 papers were included in the set for full reading. In the *Backward Snowballing* step, we extracted the references cited by these 12 publications, resulting in a total of 396 references being identified. The titles of the references were read, and 107 were selected; after removing duplicates, this resulted in 72 unique references. After that, the abstracts were read, which reduced this set to 48 candidate papers. Of these, 44 full texts were available (one was already in the set, and three were unavailable for download). Ultimately, after an initial review, 21 papers were deemed relevant and included in the set. Thus, by the end of the mapping process, 33 (12 + 21) relevant papers were selected and read in full. We aim to correlate these publications with adherence to MAOP dimensions and demonstrate the capability to represent a physical agent's body (Table 1).

We observe that several agent-oriented BDI frameworks do not natively provide direct access to physical resources (e.g., 3APL, eJason, Jack, Jadex, and Jason). Other BDI frameworks (JasonEmbedded, lowPowerBDI, and PROFETA) provide native access to physical resources but do not offer an abstraction for modeling a shared environment where agents can perceive or act.

Some frameworks, such as 3APL-M, AFME, GOAL, SOIFRA, or even Jason, which utilizes middleware or an extended agent architecture, provide direct access to physical resources and an abstraction to model a shared environment, but do not support MAOP environmental dimension. It is the case of using Jason with Javino [50] and Javic [37], which are middleware to facilitate the integration of agent frameworks with low-level hardware devices with ATMEGA and PIC microcontrollers, respectively. The same is valid for ARGO [67] and Jason-ROS [82]. ARGO is a customized architecture designed to facilitate the programming of robotic agents

Table 1: Adherence of BDI Frameworks to MAOP Dimensions and Physical Agent Body Representation.

Framework	Related Works	MAOP			Physical resources
		ORG	AGENT	ENV	
3APL	[40]	╳	★	╳	╳
eJason	[30]	╳	★	╳	╳
Jack	[16]	╳	★	≃	╳
Jadex	[70, 79]	╳	★	≃	╳
Jason	[11]	╳	★	≃	╳
jasonEmbedded	[68, 85]	╳	★	╳	♦
lowPowerBDI	[93]	╳	★	╳	♦
PROFETA	[32]	╳	★	╳	♦
3APL-M	[49]	╳	★	≃	♦
AFME	[59]	╳	★	≃	♦
GOAL	[95]	╳	★	≃	♦
SOIFRA	[5]	╳	★	≃	♦
Jason + Javi.*	[37, 50]	╳	★	≃	♦
Jason + ARGO	[65–67]	╳	★	≃	♦
Jason + Jason–ROS	[82]	╳	★	≃	♦
JaCa.*	[21, 22, 52, 80]	╳	★	★	♦
SPADE	[64]	╳	★	★	♦
JaCaMo	[3, 15]	★	★	★	╳
JaCaMo + ARGO	[13, 46, 51, 84]	★	★	★	♦
JaCaMo + JasonArchEmb	[54, 55]	★	★	★	♦
JaCaMo + Jason–ROS	–	★	★	★	♦
JaCaMo + Javi.*	–	★	★	★	♦

[★] MAOP adherent; [≃] with representation, but not adherent to MAOP; [♦] native support; [╳] without representation or support.

in Jason, allowing the agent to interact directly with microcontrollers without explicitly modeling an environment dimension. The Jason-ROS, in turn, is an architecture that complies with Jason, allowing the manipulation of sensors and actuators using the Robot Operating System (ROS) [71]. Also, it was found that some JaCa solutions [21, 22, 52, 80] and the SPADE Platform [64] that provides native access to physical resources; however, they do not support the MAOP organizational dimension.

JaCaMo [10] is the only framework that complies with the three MAOP dimensions. However, it does not access physical resources directly, but supports integration with some middleware or agent architecture to fulfill this need. This framework can be integrated with ARGO agents to manage hardware components or even the architecture JasonArchEmb [54] that allows BDI agents to access hardware interfaces in a layered structure (named by the authors as head-body). However, this approach is brain-centric: the body is treated merely as an interface to hardware controllers rather than as a cognitive component. Furthermore, JaCaMo can be used along with Jason-ROS and the middleware Javino and Javic. These solutions comply with MAOP due to JaCaMo and also represent physical resources natively, making them candidates to create a body representation in practice.

In summary, although the mapping indicates several BDI frameworks, few of them provide direct access to physical resources with comprehensive support for the three MAOP dimensions. However, none offer a unified model that explicitly represents the agent's body with specific perceptions and organizational abstraction. It reinforces the theoretical and practical gap, motivating the development of an extension of the MAOP paradigm to integrate embodied cognition and bodily representation; as well as the comparison with the four solutions identified that fulfill the three MAOP dimensions and interact with physical devices.

4 MY BODY, MY PERCEPTIONS: AN EMBODIED COGNITION APPROACH

BDI-based MAS are rooted in the *computational metaphor*, a classical cognitivist view that conceives the mind as manipulating abstract symbolic representations and that has long shaped cognitive science [33]. However, when dealing with BDI-based embedded systems, the cognition cannot be adequately modeled by the classical approach alone, because its agents are hybrid hardware/software entities coupled with sensors and actuators [51, 68], and their operation is inseparable from the physical environment in which they evolve (their *physical reality*) [45]. So, their intelligent behavior emerges from perception, action, and the morphology and sensorimotor constraints of their physical bodies [4, 99]. Nevertheless, since different perspectives and fields of cognitive science can be explored and applied in AI research [86], an alternative to model these systems can be found in Embodied Cognition. This approach contrasts with the *Brain in a Vat* analogy, which posits a disembodied brain receiving simulated sensory inputs [36], rejecting mind–body dualism and emphasizing that cognition cannot be fully understood through a *brain-centric* view alone, since it assumes that the body and its interactions with the environment are constitutive of cognition but also constraining how the environment can be perceived [81, 96, 99]; advocating that the mind's understanding of the environment depends on the nature of its body; minds in different bodies perceive the same environment differently [81].

In this section, we reconceptualize embodied agents as *mechatropsychosocial* entities, integrating mechatronic, psychological, and social aspects within the same architecture. It defines three new sources of perceptions — *interoception*, *exteroception*, and *proprioception* — grounded in embodied cognition, enabling agents to reason about both internal and external states.

4.1 I Am My Body

Although the current agent-environment relationship is suitable for agent-based systems, we assert that an alternative representation of the body and perceptions is necessary to enhance the agent's reasoning and expressiveness, improving the MAOP to consider the existence of an agent's physical body that can be understood as the embodiment of the agent, or as a set of apparatuses (e.g., locomotion, localization, cleaning) composed of elements (e.g., motors, obstacle sensors, compass) that together constitute the autonomous entity.

Even though the A&A model has been predicting the *Agent Body*, this view aligns with the situatedness idea rather than with embodiment, as observed in Ricci et al. [78], who proposed it as a control interface that the agent's mind uses to manipulate *artifacts* in the environmental dimension. Furthermore, for each agent that joins this *workspace*, a relationship between the agent and the workspace is created, as a body [75]. As an agent can be inserted into multiple *workspaces*, it can simultaneously possess multiple distinct bodies, analogous to the *Brain in a Vat* scenario where the mind is decoupled from the physical body.

For example, consider a team of mobile vacuum-cleaning robots, based on the BDI model, whose mission is to clean the university campus. Upon entering a given room (*workspace*), each robot needs to check whether there is an air conditioner (*artifact*) and whether it is turned on; the air conditioner must be switched off before initiating the cleaning task and reactivated once the robot leaves the room. In the A&A model, the perceptions of environmental artifacts are shared among multiple agents; however, the motor, obstacle sensors, internal operating temperature, and the state of dust capacity provide internal perceptions of the robot itself. In the MAOP, the mobile robot is currently modeled not as a body, but rather as an artifact. Therefore, we advocate that the MAOP should be extended to represent both an environmental artifact (the air conditioner) and an agent body (the robot).

An Embedded MAS needs to be understood differently because the agent's body contributes more than merely causally to cognitive processes. It aligns symbolic deliberation with the embodied, sensorimotor dynamics of intelligent behavior [34]. So, we propose the neologism *mechatropsychosocial*, defined as:

Definition 4.1 (Mechatropsychosocial). of, relating to, or concerned with the mechatronic, psychological, and social aspects of autonomous entities, in contrast to the classical cognitivist (*brain-centric*) model.

This *mechatropsychosocial* entity is composed of three aspects: (1) a *mechatronic* with mechanical and electronic components to perceive and act in the physical world; (2) a *psychological*, the mind of the agents that live inside and control the physical body; (3) a *social* concerning the interaction between internal and external agents, or even with humans.

This neologism is motivated by the biopsychosocial term introduced by Engel [31], which opposes the biomedical model with emphasis on anatomy and physical science, based on the mind–body dualism — the notion of the body as a machine and disease as a "machine failure," with the physician's task being to "repair the machine". Similarly, we assert that the Environment Dimension in MAOP should be reconsidered and expanded to encompass the body. However, just as Engel's biopsychosocial model does not replace the biomedical one, but rather complements and expands it by integrating psychological and social dimensions that enable a more comprehensive understanding of the patient [1], we also present a complementary view where *artifacts* are necessary as well as the *body*, since it constitutes the agent itself. Therefore, they are not exclusive in a MAS.

4.2 My Perceptions

According to Wooldridge [98], a cognitive agent based on the BDI model has databases and functions to support its reasoning process. Considering this, two components help in the process of perceiving the Environment dimension: the *Belief Revision Function* (BRF) is a function that receives new perceptions from the environment and the beliefs stored by the agent, to update the belief set and reflect the current state of the environment; and the *Beliefs Base* (BB) is a representation of what the agent knows about the environment, itself, or other agents.

Each BDI agent has its own BB where its beliefs are stored. Each belief has an additional meta-information that facilitates the description of the agent's behavior. According to Bordini and Hübner [11], the special meta-information $[source(S)]$ establishes that every belief has an origin, which can be one of three possible sources ($S = self|Sender|percept$). When an agent creates a mental note, the source is $self$; when a belief is received from another agent in the system, the source is the name of the agent who sent it. For example, in a system with two agents, the name of the *Sender* is recorded in the receiver's belief; finally the *percept* identify that a belief's origin is from a non-agent element of a MAS, typically regarded as part of the environment.

This annotation mechanism provides an elegant notation to make explicit the sources of an agent's beliefs, with advantages in expressiveness and readability. This source of information is also used by the agent's reasoning (e.g., plans selection), facilitating the description of the agent's behavior, and allowing trust levels to be noted in beliefs [11, 12].

Currently, all information coming from the environment is canonically defined as *percept*. However, a *mechatropsychosocial* entity deals with information and interactions that influence its behavior and can expose its mechatronic body to stress. It can be disturbed by aspects of the *physical reality*, which is formed by *elements* that can present failure. For example, considering a human body (biopsychosocial entity), the visual apparatus has some elements (e.g., eye, eyelids, lacrimal glands, optic nerve), where the eyes can be understood as a sensor that provide optical perception. According to Tossato [89], the perception of an object seen in the environment is given by three phenomena: a physical one, the relationship between the object, the eye, and the light; a physiological and anatomical one, the eye's internal mechanism for capturing raw data; and a mental one, the mental representation of the object seen.

Then, the eyes (as an element of visual apparatus) can hurt in different ways, and each pain tells something different about what's happening. When the pain is due to a very bright light, that pain is kind of an alert about the external environment. Excessively bright light can even make the eyes water to protect themselves. On the other hand, the pain of tired eyes, like the one we feel after hours in front of a computer, is internal information about the state of the eye itself. Thus, the pain caused by a bright light is a reaction to an external stimulus, while the pain of tiredness is a warning that comes from within our own body, indicating that the eyes need a rest [23, 89].

Considering that information about the body could be useful and necessary for the agent (for example, in identifying the need for preventive maintenance), a possible solution to differentiate between the environment and the body is to seek a conceptually grounded way to distinguish different sources of sensory information in biological and psychological sciences — particularly those discussed by Ceunen et al. [18] and by Tuthill and Azim [90] — as they offer precise descriptions of how real organisms perceive internal, positional, and external signals. Therefore, we advocate the need to extend the belief source meta-information, from only $[source(perception)]$ to three different sources into the agent's BB:

- $[source(exteroception)]$: the sensory perception of exogenous stimuli captured by the components of the body and coming from the surroundings (environment) [18];
- $[source(proprioception)]$: is the perception of inhabiting a body, its position, and its movement (kinesthesia) [90];
- $[source(interoception)]$: is the perception of the body's state in the phenomenological experience (internal state) [18].

4.3 $MAOP^{+b}$: MAOP plus Body Representation

Aiming to represent a *mechatropsychosocial* entity, we propose resignifying the MAOP, adding the agent bodies' representation in the endogenous environmental dimension. As shown in Figure 1, we consider an agent's body as a set of apparatuses in the exogenous dimension. Each *Apparatus* is formed by a set of sensors or end-effectors, named *Elements*. These produce interoceptions, proprioceptions, or exteroceptions. Given by:

$$MAOP^{+b} = (\mathcal{O}, \mathcal{A}, \mathcal{W}, \mathcal{B}) \tag{1}$$

where:

- $\mathcal{O}$: It is a set of *organizations*, which are a normative and structural coordination framework that provides a *social cohesion*;
- $\mathcal{A}$: It is a set of *agents*, which are autonomous entities that exhibit a *rationality* characteristic, through a cognitive model;
- $\mathcal{W}$: It is a set of *workspaces*, which are logical containers, that allows a notion of *situatedness*, defining a scope where agents can share to carry out their activities;
- $\mathcal{B}$: It is a set of *bodies*, which are logical containers that allow the notion of *embodiment*, defining the agent' physical reality that is strongly coupled with the real world.

In this extension, bodily elements generate percepts that are integrated directly into the agent's mental state, allowing the physical body to actively influence deliberation rather than functioning only as an interface. This means that the body becomes part of the agent's cognitive condition, shaping how it interprets situations and selects plans. In this context, we proposes that a multi-agent-based embedded system, be defined as:

$$MAS_E = (\mathcal{A}, \mathcal{D}, \mathcal{B}) : \mathcal{D} \subseteq \mathcal{A} \ \wedge \ |\mathcal{D}| \geq 1 \ \wedge \ \forall A_j \in \mathcal{D}, \exists B_m \tag{2}$$

where:

- $\mathcal{A}$ is the group of agents of the MAS $(A_1, A_2, ..., A_N)$;
- D is a subgroup of agents that are embodied $(D_1, D_2, ..., D_n)$;
- $\mathcal{B}_i$ is an abstraction of the set that constitutes the agent's body, defined as $B_i = \{S_i^1, S_i^2, ..., S_i^n\}$;
- S_{ij} is an *apparatus* – an aggregator of a set of elements – defined as $S_{ij} = \{E_{ij}^1, E_{ij}^2, ..., E_{ij}^n\}$;
- E_{ijk} is an *element* – a sensor or an end-effector part.

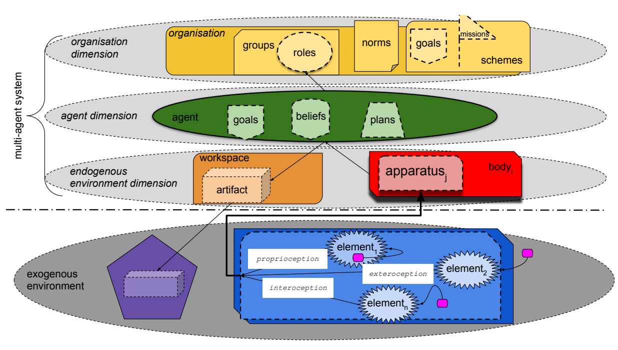

Figure 1: Dimensions of the MAOP paradigm, considering the agent physical body.

4.4 $JACAMO^{+b}$: Instantiation of the $MAOP^{+b}$

JaCaMo is composed of three different technologies: Jason [11], CArtAgO [78], and Moise [44]; however, there are also three other important components into the core: JaCa [73], a framework that enables the interaction of BDI agents with artifacts; ORA4MAS [48], an artifact-based organizational framework that uses the Moise+ model; and Normative Programming Language (NPL) [43], a translator from Organizational Modeling Language (OML) which is interpreted to execute the organization at runtime.

When using the JaCaMo framework, the description of the organizational dimension is loaded by Moise, translated into NPL, and then loaded by the organizational artifacts. In summary, ORA4MAS is the execution infrastructure whose artifacts use an NPL interpreter to bring the high-level organizational specification to life, managing the creation of obligations and enforcing the prohibition of prohibited actions to ensure that agents comply with the organization's rules. Moreover, Jason loads the agents' plans, beliefs, and goal descriptions. In turn, the description of workspaces and artifacts is loaded by CArtAgO. Thereby, when an agent acts as the artifact, it occurs through the JaCa integration. Still, before that, the NPL interpreter is triggered (by ORA4MAS) to check whether this action complies with the organization's norms.

Then, we instantiate the $MAOP^{+b}$ in an extended version of JaCaMo, adding a new component to the framework core, the *ageNt Embodied Cognition development Kit* (NECK)[1]. This component accepts physical body descriptions and loads the agents' bodies from the project specification. An overview of the $JACAMO^{+b}$, is shown in Figure 2. The body description receives the name of the embodied agent and a set of apparatus that constitute the agent's body, with their respective implementation. NECK continuously collects data from the elements that compose each *apparatus* and converts them into percepts automatically injected into the agent's belief base at each reasoning cycle, thereby implementing a passive perception approach.

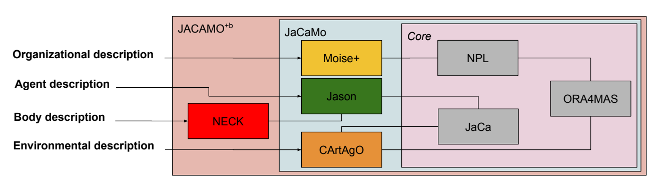

Figure 2: Extension of JaCaMo for representing the agent's physical body.

Additionally, it was necessary to adjust the grammar of the JCM project by adding body and apparatus definitions. In Code 1, an example of a JCM^{+b} project specification. Now, every agent can have a body, and during the project loading, the interpreter checks whether a body is specified for that agent (lines 5-7). If so, the body with all assembled apparatus is bound to the agent. Therefore, perceptions from the body are transmitted directly to the agent's mind, updating the BB at each reasoning cycle. Finally, an internal action *.act* was added so that the agent can act directly in the physical world.

[1] https://jacamob.chon.group

```
mas jcmWithBodyProjectName {
    agent agentName: agt1.asl {
        focus: workspaceName.a1
    }
    body agentName {
        apparatus ap1 : neck.DefaultApparatus("/dev/ttyUSB0")
    }
    workspace workspaceName {
      artifact a1 : environmental.Implementation
    }
    organisation orgName: org.xml {
      group g1 : groupName {
        players: agentName roleName
      }
    }
}
```

Listing 1: $JACAMO^{+b}$ **project specification.**

5 EVALUATION

To demonstrate the adequacy of the $JACAMO^{+b}$, we adopted the *Descriptive Evaluation* method, which makes it possible to demonstrate the utility of a DSR artifact (construct, model, or instantiation) [26]. This method enables the observation and qualitative analysis of the proposed model's behavior in comparison with others in the literature, specifically the four MAOP-compliant solutions (ARGO, Jason-ROS, JasonArchEmb, Javino) that provide physical resource support for BDI agents in JaCaMo. It will take place through an *Illustrative Scenario* – a DSR descriptive evaluation technique [69] that guides the construction of a synthetic scenario to demonstrate the utility and effectiveness of this proposal in relation to the related works.

We consider the scenario of a BDI agent-based system embedded in a robotic vacuum cleaner, responsible for autonomously cleaning a campus. When physically entering a laboratory, the agent must join the laboratory's workspace and access the environmental artifacts. Subsequently, it must start cleaning the laboratory floor. Since dust can clog the internal robot ventilation, causing battery overheating and damage to its functioning, the agent needs to monitor its internal temperature. If the agent perceives that the internal robot temperature exceeds the manufacturer's specification, it must turn off the cleaning mechanism and raise an overheating alert, thereby avoiding exceeding the maximum operating temperature.

The considered successful scenario is that the agent avoids burning out the equipment, given that the robot's temperature increases over time. However, as the air conditioner (artifact) in the laboratory (workspace) and the vacuum (apparatus) inform percepts as *powerStatus* and *temperature*, we aim to verify if the agent can distinguish correctly its internal temperature from the ambient temperature. Additionally, we will compare the evaluated solutions to determine if they have an agent's body representation, how the agent understands its body percepts, and if the absence of body representation affects the agent's behavior.

To provide a synthetic situation necessary for evaluating each DSR artifact, we opt to use two software simulators that replicate the behavioral patterns of the physical devices. The first simulates the laboratory's air conditioner (as a MAOP artifact) using a Web of Digital Twins approach that models the properties, relationships, and events of physical assets [72]. The second simulates the robot's microcontroller (as a $MAOP^{+b}$ agent-body apparatus), using a simulated exogenous environment that replaces physical devices, without requiring changes to the agent's specification [35].

To provide the qualitative analysis, we consider three criteria: *i)* whether the instantiation is functional and fully compliant with MAOP; *ii)* whether it allows the representation of an agent's body; *iii)* whether it allows the representation of the agent's bodily perceptions. Also, the evaluation was conducted in three stages, one for each criterion: In the first step, the availability and compatibility of the related works were evaluated, and a MAS was developed to fulfill the scenario; In the second, we look at how each of them represents the agent's body in the MAS; In the last, it was observed how the agent perceives its body while running the MAS. The evaluation summary is presented in Table 2.

In the first stage, five different BDI-agent embedded systems were developed: one using the $JACAMO^{+b}$; one using JaCaMo and Javino; and three of them using specialized BDI agent architectures for JaCaMo (ARGO, JasonROS, and JasonArchEmb). During this stage, $JACAMO^{+b}$, Javino and a JasonROS extension (EMAS Jason-ROS, proposed by Silvestre et al [83]) demonstrated full compatibility with the most recent version of JaCaMo. However, it was noted that JasonArchEmb is abandoned and not compliant with the latest versions of JaCaMo. When attempting to run the MAS, an incompatibility of the architecture with the framework message is displayed, and the system does not execute. Furthermore, although the ARGO architecture remains active and the MAS is running, it presents a conflict with CArtAgO (environment dimension); the architecture does not allow an agent to join a workspace. Therefore, these two architectures proved unfeasible and did not proceed to the next phase.

In the second stage, it was observed that Javino does not support representing the agent's body. As shown in Figure 3a, the microcontroller is described as a workspace artifact (lines 9-11), similar to the air conditioner (lines 6-8). Additionally, it was noted that JasonROS allows an implicit body representation, as depicted in Figure 3b. It is necessary to define the *agent class* and *agent architecture* in the MAS project (lines 4-5), and a manifest file (with the same name as the agent in the /src/agt folder). On the other hand, as shown in Figure 3c, when using the $JACAMO^{+b}$, the agent has an explicit description of the body (lines 7-9). Additionally, it is not tied to a specific technology (e.g., Javino or ROS), allowing customization (line 8); unlike the other solutions.

In the final stage, the scenario was implemented and run using the JaCaMo+Javino, Jason-ROS, and $JACAMO^{+b}$. To ensure a fair comparison, the same agent program (plans, goals, and decision

Table 2: Descriptive Evaluation with Illustrative Scenario comparing the proposal of this article with other proposals from the literature.

	Full MAOP adherence	Agent's body representation	Agent's body percepts
JasonArchEmb [54]	╳	∅	∅
ARGO [67]	╳	∅	∅
Javino [50]	★	╳	╳
EMAS Jason-ROS [83]	★	≃	╳
$JACAMO^{+b}$	★	★	★

[★] valid; [╳] absent; [≃] non-explicit; [∅] = not applicable

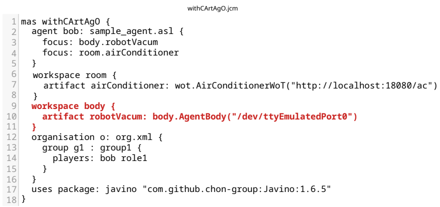

(a) Body's representation using JaCaMo+Javino.

(b) Body's representation using JaCaMo+JasonROS.

(c) Body's representation using $JACAMO^{+b}$.

(d) Belief Base using JaCaMo+Javino.

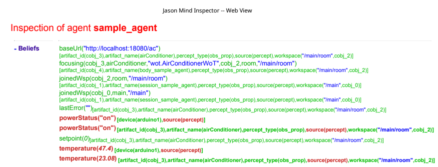

(e) Belief Base using JaCaMo+JasonROS.

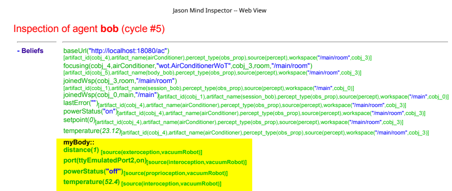

(f) Belief Base using $JACAMO^{+b}$.

Figure 3: Overview of Descriptive Evaluation and Illustrative Scenario execution.

rules) was used in all three implementations; only the way percepts are delivered by each framework differs. It was observed that the first two solutions failed in the proposed scenario, resulting in damage to the robot vacuum. This failure occurs because the temperature percept from the air conditioner and the temperature percept from the robot's internal hardware are both delivered as identical beliefs with the same predicate and the same $source(percept)$ annotation, causing the most recent percept to overwrite the previous one in the belief base and preventing the agent from distinguishing ambient temperature from its own internal temperature. When inspecting the agents' belief base in these MAS, we notice that they do not have an explicit representation of body perceptions (Figures 3d-3e). The perceptions *powerStatus* and *temperature* from both the air conditioner and the microcontroller had the same $source(percept)$, being indistinguishable between the robot's functioning and temperature or those of the air conditioning, in the BB. The $JACAMO^{+b}$, in turn, successfully executed the scenario, and the agent stopped the vacuum before the temperature exceeded the manufacturer's defined limit. When inspecting the agent mind (Figure 3f), there is an explicit representation of body perceptions using the $source(interoception|exteroception|proprioception)$, and the *namespace* – an AOP Modularity Approach [62] – called *myBody::*.

The successful demonstration that the proposed extension met your utility and effectiveness criteria. Furthermore, we have responded to the research questions: 1) How can a cognitive agent perceive its own body? 2) Is it possible to extend the MAOP to represent a notion of the agent's body? The first was achieved through an embodied cognitive approach that provides a differentiation of the body from environmental percepts, directly into the Belief Base, allowing a cognitive agent to perceive its own body correctly. The second, over an intervention into the MAOP environmental dimension and adjusting the grammar of multi-agent system specification, it was possible to add the notion of the agent's body.

5.1 Reproducibility

To ensure research replication by providing a detailed understanding of the artifacts produced and the evaluation performed we provide a specific page[2] containing general information about tests, source codes, and videos.

6 CONCLUSION

This work discussed the limitations of traditional computational paradigms for representing agents' physical bodies in BDI-based multi-agent systems. Furthermore, it revealed the inadequacies of current MAOP-compliant frameworks in both representing and managing agents' physical bodies. Inspired by the theory of Embodied Cognition, we propose a new model that aims to enhance the BDI model by incorporating new sources of perception, including interoception, exteroception, and proprioception. This change enabled agents to better perceive and reason about their bodies, with advantages in expressiveness and readability, in addition to facilitating the developer's description of the agent's behavior.

The proposed model was implemented as an extension of the JaCaMo framework, and the evaluation performed demonstrated the usefulness and effectiveness of the proposed artifact. It not only confirmed that the artifact supported the problem solution but also complied with the DSR design evaluation guideline. The successful execution of the scenario and the prevention of damage to the vacuum cleaner robot validated the final instantiation, demonstrating that the artifact meets the problem's requirements and constraints, operating more efficiently and effectively than alternatives available in the literature. Future work will explore more complex scenarios, with multiple agents and multiple bodies. It will pave the way for further improvements in MAS embodied cognition, such as advanced fault tolerance mechanisms and the ability to exchange physical components between agents at runtime.

[2]https://papers.chon.group/AAMAS/2026/myBodyMyPercepts/

REFERENCES

[1] Rolf H. Adler. 2009. Engel's biopsychosocial model is still relevant today. *Journal of Psychosomatic Research* 67, 6 (Dec. 2009), 607–611. https://doi.org/10.1016/j.jpsychores.2009.08.008

[2] Carlos S. Alvarado. 2006. Human Radiations: Concepts of Force in Mesmerism, Spiritualism and Psychical Research. *JSPR* 70.3, 884 (2006), 138–162. https://scholar.google.com/scholar?cluster=12377861775120697234

[3] Cleber Jorge Amaral and Jomi Fred Hübner. 2020. Jacamo-Web is on the Fly: An Interactive Multi-Agent System IDE. In *EMAS 2019*. Vol. 12058. Springer, Cham. https://doi.org/10.1007/978-3-030-51417-4_13

[4] Michael L. Anderson. 2003. Embodied Cognition: A field guide. *Artificial Intelligence* 149, 1 (2003), 91–130. https://doi.org/10.1016/S0004-3702(03)00054-7

[5] Willson Amalraj Arokiasami, Prahlad Vadakkepat, Kay Chen Tan, and Dipti Srinivasan. 2016. Interoperable multi-agent framework for unmanned aerial/ground vehicles: towards robot autonomy. *Complex & Intelligent Systems* 2, 1 (March 2016), 45–59. https://doi.org/10.1007/s40747-016-0014-8

[6] Kevin Beaufils and Alexis Berland. 2022. Avatar embodiment: from cognitive self-representation to digital body ownership. *Hybrid* 9 (Nov. 2022). https://doi.org/10.4000/hybrid.2664

[7] Michael Beetz, Lorenz Mösenlechner, and Moritz Tenorth. 2010. CRAM — A Cognitive Robot Abstract Machine for everyday manipulation in human environments. In *IROS 2010*. IEEE, Taipei. https://doi.org/10.1109/IROS.2010.5650146

[8] Olivier Boissier, Rafael H. Bordini, Jomi Fred Hübner, and Alessandro Ricci. 2020. *Multi-agent oriented programming: programming multi-agent systems using JaCaMo*. The MIT Press, Cambridge, Massachusetts.

[9] Olivier Boissier, Rafael H. Bordini, Jomi F. Hübner, Alessandro Ricci, and Andrea Santi. 2013. Multi-agent oriented programming with JaCaMo. *Science of Computer Programming* 78, 6 (2013). https://doi.org/10.1016/j.scico.2011.10.004

[10] Olivier Boissier, Jomi F. Hübner, and Alessandro Ricci. 2016. The JaCaMo Framework. In *Social Coordination Frameworks for Social Technical Systems*. Vol. 30. Springer, Cham, 125–151. https://doi.org/10.1007/978-3-319-33570-4_7

[11] Rafael H. Bordini and Jomi F. Hübner. 2006. BDI Agent Programming in AgentSpeak Using Jason. In *Computational Logic in Multi-Agent Systems*. Vol. 3900. Springer Berlin Heidelberg, 143–164. https://doi.org/10.1007/11750734_9

[12] Rafael H. Bordini, Jomi F. Hübner, and Renata Vieira. 2005. Jason and the Golden Fleece of Agent-Oriented Programming. In *Multi-Agent Programming*. Vol. 15. Springer US, Boston, MA, 3–37. https://doi.org/10.1007/0-387-26350-0_1

[13] Fabian Cesar Brandão, Maria Alice Trinta Lima, Carlos Eduardo Pantoja, Jean Zahn, and José Viterbo. 2021. Engineering Approaches for Programming Agent-Based IoT Objects Using the Resource Management Architecture. *Sensors* 21, 23 (Dec. 2021), 8110. https://doi.org/10.3390/s21238110

[14] Michael E. Bratman, David J. Israel, and Martha E. Pollack. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence* 4, 3 (Sept. 1988), 349–355. https://doi.org/10.1111/j.1467-8640.1988.tb00284.x

[15] Samuele Burattini, Alessandro Ricci, Simon Mayer, Danai Vachtsevanou, Jérémy Lemee, Andrei Ciortea, and Angelo Croatti. 2022. Agent-Oriented Visual Programming for the Web of Things. In *Engineering Multi-Agent Systems 2022*. http://www.alexandria.unisg.ch/266966/

[16] Paolo Busetta, Ralph Rönnquist, Andrew Hodgson, and Andrew Lucas. 1999. Jack intelligent agents-components for intelligent agents in java. *AgentLink News Letter* 2, 1 (1999), 2–5. https://scholar.google.com.br/scholar?cluster=10359603077746532284

[17] Angelo Cangelosi, Josh Bongard, Martin H. Fischer, and Stefano Nolfi. 2015. Embodied Intelligence. In *Springer Handbook of Computational Intelligence*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-43505-2_37

[18] Erik Ceunen, Johan W. S. Vlaeyen, and Ilse Van Diest. 2016. On the Origin of Interoception. *Frontiers in Psychology* 7 (May 2016). https://doi.org/10.3389/fpsyg.2016.00743

[19] Uchendu Eugene Chigbu, Sulaiman Olusegun Atiku, and Cherley C. Du Plessis. 2023. The Science of Literature Reviews: Searching, Identifying, Selecting, and Synthesising. *Publications* 11, 1 (Jan. 2023), 2. https://doi.org/10.3390/publications11010002

[20] Ron Chrisley. 2003. Embodied artificial intelligence. *Artificial Intelligence* 149, 1 (Sept. 2003), 131–150. https://doi.org/10.1016/S0004-3702(03)00055-9

[21] Angelo Croatti and Alessandro Ricci. 2020. The JaCa-Android Framework for Programming BDI-Based Personal Agents on Mobile Devices. In *PAAMS 2020*. Springer. https://doi.org/10.1007/978-3-030-49778-1_7

[22] Angelo Croatti and Alessandro Ricci. 2020. Mobile Apps as Personal Assistant Agents: the JaCa-Android Framework for programming Agents-based applications on mobile devices. *Autonomous Agents and Multi-Agent Systems* 34, 2 (Oct. 2020), 48. https://doi.org/10.1007/s10458-020-09474-7

[23] András Csillag. 2005. The Organ of Vision. In *Atlas of the Sensory Organs*. Humana Press, Totowa, NJ, 85–164. https://doi.org/10.1385/1-59259-849-8:085

[24] Darryl Davis and James Gwatkin. 2010. robo-CAMAL: A BDI Motivational Robot. *Paladyn, Journal of Behavioral Robotics* 1, 2 (June 2010), 116–129. https://doi.org/10.2478/s13230-010-0010-4

[25] Lavindra De Silva, Felipe Meneguzzi, and Brian Logan. 2020. BDI agent architectures: a survey. In *IJCAI'20*. ACM, Yokohama, Japan, 4914–4921. https://dl.acm.org/doi/abs/10.5555/3491440.3492124

[26] Aline Dresch, Daniel Pacheco Lacerda, and José Antônio Valle Antunes. 2015. Design Science Research. In *Design Science Research*. Springer International Publishing, Cham, 67–102. https://doi.org/10.1007/978-3-319-07374-3_4

[27] B.R. Duffy, G.M.P. O'Hare, A.G. Campbell, J. Stafford, and M.J. O'Grady. 2005. NEXUS: Fusing the Real & the Virtual. In *EUROCON 2005 - The International Conference on "Computer as a Tool"*. IEEE, Belgrade, Serbia and Montenegro, 716–719. https://doi.org/10.1109/EURCON.2005.1630031

[28] B.R. Duffy, G.M.P. O'Hare, A.N. Martin, J.F. Bradley, and B. Schon. 2003. Agent chameleons: agent minds and bodies. In *Proceedings 11th IEEE International Workshop on Program Comprehension*. IEEE Comput. Soc, New Brunswick, NJ, USA, 118–125. https://doi.org/10.1109/CASA.2003.1199312

[29] B. R. Duffy, G. M. P. O'Hare, R. P. S. O'Donoghue, C. F. B. Rooney, and R. W. Collier. 2000. Reality & Virtual Reality in Mobile Robotics. In *Managing Interactions in Smart Environments*. Springer. https://doi.org/10.1007/978-1-4471-0743-9_13

[30] Álvaro Fernández Díaz, Clara Benac Earle, and Lars-Åke Fredlund. 2013. eJason: An Implementation of Jason in Erlang. In *Programming Multi-Agent Systems*. Springer Berlin Heidelberg, 1–16. https://doi.org/10.1007/978-3-642-38700-5_1

[31] George L. Engel. 1977. The Need for a New Medical Model: A Challenge for Biomedicine. *Science* 196, 4286 (1977). https://doi.org/10.1126/science.847460

[32] Loris Fichera, Fabrizio Messina, Giuseppe Pappalardo, and Corrado Santoro. 2017. A Python framework for programming autonomous robots using a declarative approach. *Science of Computer Programming* 139 (June 2017), 36–55. https://doi.org/10.1016/j.scico.2017.01.003

[33] Kate Finley. 2025. Embodied Cognition and the Grip of Computational Metaphors. *Ergo* 12 (2025). https://doi.org/10.3998/ergo.7136

[34] Dario Floreano, Francesco Mondada, Andres Perez-Uribe, and Daniel Roggen. 2004. Evolution of Embodied Intelligence. In *Embodied Artificial Intelligence*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-27833-7_23

[35] Bruno Freitas, Nilson Lazarin, Carlos Pantoja, and José Viterbo. 2025. Integrating Simulated Exogenous Environments to Support the Learning Process of the Embedded MAS Approach. In *33rd WEI*. https://doi.org/10.5753/wei.2025.9115

[36] Cathy Gere. 2004. The brain in a vat. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 35, 2 (June 2004), 219–225. https://doi.org/10.1016/j.shpsc.2004.03.002

[37] Joao Victor Guinelli and Carlos Pantoja. 2016. A Middleware for Using PIC Microcontrollers and Jason Framework for Programming Multi-Agent Systems. In *Anais do WPCCG*, Vol. 1. UTFPR. http://www.wpccg.pro.br/volume001.html

[38] Nesrine Hafiene, Luis G. Nardin, Antoine Zimmermann, and Olivier Boissier. 2023. Governing Online Forum Interactions with Multi-agent System: A Reddit Use Case with the JaCaMo Platform. In *PAAMS 2023*. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-37616-0_10

[39] Mahdi Hannoun, Olivier Boissier, Jaime S. Sichman, and Claudette Sayettat. 2000. MOISE: An Organizational Model for Multi-agent Systems. In *IBERAMIA 2000, SBIA 2000*. Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-44399-1_17

[40] Koen V. Hindriks, Frank S. De Boer, Wiebe Van Der Hoek, and John-Jules Ch. Meyer. 1999. Agent Programming in 3APL. *Autonomous Agents and Multi-Agent Systems* 2, 4 (1999), 357–401. https://doi.org/10.1023/A:1010084620690

[41] Jomi F. Hubner, Jaime S. Sichman, and Olivier Boissier. 2007. Developing organised multiagent systems using the MOISE+ model: programming issues at the system and agent levels. *International Journal of Agent-Oriented Software Engineering* 1, 3/4 (2007), 370. https://doi.org/10.1504/IJAOSE.2007.016266

[42] Josie Hughes, Arsen Abdulali, Ryman Hashem, and Fumiya Iida. 2022. Embodied Artificial Intelligence: Enabling the Next Intelligence Revolution. *IOP Conference Series: Materials Science and Engineering* 1261, 1 (Oct. 2022), 012001. https://doi.org/10.1088/1757-899X/1261/1/012001

[43] Jomi F. Hübner, Olivier Boissier, and Rafael H. Bordini. 2010. A Normative Organisation Programming Language for Organisation Management Infrastructures. In *COIN 2009*. Springer. https://doi.org/10.1007/978-3-642-14962-7_8

[44] Jomi Fred Hübner, Jaime Simão Sichman, and Olivier Boissier. 2006. S-Moise+ : A Middleware for Developing Organised Multi-agent Systems. In *AAMAS 2005*. Vol. 3913. Springer Berlin Heidelberg, 64–77. https://doi.org/10.1007/11775331_5

[45] Jean-Paul Jamont and Michel Occello. 2008. A multiagent method to design hardware/software collaborative systems. In *12th CSCWD*. IEEE, Xian, China, 361–366. https://doi.org/10.1109/CSCWD.2008.4537007

[46] Vinicius Souza De Jesus, Nilson Lazarin, Carlos Pantoja, Fabian Brandão Manoel, Gleifer Alves, and José Viterbo. 2023. A middleware for providing communicability to Embedded MAS based on the lack of connectivity. *Artificial Intelligence Review* 56, S3 (2023), 2971–3001. https://doi.org/10.1007/s10462-023-10596-z

[47] Burak Karaduman, Baris Tekin Tezel, and Moharram Challenger. 2023. Rational software agents with the BDI reasoning model for Cyber–Physical Systems. *Engineering Applications of Artificial Intelligence* 123 (2023). https://doi.org/10.1016/j.engappai.2023.106478

[48] Rosine Kitio, Olivier Boissier, Jomi Fred Hübner, and Alessandro Ricci. 2008. Organisational Artifacts and Agents for Open Multi-Agent Organisations: "Giving the Power Back to the Agents". In *COIN 2007*. Vol. 4870. Springer Berlin Heidelberg, 171–186. https://doi.org/10.1007/978-3-540-79003-7_13
[49] Fernando Koch. 2005. 3APL-M platform for deliberative agents in mobile devices. In *AAMAS '05*. ACM, The Netherlands. https://doi.org/10.1145/1082473.1082827
[50] Nilson Lazarin and Carlos Pantoja. 2015. A Robotic-agent Platform For Embedding Software Agents using Raspberry Pi and Arduino Boards. In *9th WESAAC* (Niterói/RJ). SBC, 13–20. https://doi.org/10.5753/wesaac.2015.33308
[51] Nilson Mori Lazarin, Carlos Eduardo Pantoja, and José Viterbo. 2024. Dealing with the Unpredictability of Physical Resources in Real-World Multi-agent Systems. In *ICAART 2023*. Springer, Cham. https://doi.org/10.1007/978-3-031-55326-4_3
[52] Fabian Cesar Pereira Brandão Manoel, Carlos Eduardo Pantoja, Leandro Samyn, and Vinicius Souza de Jesus. 2020. Physical Artifacts for Agents in a Cyber-Physical System: A Case Study in Oil & Gas Scenario (EEAS). In *SEKE 2020*. KSI Research Inc., 55–60. https://doi.org/10.18293/SEKE2020-154
[53] Alan Martin, Brian Duffy, Gregory O'Hare, Bianca Schoen-Phelan, and John Bradley. 2004. Empowering Agents within Virtual Environments. In *(ICS'04)*. University of Ulster at Magee, Londonderry. https://doi.org/10.21427/FFZC-5Z45
[54] Marcelo S. Menegol, Jomi F. Hübner, and Leandro B. Becker. 2018. Coordinated UAV Search and Rescue Application with JaCaMo. In *PAAMS 2018*. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-94580-4_33
[55] Marcelo S. Menegol, Jomi F. Hübner, and Leandro B. Becker. 2018. Evaluation of Multi-agent Coordination on Embedded Systems. In *PAAMS 2018*. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-94580-4_17
[56] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. Von Hofsten, K. Rosander, M. Lopes, J. S-Victor, A. Bernardino, and L. Montesano. 2010. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks* 23, 8-9 (2010). https://doi.org/10.1016/j.neunet.2010.08.010
[57] R. Montufar-Chaveznava and J.A. Mendez-Polanco. 2005. Multiagent Architecture for Telerobotics. In *CONIELECOMP'05*. https://doi.org/10.1109/CONIEL.2005.52
[58] Wagner Francisco Mourão. 2005. Coesão social e individualização: análises e interpretações da obra de Emile Durkheim - da divisão do trabalho social. *Cadernos de Campo: Revista de Ciências Sociais* 11 (2005). https://periodicos.fclar.unesp.br/cadernos/article/view/10552
[59] C. Muldoon, G. M. P. O'Hare, R. Collier, and M. J. O'Grady. 2006. Agent Factory Micro Edition: A Framework for Ambient Applications. In *ICCS 2006*. Vol. 3993. Springer Berlin Heidelberg, 727–734. https://doi.org/10.1007/11758532_95
[60] Thomas M. Olshewsky. 1976. On the Relations of Soul to Body in Plato and Aristotle. *Journal of the History of Philosophy* 14, 4 (Oct. 1976), 391–404. https://doi.org/10.1353/hph.2008.0163
[61] Chidiebere Onyedinma, Patrick Gavigan, and Babak Esfandiari. 2020. Toward Campus Mail Delivery Using BDI. *Journal of Sensor and Actuator Networks* 9, 4 (Dec. 2020), 56. https://doi.org/10.3390/jsan9040056
[62] Gustavo Ortiz-Hernández, Jomi Fred Hübner, Rafael H. Bordini, Alejandro Guerra-Hernández, Guillermo J. Hoyos-Rivera, and Nicandro Cruz-Ramírez. 2016. A Namespace Approach for Modularity in BDI Programming Languages. In *EMAS 2016*. Springer, Cham, 117–135. https://doi.org/10.1007/978-3-319-50983-9_7
[63] Greg M. P. O'Hare, Rem Collier, J. Conlon, and S. Abbas. 1998. Agent Factory : An Environment for Constructing and Visualising Agent Communities. In *9th AICS*. UCD, Dublin. http://hdl.handle.net/10197/4410
[64] Javier Palanca, Jaime Rincon, Vicente Julian, Carlos Carrascosa, and Andrés Terrasa. 2022. Developing IoT Artifacts in a MAS Platform. *Electronics* 11, 4 (Feb. 2022), 655. https://doi.org/10.3390/electronics11040655
[65] Carlos Pantoja, Vinicius Souza De Jesus, Fabian Manoel, and José Viterbo. 2018. A Heterogeneous Architecture for Integrating Multi-Agent Systems in AmI Systems (S). 242–285. https://doi.org/10.18293/SEKE2018-211
[66] Carlos Pantoja, Heder Soares, José Viterbo, and Amal El Fallah Seghrouchni. 2018. An Architecture for the Development of Ambient Intelligence Systems Managed by Embedded Agents. 215–249. https://doi.org/10.18293/SEKE2018-110
[67] Carlos Pantoja, Márcio Stabile, Nilson Lazarin, and Jaime Sichman. 2016. ARGO: An Extended Jason Architecture that Facilitates Embedded Robotic Agents Programming. In *EMAS 2016*. Springer. https://doi.org/10.1007/978-3-319-50983-9_8
[68] Carlos Eduardo Pantoja, Vinicius Souza De Jesus, Nilson Mori Lazarin, and José Viterbo. 2023. A Spin-off Version of Jason for IoT and Embedded Multi-Agent Systems. In *BRACIS 2023*. Springer. https://doi.org/10.1007/978-3-031-45368-7_25
[69] Ken Peffers, Marcus Rothenberger, Tuure Tuunanen, and Reza Vaezi. 2012. Design Science Research Evaluation. In *DESRIST 2012*. Vol. 7286. Springer Berlin Heidelberg, 398–410. https://doi.org/10.1007/978-3-642-29863-9_29
[70] Alexander Pokahr, Lars Braubach, and Winfried Lamersdorf. 2005. Jadex: A BDI Reasoning Engine. In *Multi-Agent Programming*. Vol. 15. Springer US, Boston, MA, 149–174. https://doi.org/10.1007/0-387-26350-0_6
[71] Morgan Quigley, Ken Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and Andrew Ng. 2009. ROS: an open-source Robot Operating System. In *ICRA 2009*. Kobe, Japan. https://scholar.google.com/scholar?cluster=143767492575573826
[72] Alessandro Ricci, Angelo Croatti, Stefano Mariani, Sara Montagna, and Marco Picone. 2022. Web of Digital Twins. *ACM Transactions on Internet Technology* 22, 4 (Nov. 2022), 1–30. https://doi.org/10.1145/3507909
[73] Alessandro Ricci, Michele Piunti, L. Daghan Acay, Rafael H. Bordini, Jomi F. Hübner, and Mehdi Dastani. 2008. Integrating heterogeneous agent programming platforms within artifact-based environments. In *AAMAS '08*. IFAAMAS, Richland, SC, 225–232. https://dl.acm.org/doi/10.5555/1402383.1402419
[74] Alessandro Ricci, Michele Piunti, and Mirko Viroli. 2011. Environment programming in multi-agent systems: an artifact-based perspective. *Autonomous Agents and Multi-Agent Systems* 23, 2 (2011). https://doi.org/10.1007/s10458-010-9140-7
[75] Alessandro Ricci, Michele Piunti, Mirko Viroli, and Andrea Omicini. 2009. Environment Programming in CArtAgO. In *Multi-Agent Programming*. Springer US, Boston, MA, 259–288. https://doi.org/10.1007/978-0-387-89299-3_8
[76] Alessandro Ricci, Andrea Santi, and Michele Piunti. 2012. Action and Perception in Agent Programming Languages: From Exogenous to Endogenous Environments. In *ProMAS 2010*. Springer. https://doi.org/10.1007/978-3-642-28939-2_7
[77] Alessandro Ricci, Mirko Viroli, and Andrea Omicini. 2006. Programming MAS with Artifacts. In *ProMAS 2005*. Springer. https://doi.org/10.1007/11678823_13
[78] Alessandro Ricci, Mirko Viroli, and Andrea Omicini. 2007. CArtA gO: A Framework for Prototyping Artifact-Based Environments in MAS. In *Environments for Multi-Agent Systems III*. Springer. https://doi.org/10.1007/978-3-540-71103-2_4
[79] Sebastian Rockel, Denis Klimentjew, and Jianwei Zhang. 2012. A multi-robot platform for mobile robots - A novel evaluation and development approach with multi-agent technology. In *MFI 2012*. IEEE, Hamburg, Germany, 470–477. https://doi.org/10.1109/MFI.2012.6343020
[80] Andrea Santi, Marco Guidi, and Alessandro Ricci. 2011. JaCa-Android: An Agent-Based Platform for Building Smart Mobile Applications. In *LADS 2010*. Springer Berlin Heidelberg, 95–114. https://doi.org/10.1007/978-3-642-22723-3_6
[81] Lawrence Shapiro and Steven A. Stolz. 2019. Embodied cognition and its significance for education. *Theory and Research in Education* 17, 1 (March 2019), 19–39. https://doi.org/10.1177/1477878518822149
[82] Gustavo R. Silva, Leandro Becker, and Jomi F. Hübner. 2020. Embedded Architecture Composed of Cognitive Agents and ROS for Programming Intelligent Robots. *IFAC-PapersOnLine* 53, 2 (2020). https://doi.org/10.1016/j.ifacol.2020.12.2718
[83] Iago Silvestre, L. Becker, M. Fisher, Jomi Hübner, and Maiquel De Brito. 2025. Enhanced agent-oriented programming for robot teams. *Engineering Applications of Artificial Intelligence* 158 (2025). https://doi.org/10.1016/j.engappai.2025.111390
[84] Lucas F. Souza de Castro, Fabian Manoel, Vinicius De Jesus, Carlos Pantoja, Andre Borges, and Gleifer Alves. 2022. Integrating Embedded Multiagent Systems with Urban Simulation Tools and IoT Applications. *Revista de Informática Teórica e Aplicada* 29, 1 (2022), 81–90. https://doi.org/10.22456/2175-2745.110837
[85] Vinicius Souza De Jesus, Nilson Mori Lazarin, Carlos Eduardo Pantoja, Gleifer Vaz Alves, Gabriel Ramos Alves De Lima, and Jose Viterbo. 2023. An IDE to Support the Development of Embedded Multi-Agent Systems. In *PAAMS 2023*. Springer, Cham. https://doi.org/10.1007/978-3-031-37616-0_29
[86] Paul Thagard. 2023. Cognitive Science. In *The Stanford Encyclopedia of Philosophy* (winter 2023 ed.). Metaphysics Research Lab, Stanford University, Online. https://plato.stanford.edu/archives/win2023/entries/cognitive-science/
[87] Jorge A. Torres, Luciana P. Nedel, and Rafael H. Bordini. 2003. Using the BDI Architecture to Produce Autonomous Characters in Virtual Worlds. In *IVA 2003*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-39396-2_32
[88] G.J. Tortora and B.H. Derrickson. 2017. *Introduction to the Human Body*. Wiley.
[89] Claudemir Roque Tossato. 2005. The function of human eye in the optics at the end of XVIth century. *Scientiae Studia* 3, 3 (2005), 415–441. https://doi.org/10.1590/S1678-31662005000300004
[90] John C. Tuthill and Eiman Azim. 2018. Proprioception. *Current Biology* 28, 5 (March 2018), R194–R203. https://doi.org/10.1016/j.cub.2018.01.064
[91] Christoph Urban. 2001. PECS. In *Deformable Avatars*. Springer US, Boston, MA, 206–216. https://doi.org/10.1007/978-0-306-47002-8_18
[92] Christoph Urban and Bernd Schmidt. 2001. PECS – Agent-Based Modelling of Human Behaviour. In *AAAI Fall Symposium - Technical Report*, Vol. FS-01-02. 141 – 146. https://aaai.org/papers/0027-fs01-02-027/.
[93] Danai Vachtsevanou, Jannik William, Matuzalém M. Dos Santos, Maiquel De Brito, Jomi Fred Hübner, Simon Mayer, and Andres Gomez. 2023. Embedding Autonomous Agents into Low-Power Wireless Sensor Networks. In *PAAMS 2023*. Springer. https://doi.org/10.1007/978-3-031-37616-0_31
[94] David Vernon. 2010. Enaction as a Conceptual Framework for Developmental Cognitive Robotics. *Paladyn, Journal of Behavioral Robotics* 1, 2 (June 2010), 89–98. https://doi.org/10.2478/s13230-010-0016-y
[95] Changyun Wei and Koen V. Hindriks. 2013. An Agent-Based Cognitive Robot Architecture. In *ProMAS 2012*. https://doi.org/10.1007/978-3-642-38700-5_4
[96] Margaret Wilson. 2002. Six views of embodied cognition. *Psychonomic Bulletin & Review* 9, 4 (Dec. 2002), 625–636. https://doi.org/10.3758/BF03196322
[97] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the EASE '14*. ACM, London England United Kingdom, 1–10. https://doi.org/10.1145/2601248.2601268
[98] Michael Wooldridge. 2002. Intelligent Agents: The Key Concepts. In *ACAI 2001*. Springer. https://doi.org/10.1007/3-540-45982-0_1
[99] Tom Ziemke. 2004. Embodied AI as Science: Models of Embodied Cognition, Embodied Models of Cognition, or Both? In *Embodied Artificial Intelligence*. Springer Berlin Heidelberg, 27–36. https://doi.org/10.1007/978-3-540-27833-7_2